Agent-Based Modeling of Resilience in Smallholder Agriculture: Toward Robust Models and Equitable Outcomes

by

Timothy Glynne Williams

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in the University of Michigan
2021

Doctoral Committee:

Professor Seth D. Guikema, Chair
Professor Arun Agrawal
Professor Daniel G. Brown
Professor Mark P. van Oyen

Timothy Glynne Williams

tgw@umich.edu

ORCID iD: 0000-0001-5642-728X

Acknowledgments

I am sincerely grateful to the people, places, and experiences that have shaped my thinking and supported me throughout my PhD. This dissertation would not exist without this guidance and support.

I would first like to acknowledge the privilege that I have experienced throughout my life as a White, European settler in New Zealand, and how this privilege is derived from the colonialist dispossession of Māori from their land rights and culture. It is a lifelong process to understand and learn how to work to reform these effects. I endeavor to use my research both as an instrument for personal learning and for dismantling unjust systems.

In Māori culture, it is customary to introduce oneself through the connections one shares with the world. I feel this is an appropriate place to start, as when I reflect on my time in Michigan, it is the *place* that comes to the front of mind. The Huron is my river. Thank you for providing water, enjoyment, and tranquility to Ann Arbor. Bird Hills is my mountain. Thank you for raising my heart rate and giving perspective when I run. Leslie Park is my forest. Thank you for being there throughout Summer, Fall, Winter, and Spring.

Of the people I would like to thank, first and foremost is my advisor, Dr. Seth Guikema. Thank you for supporting me to come to the United States. Your guidance over the past five years has been instrumental to my growth as a researcher. You have imparted to me your broad methodological understanding, taught me what research is and what it is not, and contributed many ideas to the chapters of this dissertation. Further, my involvement with the INFEWS project ignited my interest in food and agricultural systems, which is now central to my academic and personal identity. Thank you for your patience and support, and for the freedom you allowed me in pursuing my interests.

Thanks to my committee members. To Dr. Daniel Brown, for your modeling wisdom and detailed feedback on my paper drafts (I often find myself asking "what would Dan say?"). Thanks also for connecting me to job opportunities, both in Germany and The Netherlands. To Dr. Arun Agrawal, for your ability to get right to the essence of a problem. You have nudged my research in subtle but profound ways. And to Dr. Mark van Oyen, for being there as a source of support over the years, despite being less directly involved in the content of the research.

Thanks to the many people, networks, and communities that have supported my intellectual growth. Thanks to the Michigan University-wide Sustainability and Environment (MUSE) Initiative, for providing a home for my interests and teaching me to relate to other disciplines and ways of thinking. The MUSE community has been formative in my development as an interdisciplinary researcher, and I hope MUSE continues to thrive in years to come. Thanks to Stack Exchange for answering innumerable questions on R and Python. Thanks to Dr. Bilal Butt; your course on Conservation and Development was perhaps the most influential course I took during my PhD, and taught me to un-learn prevailing narratives, profoundly broadening my understanding of the world. Thanks to Dr. Birgit Müller, for pushing me to sharpen my hypotheses. And thanks to all of my other collaborators and coauthors—you make research fun.

Thanks to the communities that have supported me socially and spiritually. Thanks to the Guikema Research Group and IOE PhD Student community, for the lunchtime conversations, coffee trips, and generally making it a joy to come into the office. Thanks to my housemates at North Division Street, Peach Street, and Holsteinstraße, for sharing meals and making it a joy to come home from the office. Thanks to the Ann Arbor biking and food scenes, for providing local community beyond the university. Finally, thanks to the Southern Michigan Orienteering Club, for facilitating the hours I have spent wandering (sometimes aimlessly) in Michigan's forests, grasslands, and wetlands (sometimes begrudgingly). Orienteering has been a stellar and unique way to experience the Great Michigan Outdoors.

Thanks to the people and institutions that enabled me to be here. Thanks to my Dad, Jon, for giving me an analytical mind and my Mum, Ruth, for compassion and caring. These are critical parts of my personality and have shaped my interests in engineering and the environment. Thanks also to the University of Canterbury for providing an excellent educational grounding.

Thanks to my two closest friends, who patiently provided support as I rode the highs and lows of my PhD. Thanks to Tom Logan. For all outside adventures under the sun; for past, present, and future research collaborations; and for your unwavering solidarity. And thanks for introducing me to Anne Elise. Thanks to Anne Elise Stratton. For opening my eyes to ecological and political components of food systems; for adventures, gardening, and cooking; and for building a home and a life with me. And thanks to your wonderful family, who has welcomed me as one of their own.

Last, but certainly not least, I gratefully acknowledge the funding that has supported this research. The National Science Foundation supported me during the majority of my PhD through an INFEWS/T1 grant [1639214] and a CNH grant [1617364]. During a three-month period in Leipzig, Germany, I was funded by the German Federal Ministry of Education and Research [BMBF-

01LN1315A] within the Junior Research Group POLISES. During the last stage of my PhD, I was supported by the Rackham Predoctoral Fellowship and the John R Templin Scholarship.

Ngā mihi nui ki a koe. [Thank you]

TABLE OF CONTENTS

List of Figures

List of Tables

List of Appendices

Abstract

Smallholder farmers constitute one of the world's most vulnerable populations. Moreover, rising socioeconomic inequalities and biophysical degradation threaten to increase this vulnerability. There is therefore a pressing need to build resilience in smallholder agriculture. Socio-environmental systems (SES) modeling can support this goal, yet confronts two challenges that may limit its usefulness for informing agricultural development. First, as agricultural systems are highly heterogeneous and our ability to model them is imperfect, there is a risk that model-based recommendations inadvertently increase vulnerability. Second, there exist a range of approaches to agricultural development that prioritize distinct objectives (e.g., market integration versus social equity), and conflicts between these approaches could undermine progress toward more resilient futures. To build smallholder resilience therefore requires an integrated perspective on development as well as robust methodologies for comparing and integrating alternative development strategies. This dissertation uses agent-based modeling (ABM) to help address these challenges.

The first contribution of this dissertation is a set of methodological advances that improve the robustness of model-based policy analysis. These advances question two analytical norms within SES modeling. The first is a lack of attention to equity. For instance, by disregarding heterogeneity in outcomes, model-based recommendations may benefit the well-off at the expense of the vulnerable and thereby perpetuate inequity. Chapters two and three address this issue, first by establishing a conceptual framework for the equity-ABM interface and then by applying an agent-based model to examine equity in the effects of resilience-enhancing strategies. The second analytical norm that this dissertation questions is the use of a single, "best-fit" model to assess policy effects; due to our incomplete understanding of complex SES, *multiple* plausible models may exist. This common condition is known as equifinality, but it is not often considered in SES

modeling or policy analysis. To attend to this challenge, chapter four develops an approach for identifying a set of diverse model calibrations and using these to achieve a more robust policy analysis. Together, these methodological advances facilitate more robust and equitable policy assessments, in agricultural systems and beyond.

The second principal contribution of this dissertation is substantive. Emerging from the modeling of smallholder resilience, I find complementarity between disparate agricultural development approaches. For instance, chapter five compares the effects of legume cover cropping (a form of ecological farm management) and microinsurance (a financial institutional support) on smallholder climate resilience. Although these approaches are traditionally promoted by distinct academic communities and development organizations, the results show that, when implemented together, they are highly complementary. Next, chapter six investigates the potential for contract farming to overcome the negative effects of large-scale land acquisitions on smallholder food security. Results suggest that preserving smallholder autonomy through contract farming can simultaneously improve smallholder food security and agricultural production, thereby better aligning the preferences of developers and smallholders. Thus, these chapters together suggest the benefits of reconciling perspectives on and approaches to agricultural development.

As a whole, this dissertation advances the application of agent-based modeling and resilience thinking in smallholder agriculture. Beyond agricultural applications, it lays the groundwork for identifying robust and equitable development strategies in SES.

Chapter 1

Introduction

1.1 Fostering resilience in smallholder agriculture

Approximately two billion people worldwide rely on agriculture to support their livelihoods (Hazell et al., 2010b). The vast majority of these are smallholder, subsistence farmers living in low- and middle-income countries. Unfortunately, these people constitute some of the world's most vulnerable populations, whose livelihoods are threatened by a wide (and increasing) array of challenges. These challenges range from direct stresses to the production system, such as soil degradation and drought, to large-scale drivers affecting the broader socio-political context in which agriculture operates, such as land access and globalization.

In the face of these interacting stresses, agricultural development will be central to global poverty alleviation going forward. It is widely acknowledged that investments into agriculture are some of the most effective toward this end (Webb and Block, 2012; Pray et al., 2017). There exist, however, diverging perspectives on how agriculture should develop so as to simultaneously support production, livelihoods, and the environment (Wiggins et al., 2010; Gaffney et al., 2019; Pretty et al., 2018; Bommarco et al., 2013). Under some future visions, smallholder agriculture becomes obsolete and is replaced by large-scale, mechanized farming systems with high levels of external inputs. This most closely describes the development trajectories seen in the Global North. Alternative perspectives take a more smallholder-centric view, with a range of opinions regarding the relative importance of external inputs and internal ecological processes for crop nutrient provisioning (Bernard and Lux, 2017).

Although the specifics of smallholder agriculture's future remain uncertain, "resilience thinking" provides an encompassing and normative analytical lens through which to evaluate smallholder vulnerability and agricultural development. In this context, resilience describes the capacity of a smallholder agricultural system to persist in the face of sudden and incremental change, to adapt to sustain development within current pathways, and to transform into other emergent pathways (Folke, 2016). Agricultural development strategies can act at different scales and through different components of the smallholder system (e.g., by building soil fertility versus providing non-farm employment opportunities) to foster resilience across these dimensions (Hansen et al., 2019).

1.2 Agent-based modeling of smallholder resilience

It is difficult to evaluate how development strategies might affect smallholder resilience. First, multidimensional outcomes and the unprecedented scales of intervention push the limits of historical data, rendering empirical methods inappropriate (Egli et al., 2018). Second, agricultural systems are complex (Liu et al., 2007); they evolve through the dynamic interplay of heterogeneous human decision-makers (e.g., farmers, policymakers, consumers) and non-human processes (e.g., water cycles, soil fertility, pollinators) (Peterson et al., 2018). This socio-environmental intertwinedness generates path dependencies that affect future system states (Haider et al., 2018). Together, these characteristics can lead to non-linear responses that are difficult to predict without a nuanced understanding of underlying processes and proximate drivers of change.

Agent-based modeling is a simulation-based approach for modeling complex systems. Agent-based models (ABMs) simulate an interacting, heterogeneous population of autonomous actors. They can be applied to investigate how both top-down interventions and bottom-up processes generate emergent, system-level outcomes, as well as how effects are distributed throughout a population. Due to these features, ABMs have been extensively applied to assess the effects of policies and interventions in agricultural systems (Kremmydas et al., 2018) and can be useful for characterizing system resilience (Egli et al., 2018).

Moreover, ABMs exhibit a relatively underexplored potential to act as boundary objects across normative visions for agricultural development. Different research communities advocate for distinct development paradigms, for instance technology-oriented versus ecologically-oriented. Due to their different epistemological foundations, these paradigms can be difficult to compare under a common framework. ABMs can be used to examine synergies and tradeoffs between outcomes (e.g., productivity, profit, soil degradation), both over time and between different types of actors. Thus, ABMs facilitate the comparison of alternative development perspectives.

1.3 The need for robust and equitable resilience assessments

Given the current levels of poverty and alarming future projections for socio-environmental change, timely action to improve resilience in smallholder systems is necessary. However, due to the complexity and uncertainty inherent to agricultural systems, models may mis-estimate the true effects of proposed resilience-enhancing strategies. In extreme cases, model-based recommendations therefore have the potential to inadvertently increase vulnerability, i.e., to be inequitable or "maladaptive" (Barnett and O'Neill, 2010). This could be particularly problematic if interventions are difficult or costly to counteract or affect other system components in unexpected ways (Leclère et al., 2014).

There are several analytical norms within quantitative resilience analysis that could give rise to maladaptive recommendations. First, most model-based assessments examine the *average* effect of an intervention across the population. It is possible, however, that increasing resilience for one population group adversely impacts those most at risk (Barnett and O'Neill, 2010; Miller et al., 2010). Thus, there is a need to integrate distributional equity more thoroughly into resilience analysis. Second, given that many environmental processes operate at different temporal scales to human decision-making (Rodríguez et al., 2006), the time horizon used for policy assessment can exert a strong influence on model-based recommendations. Yet, most modeling studies do not consider the impacts of the time horizon. Finally, due to the complexity of agricultural systems and the rising "complicatedness" of ABMs (Sun et al., 2016), it is possible that there exist *multiple* plausible system representations. This condition is known as equifinality (Oreskes et al., 1994). Yet, most modeling studies utilize a single model configuration. Maladaptation could arise if multiple plausible descriptions exist and these lead to qualitatively different outcomes. In all these cases, modeling advances are needed to address these challenges and facilitate more robust resilience assessment.

1.4 Dissertation outline

In this dissertation, I employ a complex systems perspective to address the overarching question of how to improve resilience in smallholder agricultural systems. The dissertation chapters, which each act as stand-alone research articles, provide substantive insights into the relative merits of strategies for enhancing smallholder resilience. In seeking to answer this overarching question, however, shortcomings within predominant methodological practices became apparent. Correspondingly, the dissertation also demonstrates methodological contributions to ABM development and analysis, which are relevant beyond agriculture. The main body of the dissertation is split into the following five chapters.

1.4.1 On equity in agent-based modeling

Advancing equity is a cross-cutting challenge to society, science, and policy. ABMs are increasingly applied as scientific tools to advance system understanding, inform decision-making, and share knowledge. Yet, equity has not been thoroughly integrated or discussed within the agent-based modeling community. This hinders the potential for agent-based modeling to be used in future research aimed at improving equity.

In chapter 2, I develop a framework that elucidates the links between agent-based modeling and equity. The framework positions the modeler as a filter and a lens through which knowledge is projected into and out of the model. To operationalize the framework, I outline three "action pathways" for advancing equity through ABM: (1) engage stakeholders and society, (2) recognize modeler positionality and bias, and (3) assess equity with agent-based models. The framework and examples can be used as guidance in future modeling efforts, so that agent-based modeling can play a larger role in creating a more equitable future.

1.4.2 Resilience and equity

Strategies aiming to increase climate resilience in smallholder agricultural systems may not equally benefit all groups of the smallholder population. To reduce the potential for aggravating existing vulnerabilities, resilience analyses need to acknowledge the possibility for inequities in the effects of resilience-enhancing strategies (RESs). However, distributional effects are seldom considered in quantitative resilience analysis.

In chapter 3, I develop, validate, and apply a household-level ABM to explore the equity of climate RESs in an Ethiopian smallholder farming system. The strategies include the provision of seasonal climate forecasts, which allow households to make better-informed management decisions, and an increase in the availability of non-farm wage labor, which can increase income and purchasing power. Given the different mechanisms through which these two strategies act, heterogeneous households may respond differently (Kansiime et al., 2018), leading to asymmetries in resilience between groups or even reinforcing poverty (Miller et al., 2010; Béné et al., 2012), amounting to maladaptation.

Results reveal that different measures of resilience lead to divergent assessments of equity in policy effects. In particular, in the wake of a drought, both RESs benefit the moderately vulnerable households at the expense of the more vulnerable households—–i.e., they are inequitable. These results demonstrate that a pure focus on poverty reduction may be insufficient to promote equitable development. Given the prevalence of climate shocks in smallholder systems, future studies of resilience should therefore jointly consider both poverty reduction and shock recovery, as well as the potential for inequity in the effects of RESs.

1.4.3 Assessing model equifinality

Equifinality describes a situation where there exist multiple plausible explanations for a single outcome (Axtell and Epstein, 1994; Oreskes et al., 1994; Beven, 2006). Equifinality can arise when our understanding of underlying processes is incomplete and there are few data against which to validate hypothesized process-based descriptions. Although this condition is prevalent in socio-environmental systems, model calibration most frequently seeks to identify a single, "best-fit" model, thus not allowing for equifinality. If other feasible calibrations lead to qualitatively different model behavior, prioritizing policy from a single, best-fit calibration could amount to maladaptation.

In chapter 4, I develop and demonstrate an approach for ABM calibration and analysis that (1) identifies multiple model calibrations and (2) assesses the implications for policy analysis. The optimization-based approach maximizes diversity in the model parameters and/or structural configurations to efficiently represent any equifinality in the model set. I apply the approach to the ABM developed in Chapter 3, in order to explore the robustness of the resilience assessment to model equifinality.

Case study results demonstrate consistent policy effects over the set of diverse model calibrations, enabling stronger conclusions than a single model analysis. More generally, the approach facilitates more nuanced policy assessments in socio-environmental systems, because it identifies the conditions under which policies may or may not be robust.

1.4.4 Ecological and financial strategies

There exist different and, at times, conflicting perspectives on how to best support climate resilience in smallholder agricultural systems. Institutional interventions such as microinsurance schemes have recently gained traction as tools for agricultural development and poverty reduction (Hazell et al., 2010a; SwissRe, 2013; Kramer et al., 2019). Simultaneously, there is an increasing drive for ecological intensification to sustain or enhance both livelihoods and natural resources (Bommarco et al., 2013; FAO, 2018; HLPE, 2019). Such approaches are traditionally advocated for by different communities, with often strong ideological disagreements. However, given the different mechanisms through which these two approaches act, they may in fact be *complementary* when considered together.

In chapter 5, I examine the potential complementarities between selected ecological and financial approaches for supporting smallholder climate resilience. To do this, I develop a social-ecological simulation model of mixed crop-livestock smallholder farming. I apply the model to examine how different combinations of legume cover cropping (ecological) and index-based crop insurance (financial) affect smallholder farmers' income over time and in the wake of droughts.

The results underscore the complementary roles that ecological and financial strategies could play in building smallholder resilience. Specifically, I find that microinsurance always provides larger benefits during and in the wake of a drought, while cover cropping progressively reduces poverty in the medium- to long-term. The stylized model constitutes an important social-ecological foundation for future empirical research to inform agricultural innovation and development priorities.

1.4.5 Large-scale land acquisitions

Over the past decade, the Global South has experienced a rapid increase in large-scale investment in agricultural land (Deininger and Byerlee, 2011). Unfortunately, these large-scale land acquisitions (LSLAs) frequently generate tradeoffs between agricultural production and smallholder livelihoods (Müller et al., 2021). LSLAs remain a prevalent global phenomenon and thereby a risk to future smallholder livelihoods. Yet, literature to date has focused almost exclusively on the causes and effects of LSLAs over the past 12 years. Some work has sought to identify beneficial institutional arrangements (Oberlack et al., 2016; Arndt et al., 2010; Baumgartner et al., 2015), but there have been no process-driven assessments of how to facilitate positive outcomes for both smallholders and large-scale investors. Agent-based modeling is well-poised to explore such dynamics.

For chapter 6, I examine the potential effects of contract farming (CF), an arrangement compatible with LSLAs that preserves some smallholder land rights, on smallholder food security and regional productivity. To do this, I develop an ABM of smallholder livelihoods, calibrate it using household survey data collected in four LSLA-affected areas within Ethiopia, and apply it to examine the distributional effects of various LSLA and CF arrangements. This analysis integrates some of the conceptual advances from the previous chapters (namely, equifinality and equity) and applies them to investigate a pressing threat to smallholder resilience.

Results show that contract farming can simultaneously increase commodity agricultural production and support livelihoods in mixed crop-livestock smallholder systems. Importantly, arrangements that preserve smallholders' autonomy over their land led to the strongest and most synergistic outcomes, positioning contract farming as a promising alternative to forms of intensification by dispossession enacted by LSLAs.

Chapter 2

On Equity in Agent-Based Modeling

Advancing equity is a cross-cutting challenge to society, science, and policy. Agent-based models are increasingly applied as scientific tools to advance system understanding, inform decision-making, and share knowledge. Yet, equity has not received due attention within the agent-based modeling (ABM) community. Integrating equity into ABM will both reduce the risk of ABM research inadvertently perpetuating existing inequities and harness the opportunity for ABM to be used in future research aimed at improving equity. In this chapter, we present a conceptual framework that elucidates the links between ABM and equity. The framework positions the modeler as a filter and a lens through which knowledge is projected into and out of the model. To operationalize the framework, we outline three "action pathways" for advancing equity through ABM: (1) engage stakeholders and society, (2) recognize modeler positionality and bias, and (3) assess equity with agent-based models. Within each action pathway, we provide concrete guidance, examples, and reflection questions for modelers. We hope that our framework and examples can guide future modeling efforts, so that ABM can play a larger role in creating a more equitable future.

2.1 Introduction

The world is not fair. Resources and political power are chiefly held in the hands of a select few, while vulnerable populations are left to cope with social and environmental burdens. This unfairness transcends domains, scales, and contexts. For example, climate change, a global challenge largely attributable to rich nations of the Global North, disproportionately affects poor people in the Global South (Tol et al., 2004). Moreover, these populations often have less agency and political leverage to adapt to climate change effects. Climate change, along with many other cross-cutting global challenges, is therefore closely intertwined with notions of justice, power, and distribution. Correspondingly, promoting equity to reduce inequality is central to current scientific agendas (Clark and Harley, 2020), policy priorities,[1] and development goals (UN General Assembly, 2015).

Equity, a term often used in conjunction with the concepts of justice and fairness, provides the conceptual lens for this chapter. We define equity as a moral and political ideal that aspires toward: (1) fair distributions of goods within society, (2) fair inclusion within decision-making procedures, and (3) recognition of diverse socio-cultural identities with unique needs and historical contexts (Pereira et al., 2017; Tyler, 2000; Fraser, 1995; McDermott et al., 2013). For the purposes of this chapter, we use the terms 'equity' and 'justice' interchangeably.[2] 'Fairness' adds a normative and plural angle to the dimensions of equity, meaning that there are debates around what is considered fair by different people and in different contexts (Jacobs and Wallach, 2021; Fraser, 2009).

Quantitative models have an important role to play in achieving a more equitable future. Models are both useful tools for ex-ante exploration of equity-oriented interventions and boundary objects for stimulating societal debate and co-producing knowledge. Integrating equity into modeling allows for reducing the risk of models inadvertently discriminating against marginalized populations as well as helps to identify strategies for guarding against future challenges to equity. In order to leverage these opportunities and reduce these risks, concepts of equity are being increasingly incorporated into modeling paradigms and practices. For example, there are nascent efforts to develop and mainstream notions of 'algorithmic fairness' in machine learning models, particularly in order to reduce potential for discriminatory model-based recommendations due to biases in historical data and model design (Corbett-Davies et al., 2017; Feldman et al., 2015). In parallel, there exist a growing number of conceptual frameworks linking equity theory to substantive application domains, such as transportation (Pereira et al., 2017), ecosystem services (McDermott et al., 2013), health (Rajkomar et al., 2018; Chandanabhumma and Narasimhan, 2020), urban planning (Chu and Cannon, 2021), and sustainability (Leach et al., 2018).

Agent-based modeling (ABM)[3] is applied across a comparably diverse range of topics and particularly warrants integration with equity. Agent-based models simulate an interacting, heterogeneous population of autonomous actors and can be applied to investigate how both top-down interventions and bottom-up processes generate emergent, system-level outcomes. Given these features, agent-based models are particularly suited to represent problems of (in)equity, relative to

[1] An example is United States President Biden's executive order to spend 40% of sustainability investments in disadvantaged communities and create an Office of Health and Climate Equity in the Department of Health and Human Services.

[2] We note that some scholars take the position that justice is more encompassing than equity, both in theory (Sikor et al., 2014) and in practice (Ikeme, 2003), as it can include a wider range of ethical positions as well as notions of rights and responsibility. However, in this chapter we use the word equity to describe the encompassing concept.

[3] In this chapter, we use the ABM acronym to refer to "agent-based modeling," which we conceptualize as encompassing both the agent-based model itself (the noun: 'model') and the process through which this model is developed, tested, and applied (the verb: '-ing').

other modeling approaches that work at larger scales or do not provide such a mechanistic description. Indeed, such challenges inspired some of the seminal advances within the ABM field, including Schelling's segregation model (1971) and Epstein and Axtell's Sugarscape (1996), which sought to explain the emergence of macro-level inequalities from micro-level processes. Since then, ABM has been applied in upwards of 7500 publications in a broad range of academic disciplines (Janssen et al., 2020). However, there remains a lack of integration of equity into ABM applications.[4] Given the risk of ABM inadvertently discriminating against vulnerable populations by not considering equity, as well as the time-pressing imperative to advance equity and the untapped potential for ABM to play a role in doing so, such an integration is sorely needed.

A critical precondition is a solid conceptual and practical understanding of the equity-ABM interface. In this chapter, we make two contributions to this end. To begin, in sections 2.2 and 2.3 we work toward a conceptual framework that positions equity within ABM. The framework aims to provide a generalized conceptual understanding of this intersection. Then, in section 2.4, we outline a set of practical strategies for better integrating equity within ABM research. We intend our framework and strategies to be of interest to (agent-based) modelers and non-modelers alike. For modelers, we illustrate a concrete set of approaches to draw from throughout the stages of agent-based model development. For non-modelers interested in equity, we demonstrate how ABM can uniquely contribute to questions on this theme.

2.2 Background

2.2.1 Equity

The concept of equity can be conceptualized as a triplet (Sikor et al., 2014):

$$\text{Equity} = (\text{dimensions, objects and subjects, fairness criteria}) \tag{2.1}$$

Any complete description of equity requires these three elements. Equity is frequently understood to have three principal material 'dimensions' relating to inputs, process, and outputs. First, 'recognitional equity' describes the perspectives and identities acknowledged and valued by society, particularly seeking to eliminate forms of cultural domination and discrimination (Fraser, 1995; Sikor et al., 2014). Recognitional aspects emphasize equity as a situated phenomenon that can only be understood within the culture, beliefs, practices, and institutions that guide actors (McDermott et al., 2013). Second, 'procedural equity' describes the fairness in people's inclusion in

[4] A Scopus search for the terms (TITLE-ABS-KEY ("agent-based" OR "agent based") AND TITLE-ABS-KEY (model*)) on 31 March 2021 yielded over 30,000 results. Adding the search term (*equit* OR disparit* OR *justice* OR *fairness*) reduced the results to 344 (i.e., just over 1% of publications).

and ability to influence decision-making (Arnstein, 1969; Tyler, 2000). This dimension is closely related to notions of power and agency. Third, 'distributional equity' describes the fairness in the allocation of goods in society (Rawls, 2009). Although conceptually distinct, these three dimensions are co-created and dynamically evolve as, for example, power differentials are both a cause for and a consequence of distributional inequity (Clark and Harley, 2020; Leach et al., 2018). For the purposes of this chapter, we assess equity across these three dimensions and consider inequity to exist when there is unfairness in at least one of these dimensions. Further, we consider a process (e.g., ABM) as 'equitable' or 'equity-oriented' when it improves at least one of these dimensions.

Within the material dimensions of equity, there exist the cross-cutting questions of "equity of what" and "equity between whom," i.e., the objects and subjects of equity. Such considerations are inevitably context specific, but also invariably introduce value judgements. For example, Rawls (2009) describes "equity of what" as relating to the 'primary goods' in society (e.g., health, civil rights, income, social respect), yet this has been subsequently broadened to include human 'capabilities,' which encompass the freedoms and opportunities available for people to choose and to act (Nussbaum and Sen, 1993). Capabilities aim to recognize that people have different values and capacities to utilize resources as a means to a good life. With respect to "equity between whom," people can be grouped or classified based on different dimensions of socio-cultural identity (e.g., class, occupation, gender, ethnicity, geography, sexual identity (Leach et al., 2018)). Here, it is critical to reflect on how the subjects of equity are decided (Fraser, 2009), due to risks of "othering" people or cultural appropriation (Fraser, 1995; Smith, 2013).

Finally, underlying any assessment of equity is some notion of 'fairness'. The principles used to define fairness add a moral and normative angle to equity, therefore requiring critical reflection on, and transparency around, what is considered fair (Fraser, 2009). Fairness criteria can be broadly distinguished as pertaining to either consequences or rules (McDermott et al., 2013). The most basic (and probably oldest) consequence-based notion of fairness is that of egalitarianism (Konow, 2003), i.e., desiring an equal distribution of outcomes between people or groups. Equity as equality is frequently operationalized through quantitative measures like the Gini coefficient, which describes the deviation of a distribution from perfect equality. Rules-based principles, in contrast, judge the fairness of distributional outcomes by the rules through which they arise (McDermott et al., 2013). Prevailing notions of equity, for instance, adopt a needs-based principle that seeks to favor the least advantaged in society (Rawls, 2001), so that people who need more of something have access to more of it (Leach et al., 2018). Such perspectives acknowledge that people have variable levels of need (Konow, 2003), due to inherent disadvantages suffered by different groups. Other rules-based principles rest on alternative moral notions, such as 'merit,' in which rewards should be proportional to an individual's inputs (e.g., effort) (McDermott et al., 2013). There are a range of perspectives on the conceptual boundaries of merit, for instance whether characteristics such as talent, intelligence, and educational opportunities should (or should not) affect rewards (Konow, 2003). In any case, by articulating a fairness principle, equity seeks to make explicit the values held by society. Equity is therefore unavoidably normative, as the fairness principle makes a judgement about what *should* be. Because these values are culturally mediated and can evolve over time, equity is also both plural and political (Leach et al., 2018).

2.2.2 The equity-modeling interface

Adopting equity as an analytical lens is unquestionably a difficult endeavor. It must be done with much thought and critical reflection. Yet, integrating equity into analytical research is also increasingly necessary; as modeling becomes more deeply integrated within societal decision-making, so too does its potential to do harm. For example, in a highly publicized incident, a machine learning algorithm designed to predict recidivism (a convicted criminal committing a future crime) was biased against Black people. The algorithm, which was trained on historical data and used to inform decision-making in the United States criminal justice system, was significantly more likely to falsely flag Black defendants as future criminals (Angwin et al., 2016). In so doing, it inadvertently perpetuated historical racial disparities contained within the data. Other prominent model-based discriminatory failures have arisen in healthcare, facial recognition, and advertising (Mehrabi et al., 2019).

Accordingly, the 'equity-modeling interface' is receiving increasing scrutiny within science. These efforts pertain to two principal levels: within a model itself and within the broader process of model development and application. Within models, inequity can exist in model inputs, processes, and outputs. On the inputs side, data can be biased; historical data contain historical inequities (Jacobs and Wallach, 2021) and it is possible that insufficient data exist for a socio-cultural group (i.e., they are insufficiently recognized), either due to the size of the group or inadequate data collection procedures (Rajkomar et al., 2018). Further, the variables or metrics selected as model inputs can inadvertently lead to discriminatory model solutions. For example, an algorithm trained to optimize health care *costs* as a proxy for health led to systemic bias against Black patients (Obermeyer et al., 2019). This was because Black patients, due to barriers to accessing health care, received less weight in the cost variable (i.e., at a given level of health, Blacks generate lower costs than Whites) and thereby were attributed lower risk scores. Therefore, in line with the saying "garbage in, garbage out," we could also say "inequity in, inequity out."

Within model processes, there is a fast-growing body of work in the machine learning community on algorithmic fairness, which recognizes that model formulations can inherently be prejudiced towards or against particular individuals or groups (O'Neil, 2016; Mehrabi et al., 2019). This work focuses on formalizing measures of fairness (Feldman et al., 2015), as well as integrating fairness objectives and constraints into optimization formulations (Corbett-Davies et al., 2017). In these models, there is often found to be a 'cost' of fairness, in that the fairest solution is not necessarily also optimal with respect to other decision objectives (e.g., economic cost) (Kleinberg et al., 2016).

On the outputs side, equity is most commonly measured using inequality metrics. The most well-known is the Gini index, which measures the distribution of a resource (e.g., income) across a population. Other metrics such as the Atkinson and the Kolm-Pollak indices explicitly attach a (normative) penalty to inequality (Logan et al., 2021) and therefore more fully engage with equity as defined in this chapter. Beyond such metrics, notions of distributional equity are frequently incorporated in research on disparities in health and healthcare access (Jatoi et al., 2003; Speybroeck et al., 2013) by stratifying participants based on socio-cultural identity (e.g., race) and comparing outcomes between groups.

At the second level (i.e., the broader process of model development and application), there are efforts to reduce inequity in model bias and to mitigate power asymmetries between modelers and non-modelers. Models are not neutral objects (Voinov et al., 2014; Cruz Cortés and Ghosh, 2020); many implicit value judgements are made when developing a model, and bias is introduced when mapping theoretical constructs (e.g., socioeconomic status) to observable characteristics (e.g., income) (Jacobs and Wallach, 2021). In order to reduce inequity introduced by these biases, modelers first need to recognize if there is a problem. By engaging in a reflexive process that questions assumptions and practices, modelers can understand and begin to deconstruct their biases (Grieshaber, 2010; Steger et al., 2021b). Moving beyond reflection, equity can be improved by involving stakeholders and societal groups within the model development process. When done well, this helps to include a wider range of epistemologies to reduce the effect of the modeler's individual bias (Voinov and Bousquet, 2010), as well as to reduce power asymmetries between modelers and non-modelers. There are a variety of levels of possible stakeholder engagement that respectively achieve different levels of equity (Voinov et al., 2016). Notably, if not appropriately managed, participatory modeling can perpetuate recognitional and procedural inequity by tokenizing or appropriating local or indigenous knowledge (Arnstein, 1969).

2.2.3 The equity-ABM interface

Agent-based models share many characteristics with other modeling approaches: they often aim to and are increasingly applied to inform decision-making; they take inputs, transform them, and produce outputs; and they require a large number of explicit and implicit design decisions. Thus, many of the equity considerations developed in other fields are also relevant to ABM. Yet, agent-based models exhibit unique characteristics that give rise to distinct opportunities for and risks to equity. This warrants an analysis of the 'equity-ABM interface'. As far as the authors are aware, this is the first attempt at such an analysis.

Several features of agent-based models position them uniquely to engage with equity. First, a key characteristic of agent-based models is their ability to represent heterogeneity. Heterogeneity can exist in agent characteristics (e.g., capacity or gender), experiences (e.g., success in finding employment), and outcomes (e.g., wealth), thereby facilitating assessments across all three dimensions of equity. Moreover, agent-based models can represent multiple dimensions of agent heterogeneity that can be both static (e.g., gender) and dynamic (e.g., social networks) throughout a simulation. Equity-oriented agent-based models therefore have the potential to consider nuanced socio-cultural identities (i.e., equity between whom). Second, relative to most other modeling approaches, ABM allows for very flexible behavioral representations (utilitarian, prospect theory, theory of planned behavior, heuristics, etc.) (Groeneveld et al., 2017). Agent-based models are therefore useful tools for formalizing diverse decision-making processes to facilitate assessments of procedural equity. Such procedural inequities can exist at a single scale (e.g., heterogeneous levels of agent capacity) or across scales (e.g., heterogeneous individual-level influence over system-level decisions). Third, due to their dynamic, process-based nature, agent-based models can represent complex feedbacks and interconnections between mechanisms and outcomes. They therefore facilitate dynamic, integrated equity assessments (e.g., how distributional equity acts as both an input to and an output of procedural equity (Leach et al., 2018)).

Yet, these very features also pose risks if agent-based models are not designed with equity considerations at front of mind. First and foremost, ABM is a highly flexible modeling approach, in that it allows for a near unlimited number of degrees of freedom (or 'complicatedness' in model structure (Sun et al., 2016)). Beyond the issues of parameter identifiability, developing a complicated model requires a huge number of structural decisions, each of which contains some level of value judgement on the part of the modeler. An agent-based model is therefore a projection of the modeler's worldview, which is conditioned by their particular experiences and socio-cultural identity. Agent-based models thus may exclude or misinterpret non-dominant worldviews or theories, thereby perpetuating recognitional inequity. Second, attempts to incorporate agent heterogeneity may miss important dimensions of difference. For example, it may not be sufficient to consider heterogeneity only in *who* the agents are, but also in *how* they behave and are treated by system-level rules. In extreme cases, this risks ascribing distributional inequity to innate characteristics rather than procedural inequities (e.g., ascribing disparities between Black and White citizens' income to (genetic) differences in agent characteristics), which can be deeply problematic (VanderWeele and Robinson, 2014). Thus, as is true for all modeling, improving equity within ABM must be a balancing act between simplicity and interpretability on one hand and structural realism on the other.

Current ABM research and practices engage with some components of equity more than others. Equity dimensions are often incorporated into the purpose of agent-based models, most frequently through questions of distributional equity (i.e., disparities in outcomes across agents or agent groups). Notions of procedural equity and power dynamics are explored in some cases, albeit less frequently (Lindkvist et al., 2020). Beyond the model itself, ABM development often engages in participatory methods, thereby achieving a more equitable distribution of power between modelers and stakeholders (Biommel et al., 2016; Steger et al., 2021a) as well as facilitating integration into policy- and decision-making (Will et al., 2021). Moreover, standardized ABM frameworks such as the ODD protocol[5] and TRACE documentation[6] improve equity by encouraging modelers to be explicit about their assumptions and reasons for inclusion or exclusion of particular actors and processes. Yet, rarely do modelers explicitly consider their positionality (i.e., their socio-cultural identities and how these relate to the modeled context (Holmes, 2020)) or more meta-level reflections on the moral principles implicit within model purpose and design. Looking forward, to facilitate wider uptake of equity considerations within ABM research, there is a need to clearly define the 'equity-ABM interface' as well as to identify actionable strategies for equitable ABM. These are the aims of the following sections.

2.3 Conceptual framework

We synthesize theories of equity with prevalent notions of ABM to establish a framework for the equity-ABM interface (Figure 2.1). It is grounded on the following overarching premises:

1. Agent-based models are nested within a broader scientific and socio-political context (dashed black lines, Figure 2.1).

2. The modeler,[7] in engagement with stakeholders, acts as a filter and a lens for translating knowledge to and from the model (solid black lines, Figure 2.1).

3. The three elements of equity (dimensions, subjects and objects, fairness criteria) cut across all components and actors within ABM (orange boxes, Figure 2.1).

[5]The ODD (Overview, Design concepts, and Details) protocol is a standardized method for describing agent-based models (Grimm et al., 2006, 2020). The ODD+D builds on this by adding more specificity around human decision-making (Müller et al., 2013).

[6]TRACE stands for "TRAnsparent and Comprehensive model Evaludation," where 'evaludation' describes an integrated form of model evaluation and validation (Augusiak et al., 2014). The TRACE framework is a tool for planning, performing, and documenting model evaluation and validation (Schmolke et al., 2010; Grimm et al., 2014).

[7]We refer to a single 'modeler' in our framework and action pathways, but intend it as also appropriate for collaborative modeling projects with multiple modelers. We discuss the limitations and extensions for multiple modelers in the Discussion section.

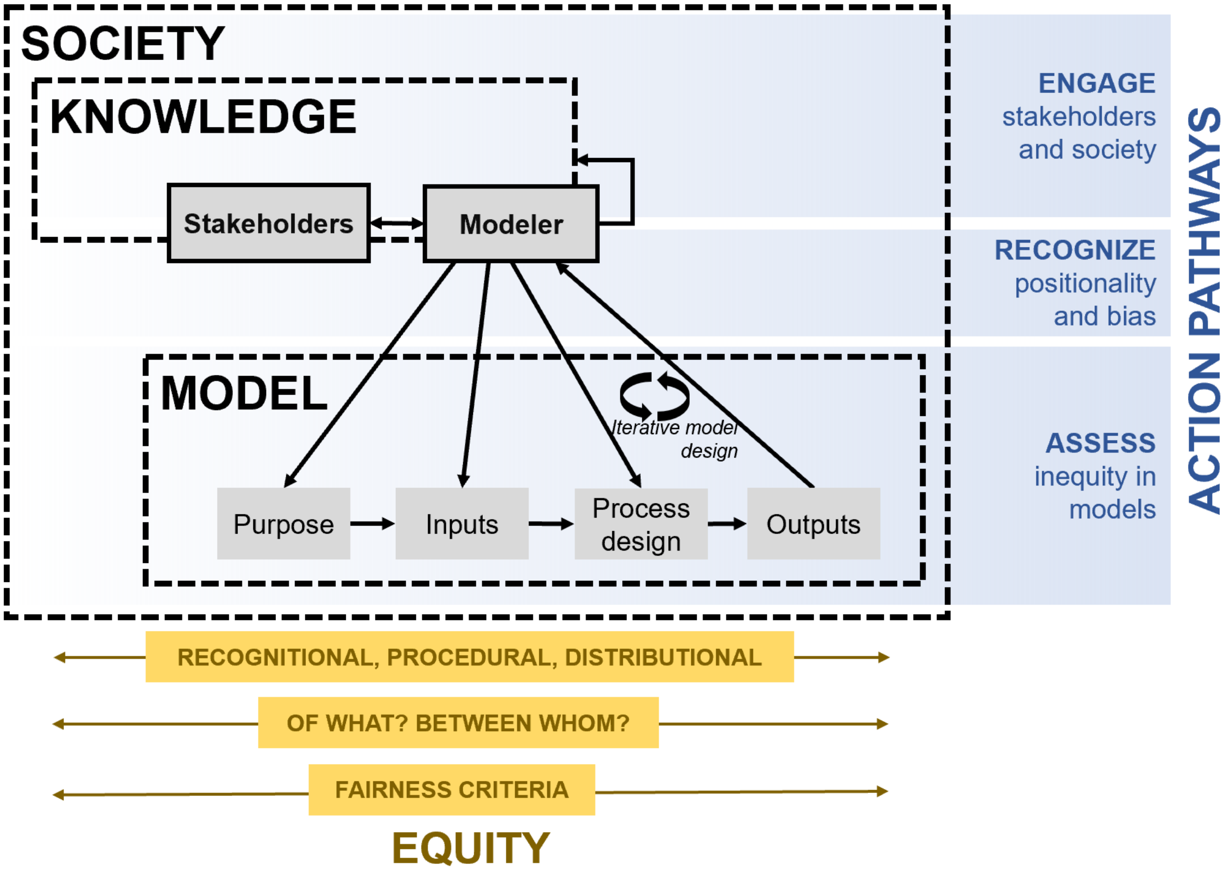

Figure 2.1: Conceptual framework for the equity-ABM interface. ABM is a societally nested process (dashed black), with the modeler as a filter and lens for translating knowledge to and from the model (solid black). Equity is a cross-cutting issue (orange) and ABM can engage with it at three principal levels (blue).

4. Advancing equity through ABM will involve (i) engaging stakeholders and society in knowledge co-production, (ii) recognizing modeler positionality and bias, and (iii) applying agent-based models to assess equity (blue boxes, Figure 2.1).

The first premise underscores that agent-based models are not isolated, technical objects. All agent-based models begin with a purpose (Grimm and Railsback, 2005), which generally should relate to a topic of interest to science and society. Predominant societal narratives can influence scientific priorities and funding, thereby motivating select model purposes. Thus, even before a model has been designed or built, ABM cannot be separated from the broader scientific and socio-political environment in which it operates.

The second premise positions the modeler as a primary locus for action on equity. The modeler is an integrator who, in an iterative fashion, formulates the model's purpose, translates it into a model-based format, filters and interprets the model outputs, and communicates these to society (Schmolke et al., 2010; Grimm and Railsback, 2005). Bias can enter through any of these processes (i.e., the black arrows linked to the modeler in Figure 2.1). Using a metaphor of the modeler

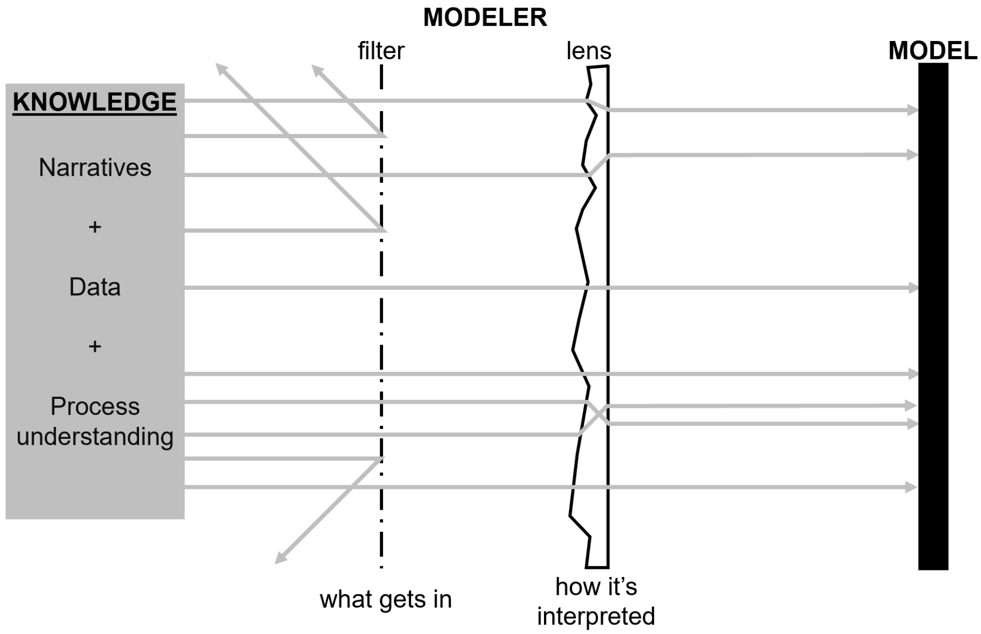

Figure 2.2: The modeler as a filter and a lens. The modeler's positionality affects both *what* knowledge and *how* knowledge is projected into the model. Engaging stakeholders can affect the filter (e.g., through broadening or narrowing the modeling scope) and the lens (e.g., through offering alternative understandings of a process).

as a filter and a lens, we conceptualize this bias in two ways: through what the modeler sees and allows to inform the model, i.e., the 'filter' that admits or blocks knowledge; and how the modeler understands the world, i.e., the 'lens' that can morph knowledge (Figure 2.2). Both of these are shaped by the modeler's positionality. Engagement with stakeholders (and their own filters and lenses) can reshape the modeler's effective filter and lens, but never eliminate them, as the modeler is generally responsible for translating stakeholder knowledge to model code and vice versa, even in highly participatory research (Voinov and Bousquet, 2010). Finally, the adjacency of the modeler and stakeholders underscores that the modeler is themselves a stakeholder in the modeling process (Voinov et al., 2014).

Through the third premise, we emphasize the cross-cutting nature of equity within ABM. Each equity element does not exclusively exist within specific components or actors, but is relevant at multiple levels. For example, procedural inequity, which describes fairness within decision-making processes, can be assessed within agent-based models by incorporating power dynamics. Yet, procedural inequities can also exist between the modeler and the stakeholders or between the stakeholders themselves (i.e., some stakeholders exert more influence over the modeling process). Similarly, the subjects of equity (i.e., between whom?) can pertain to the modeled agent identities as well as the stakeholders included within model development.

The final premise introduces three action pathways——engage stakeholders, recognize positionality, and assess inequity——for ameliorating inequities and leveraging opportunities to advance equity in ABM. First, at the highest level, modeler-stakeholder engagement can help to reduce modeler biases or misunderstandings of system contexts that could perpetuate inequity (e.g., incorrect characterization of cultural groups) as well as distribute power more equitably between parties. In extreme cases, stakeholder engagement can leverage the modeling process to affect structure and power dynamics within the stakeholder groups themselves. A further consideration within this top level is society's ability to derive benefit from the ABM process: how is the science communicated, and to whom? Second, the middle level is about how the modeler engages in self-reflexivity to deconstruct and acknowledge personal values, particularly how they may create bias throughout agent-based model development (Grieshaber, 2010). Modeler reflection also aims to make any value judgements (e.g., what is 'fair') transparent (Jacobs and Wallach, 2021). Finally, the bottom level is about how the agent-based model itself engages with equity; agent-based models can be used to conceptualize inequities, better understand their implications, and investigate mechanisms toward overcoming them (Campbell et al., 2015). The ultimate goal here is to use equity-oriented agent-based assessments to inform societal decision-making.

2.4 Applying the framework: Three action pathways

The action pathways (blue boxes in Figure 2.1) illustrate the options that are available for modelers to integrate equity within ABM research and therefore aim to make the framework operational. In this section, we discuss each of these three pathways in more detail and provide examples and concrete guidance to assist modelers in implementing them.

There are both challenges and opportunities associated with applying these pathways, which we discuss more deeply in the Discussion section. However, we flag two key ideas upfront. First, although the framework presents the action pathways as conceptually distinct, in practice they are neither mutually exclusive nor independent. For example, an equity-oriented model assessment may require stakeholder input to inform model process design and additionally requires the modeler to critically reflect on their quantification of fairness.

Second, the action pathways are not prescriptive, and all may not always be appropriate. For example, since we consider equity only in relation to human experiences, agent-based models of non-human entities (e.g., animals or plants) cannot in themselves be used to assess equity. Similarly, in theoretical or highly abstracted agent-based models it might be difficult to identify who the stakeholders are. In such cases, modelers can simply state and justify why they do not consider particular pathways relevant, or why it was impractical to attend to all aspects (e.g., due to time or financial constraints).

2.4.1 Assess equity in agent-based models

2.4.1.1 Overview

ABM can be used as a tool to mechanistically describe historical inequities as well as to assess conditions that might mitigate them going forward. For this action pathway, we draw from previous agent-based model applications to demonstrate the capabilities of agent-based modeling for assessing equity and to inspire future equity-oriented ABM research applications.

To develop a set of example applications, we conducted a keyword search in Scopus using the terms (TITLE-ABS-KEY ("agent-based" OR "agent based") AND TITLE-ABS-KEY (model*) AND TITLE-ABS-KEY (*equit* OR disparit* OR *justice* OR *fairness*)). The search terms include a variety of equity-related keywords in an attempt to capture publications from a broad range of fields. We do not intend this as a formal literature review and acknowledge that we have necessarily missed important resources in this process. However, we believe our scope to be expansive enough to reveal general trends and to establish a sufficiently large set of example applications. Within the articles returned by the Scopus search, we restricted our focus to those that apply an agent-based model and indicate clearly within the abstract some link to equity (or equality) in the model's design or application. For each valid article, we categorized its application field and the type(s) of equity assessed, including how equity was quantified and the objects and subjects of equity.

The search[8] returned 344 results, of which 141 we retained as relevant. These articles applied agent-based models in a diverse array of systems (Table 2.1), ranging from abstracted, historical analyses of socio-cultural evolution to detailed, forward-looking assessments of policies for climate change mitigation. The United States was by far the most prevalent location represented in ABMs applied to location-specific case studies. Dominant framings included environmental justice, economic inequality, and disparities in health outcomes and access to urban services. Although rarely explicitly stated, these framings most frequently conceptualized fairness as either applying preference to those with the greatest need or achieving equality in outcomes between groups. The modeled outcomes generally represented measures of social 'goods' (e.g., wealth, resources) or human 'capabilities' (e.g., access to employment, membership in social groups). The modeled dimensions of agent heterogeneity included both continuous measures of variability within a type (e.g., wealth, resources) and discrete, categorical identities (e.g., race, stakeholder type). Further descriptive statistics are included in Table 2.1 and selected examples for each equity dimension are shown in Table 2.2.

[8]Conducted on 31 March 2021.

2.4.1.2 Recognitional equity

Recognitional equity describes the identities and perspectives acknowledged and valued by society (Fraser, 1995). Accordingly, agent-based models can engage with recognitional equity by representing heterogeneity in agent attributes (i.e., identities) as well as objectives or decision-making procedures (i.e., perspectives). Many agent-based models do this. For example, models of cultural evolution often represent dynamic levels of agent altruism and cooperative tendencies to assess how specific human behaviors are selected for in human societies (and thereby recognized and valued by human societies) (Sánchez and Cuesta, 2005). Further, models of power dynamics implicitly represent how power begets recognition, for example through receiving greater weight in decision-making procedures (Orach et al., 2020). Similar ideas exist in literature and models on gendered decision-making (Villamor et al., 2014; Beal Cohen et al., 2019).

Yet, only four articles within our sample specifically engaged with recognitional equity (Table 2.1). All of these articles assessed the implications of broadening the modeling scope to include vulnerable group identities, such as the behavior of slum residents in India or the behavior of Black men who have sex with men in the USA (Table 2.2). For example, Adiga et al. (2018) show that explicitly representing the household sizes and network structures of slum dwellers in Delhi, India is necessary to achieve more equitable public health outcomes. When these characteristics are not represented, their model underestimates the risk to slum dwellers. These examples leverage ABM as a virtual laboratory to test and demonstrate the equity impacts of different recognitional assumptions. Because this requires no methodological extension beyond representing agent heterogeneity, recognitional equity could easily play a more central role in future ABM research.

2.4.1.3 Procedural equity

Procedural equity describes the fairness in peoples' inclusion in and ability to influence decision-making (Tyler, 2000). Because agent-based models explicitly represent decision-making processes and interactions across scales, they are particularly well-poised to incorporate procedural inequities. Accordingly, models focusing on procedural equity were moderately common in our search results, occurring in 41 (29%) of the retained articles.

Within these articles, we observed five principal approaches for dealing with procedural equity (Table 2.1). The most common of these was to represent heterogeneity and/or fairness in individual-level decision-making processes. This was most frequently applied in models in the "culture and game theory" category, whereby agents are specified, for example, to have heterogeneous and dynamic preferences for individual satisfaction and system-wide resource distributions ((Motchoulski, 2019); Table 2.2). Although these dynamics are solely at the level of an individual agent, they relate to procedural equity because certain agents may have more or less ability to influence system-level dynamics through their own decisions.

The second-most common approach takes the procedural focus to a slightly larger scale: group interactions. This was also most common in the "culture and game theory" category. Here, agents interact with each other more directly, for example by sharing or competing over communal resources ((Schank et al., 2015; Klein et al., 2017); Table 2.2). Power dynamics play an important role in these kinds of interactions and can be surprisingly simple to represent. For example, Mahault et al. (2017) model resource transfers within a population of agents and implement a version of the "Matthew effect," whereby the power to acquire wealth grows super-linearly with accumulated wealth. Beyond the possession of physical or material resources, other forms of capital, such the ability to influence others' decisions (Dávid-Barrett and Dunbar, 2013), can confer power and thereby represent procedural inequity.

A more complicated form of decision-making, identified in only four articles in our sample, is governance. Governance is a form of collective decision-making in which people jointly collaborate on system-level decisions. In these cases, agents may have diverging preferences, making it difficult to come to an overarching decision. Deliberation processes can help to achieve better consensus and governance outcomes (Choi and Robertson, 2013), but some cases may require a form of adjudication (an authoritative process for reaching a common decision) ((Motchoulski, 2019); Table 2.2). Different adjudicative mechanisms can be employed, such as a plurality vote (choose the most popular option) (Motchoulski, 2019) or stricter majority or supermajority requirements (50% or more in consensus) (Choi and Robertson, 2013). These mechanisms may disadvantage minority groups, due to their smaller cumulative voting power. Overall, the limited presence of governance elements within our sample reflects a general lacking within the agent-based modeling community (Lindkvist et al., 2020). Yet, these examples demonstrate that governance dynamics not only *can be* represented, but also *have been* represented in agent-based models, thus setting the stage for future governance-oriented ABM research.

Beyond the bottom-up decision-making of individual agents, the fourth and fifth approaches represent procedural equity through top-down processes. The fourth approach models decisions made by a single system-level agent, which may systematically favor or disadvantage certain population groups, for example by preferentially locating environmental dis-amenities in communities with low privilege (Eckerd et al., 2017) or allocating work to the best performing workers (Sobkowicz, 2016). Due to the focus on a single decision-maker, this approach is less unique to agent-based models and is conceptually similar to objective functions in optimization problems or aggregated actors in system dynamic or economic models. Yet, agent-based models can be used to examine both the distributional effects of top-down procedural inequities (see the following subsection) as well as explore bottom-up responses to top-down management decisions ((Farhadi et al., 2016); Table 2.2).

The fifth approach is more methodologically oriented and focuses on potential inequities introduced by the simulation procedure itself. In particular, the order in which agent processes are executed can affect both agent-level and system-level outcomes ((Page, 1997); Table 2.2). In some cases, computational approaches such as multithreading may inadvertently lead some agents to execute early within each time step (Welch and Ekwaro-Osire, 2010), thus creating an unfair playing field. Although these features pose a risk when not adequately considered, they also introduce an opportunity to represent mechanisms of power and privilege through simulation design.

2.4.1.4 Distributional equity

Distributional equity describes the fairness in the allocation of goods in society (Rawls, 2009). Any model in which agents have heterogeneous experiences therefore implicitly deals with distributional equity. This is a feature of almost all agent-based models. Perhaps given this, distributional equity was the most prominent equity dimension within our sample, being present in 117 of the retained articles (83%). We structure our discussion of these articles according to three principal modeling objectives: the conditions leading to distributional (in)equity, the distributional (in)equity itself, and its implications.

There are a wide range of reasons to explore the conditions leading to distributional inequity. For example, understanding the set(s) of equity-promoting or equity-degrading conditions can be useful for identifying equitable policies or interventions among a large set of candidates (de Wildt et al., 2020). Here, tools such as 'scenario discovery' can be used to find combinations of input parameters leading to an outcome of interest (de Wildt et al., 2020). Similarly, agent-based models have been used to identify the behavioral conditions that give rise to desirable or undesirable system-level outcomes, such as sustainable resource management (Schindler, 2012) or social stratification (Dávid-Barrett and Dunbar, 2013). In some cases, the set of hypothesized mechanisms may be insufficient to describe observed levels of distributional inequality (Goodreau et al., 2017), motivating consideration of a wider range of factors.

The approach for analyzing distributional equity itself depends primarily on the type of agent heterogeneity being modeled. In the simplest case of a single agent type, distributional equity is assessed through variability within the type, for example a Lorenz curve or Gini index to represent wealth inequality (Mahault et al., 2017). When heterogeneity exists in other agent characteristics (i.e., there is a diversity of agent types (Page, 2010)), distributional analyses generally relate an outcome measure to an agent attribute. Here, if agent attributes are discrete, a common approach is to compute disparities between the categorical classes, such as two racial groups (e.g., (Orr et al., 2014)). If, instead, the relevant agent attribute is continuous, studies frequently discretize the attribute through stratification (e.g., high-income and low-income (Auchincloss et al., 2011)).

Several articles within our sample focused on the implications of distributional inequality. Two of these were based on economic disparities and their potential to lead to conflict (Kustov, 2017) as well as instability within the US financial system (Cardaci, 2018). More generally, due to the iterative nature of agent-based models, distributional (in)equity in one time step can affect subsequent model procedures, in turn affecting subsequent distributional patterns. Thus, many other applications implicitly internally model the implications of distributional inequity.

One final consideration relevant to distributional equity is how it relates to other outcome measures. Studies frequently reported tradeoffs between equity and economic outcomes (Ponsiglione et al., 2015; Malik et al., 2015; Henry and Brugger, 2018), equity and environmental outcomes (Filatova et al., 2011), or between stakeholder objectives (Farhadi et al., 2016). Only one study in our sample reported win-win outcomes across equity and other dimensions (Bell et al., 2016). Thus, agent-based models are a useful tool for characterizing such multidimensional relationships and can be combined with other methods, such as multi-objective optimization, to identify conditions that minimize tradeoffs across multiple dimensions (Farhadi et al., 2016).

Table 2.1: Descriptive statistics for articles that apply agent-based models to topics relating to equity. The overall sample consists of 141 articles identified in a Scopus keyword search. The sample sizes do not always add up within each category because some abstracts fit within multiple labels and some abstracts were too difficult to classify for some categories.

Category	**Details**	**# of articles**
General information		
System	Built environment (e.g., transportation, housing, energy systems)	38
	Health (e.g., HIV/AIDS, infectious disease)	32
	Culture and game theory (e.g., cooperation, evolution, anthropology, ultimatum games)	27
	Environment (e.g., land-use, water systems)	21
	Economy (e.g., wealth, markets, business)	18
	Science and education (e.g., peer review, teaching)	6
	Crime (e.g., policing, incarceration)	4
	Other	9
Location	Not stated / aspatial	79
	North America	35
	Asia	11
	Europe	6
	Africa	3
	Global	3
	Oceania	2
	Central and South America	2
Journal	JASSS	10
	Lecture Notes in Computer Science	6
	Computers, Environment, and Urban Systems	4
	Other	120
Between whom? (subject)		
	Economic capital (wealth, income)	36
	Race	20
	Other forms of capital (e.g., social status, resource access)	16
	Socioeconomic (i.e., combined social and economic)	11
	Spatial (e.g., neighborhood, country)	11

Category	Details	# of articles
	Decision-making characteristics (e.g., altruism, cooperativeness)	7
	Stakeholder groups	6
	Other	9
Of what? (object)		
	Access (e.g., housing, energy, travel time)	40
	Health outcome (e.g., HIV/AIDS, influenza, diet)	23
	Wealth or income	22
	Environmental (e.g., water quality)	9
	System-level outcome (e.g., equilibrium, emergence of cooperation)	8
	Social (e.g., group membership, status, genetic selection)	7
	Other	18
Fairness criterion		
	Vulnerability/need	38
	Equality	32
	Disparity †	13
	Merit	3
Equity dimension		
Recognitional (n=4)	Implications of representing vulnerable group characteristics (e.g., incorporating gendered or race-specific decision-making and behavior)	4
Procedural (n=41)	Individual decision-making processes (e.g., agents with fairness objectives)	20
	Group interactions (e.g., cooperation, power dynamics)	14
	System-level decision-making processes (e.g., resource allocation)	5
	Governance (multiple agents collaborating on a decision)	4
	Simulation methodology (i.e., ordering of agent processes)	2
Distributional (n=117)	Distributional effects by group identity (e.g., race, spatial location)	60
	Conditions leading to inequality	40
	Distributional effects over population (e.g., Gini index)	17
	Implications of inequality	3

Table 2.2: Examples of approaches to represent or assess equity in agent-based models. We selected examples that represent the range of approaches taken within papers in the Scopus search (Table 2.1).

Category	Example	Lessons and opportunities for equity assessment
Recognitional	Adiga et al. (2018) model the spread of influenza in Delhi, India. They contrast two different network configurations: one that treats slum and non-slum regions the same, and one that represents slum-specific demographics and behaviors. They find that ignoring slum attributes can lead to a 30-55% overestimation in vaccination efficacy, thereby demonstrating the importance of **representing vulnerable group characteristics and behaviors** in both models and public health policy.	It may be important to represent heterogeneity not only in agent attributes but also their decision-making options and procedures.
	Goodreau et al. (2017) model the mechanisms leading to disparities in HIV prevalence between Black and White men who have sex with men in Atlanta, GA, USA. They find that racial assortativity (the tendency to select same-race sexual partners) does not describe the observed racial disparities. Their results demonstrate the importance of **understanding the effects of systemic biases and disparities** in other mechanisms, such as care and communication.	System-level processes may treat different socio-cultural groups unequally.

Category	Example	Lessons and opportunities for equity assessment
Procedural	Page (1997) examines the importance of **the simulation updating procedure** in affecting agent outcomes. They compare emergent system dynamics under two conditions: one where agent states are updated in a random order and one where the updating order is determined by the agents' relative utility increases (a.k.a. "incentive based asynchronous updating"). The two different conditions contribute to vastly different model dynamics.	The execution order for agents' processes makes implicit assumptions about power.
	Motchoulski (2019) examines disagreements between agents' conceptions of distributive justice and the effectiveness of different 'adjudicative' **governance mechanisms** (an authoritative process for reaching a common decision, e.g., through a third party) in resolving these disagreements. In their model, agents have **heterogeneous and dynamic preferences** toward individual satisfaction (i.e., self-interest regarding how many resources they receive) and distributive justice (i.e., the distribution of resources throughout society). Agents respond to the outcomes of adjudicative mechanisms by adjusting the relative weights placed on self-interest and justice.	Agents can have heterogeneous decision-making preferences, which can evolve throughout a simulation based on their experiences. Governance is a form of multi-level decision-making through interaction of bottom-up agent actions and a top-down governance procedure (e.g., plurality vote). Agent-based models can contrast alternative governance arrangements.
	Eckerd et al. (2017) model the emergence of environmental injustice in a community through the siting decisions of environmental amenities and dis-amenities. They vary the extent to which **top-down decisions** about amenity location are driven by cost versus community privilege (e.g., dis-amenities preferentially selecting locations with low community privilege). They find that amenities exert important influences on environmental injustice, meaning that a pure focus on the politics of dis-amenity siting is insufficient.	Agent-based models can formalize and explore the implications of bias in top-down decision-making.
	Schank et al. (2015) model the evolutionary origins of cooperation within societies. In their model, agents play the "dictator game," where one agent decides how to divide a resource with another anonymous agent. They examine the conditions under which **individual preferences toward cooperation**, which are observed empirically but conflict with some behavioral theories, are selected for in a group-based society.	Individual-level decision-making procedures can affect group structure and dynamics.
	Mahault et al. (2017) model different mechanisms for regulating **power imbalances** in an artificial society. When no constraints are imposed on wealth transfer between agents, the agent population invariably becomes polarized, whereby the agents with more opportunity accumulate all of the resources. As further constraints are imposed on wealth transfer, agents' 'frustration,' a measure of the degree to which they act to reduce unsatisfied aspirations, mediates the effects of power imbalances within the society.	Power dynamics can be represented in agent interactions

Category	Example	Lessons and opportunities for equity assessment
Distributional	de Wildt et al. (2020) examine **the conditions that give rise to inequality** through conflicts in agents' capabilities (a measure of the freedom to achieve wellbeing) in the deployment of decentralized energy systems. Their model operationalizes the capability approach of Nussbaum and Sen (1993) and their analysis examines a broad range of energy system and geographic scenarios. They find distinct classes of capability conflicts: sometimes conflicts are inherent to organizational characteristics of energy systems; sometimes conflicts are specific to a type of population (e.g., affluent); and sometimes conflicts occur between population groups.	Scenario discovery can be used to examine the conditions that give rise to inequality.
	Farhadi et al. (2016) model **tradeoffs between stakeholders' objectives** in groundwater management: reducing water deficit (the farmers), increasing equity of groundwater allocation (government executive sector), and reducing groundwater drawdown (environmental protection institutes). They model compromises between the three stakeholders and identify a Pareto-optimal set of solutions using an optimization method. They then use an agent-based model to represent how farmers respond to management decisions and adjust the optimal solutions according to these social conditions.	Distributional inequity can arise both within agent groups and between different types of agents (/stakeholders). Agent-based models can be combined with other modeling approaches that quantify tradeoffs and generate prescriptive solutions (i.e., multi-objective optimization).
	Smart (2019) models 'colorism,' the prejudice toward allocating privilege to lighter skin color, in the context of policing in the US criminal justice system. Their model contains three citizen groups: lights, mediums, and darks. They found that aggressive policies to counteract colorism yielded counterintuitive **distributional effects between groups**. Specifically, agents in the middle of the skin color spectrum experienced higher rates of incarceration. The results demonstrate the importance of broadening the description of colorism to include those in the middle of the color spectrum.	Policy effects can be stratified according to dimensions of agent heterogeneity.
	Cardaci (2018) models the **implications of inequality** for the 2007-08 financial crisis in the United States. Their model represents how inequality in income can lead to expenditure cascades that result in accumulation of household debt, increasing the fragility of the economy and paving the way for the financial collapse.	Distributional inequity can be used as an input.

2.4.2 Recognize modeler positionality and bias

2.4.2.1 Reflexive themes

In many cases of socio-cultural injustice, a necessary preliminary step toward systemic transformation is to recognize the problem. Accordingly, recognitional aspects appear at the forefront of other equity-oriented science frameworks (Chandanabhumma and Narasimhan, 2020; Leach et al., 2018), in qualitative research through 'positionality statements' (Milner, 2007; Holmes, 2020), as well as through acknowledging 'privilege' in antiracist and feminist methodologies (McIntosh, 2020). Barnaud and Van Paassen (2013) succinctly state: "if a designer does not make his biases explicit, he risks imposing them unconsciously." The issue of recognition is thus of utmost importance for rigorous and equitable model development. Yet, such recognitional features are surprisingly absent in ABM research.

We conceptualize recognitional considerations as pertaining to five distinct themes: modeler positionality, problem framing, data, process quantification, and model interpretation. The first theme (positionality) is about the modeler's identity, how this may have influenced their experience and perception of the world, and how it relates to the research context (Milner, 2007). Here, 'identity' constitutes a combination of social and cultural factors, such as race, gender, political beliefs, age, and experience. Different facets of a modeler's identity may position them as an insider or outsider with respect to the research context (Holmes, 2020), together forming a psychological 'distance' from the research phenomenon. For example, as an insider to the culture being studied, the researcher has a 'lived familiarity' and a deeper *a priori* understanding of the culture, yet may be unable to sufficiently detach themselves from the culture to study it without bias (Kusow, 2003).

The remaining four themes each correspond to an arrow entering and/or exiting the modeler in our framework (Figure 2.1), thereby underscoring the importance of reflexivity at all stages of agent-based model development and application. First, 'problem framing' describes the narratives that motivate the research and what problems and solutions these narratives include, exclude, and prioritize (Leach et al., 2010). Second, 'data' describes both how data contain embedded historical inequities and how inequitable data-gathering procedures may generate bias in a dataset (Mehrabi et al., 2019). Third, 'process quantification' encompasses bias associated with translating knowledge into a model format.[9] Fourth and finally, 'model interpretation' closes the loop by describing how the modeler's positionality affects their interpretation and communication of model outputs.

2.4.2.2 Applying the themes

In this action pathway, the modeler engages in a reflexive process of understanding, critiquing, mitigating, and acknowledging sources of potential inequity. Although this does not necessitate changes to the modeling focus or model design, recognition should be seen as an iterative and ongoing component of good modeling practice that, over time, seeks to mitigate potential inequities and therefore influences the direction of model development. Thus, it is important that such efforts do not become perfunctory or solely post-hoc considerations; recognition should be understood as a *means* for opening doors to deeper engagement with inequity, rather than an *end* in itself. A critical approach is required to disentangle how each aspect of the modeling process is subjective, what is included and/or might be excluded, who might be the winners and losers, and who has decided about this. Yet, no matter how reflexive a researcher is, it is impossible to objectively describe something as it is (Holmes, 2020) and therefore this process should seek to understand

[9] Subjectivity in quantification is a well-recognized issue in agent-based modeling, where different modelers reach different conclusions given the same set of data, or independent implementations from a single model documentation lead to different results (Zhang and Robinson, 2021).

and acknowledge the model as a non-neutral object (Voinov et al., 2014; Barnaud and Van Paassen, 2013), rather than to neutralize or neuter it.

To assist modelers in this exercise, we enumerate a set of reflection prompts for each theme (Table 2.3). We suggest two alternative approaches for operationalizing these considerations. The first is through a "positionality and bias statement," which synthesizes the insights from reflection on these five themes into a standalone document. The document may be narrative or simply a list of answers to the prompts. It can be updated over time, as necessary, and included with any written communication of the model and its outputs (e.g., an appendix to an academic research article).

The second approach integrates the questionnaire into standardized modeling protocols within the ABM field, as well as more generally into the presentation of models and results. In Table 2.3 we suggest how these reflections could be embedded into academic articles, the TRACE framework, and the ODD protocol. TRACE and ODD are frequently used to document and describe agent-based models (Grimm et al., 2020, 2014) and already encourage the modeler to reflect on model design decisions. The questions outlined in Table 2.3 could supplement the existing prompts in either of these frameworks, in order to give more targeted reflection on positionality and bias. In line with the TRACE and ODD practices, if space or structure does not allow for comprehensive exposition, we suggest briefly summarizing these themes within the main body of a manuscript and including the full set of reflection questions as an appendix or supplement. Not all themes are necessarily always relevant, for example if the model does not draw from data.

Appendix A provides an example standalone positionality and bias document, using the prompts in Table 2.3, for an agent-based model of smallholder agriculture. The model draws on secondary empirical data from Ethiopia and was used to examine the efficacy of alternative strategies for building smallholder climate resilience. The principal modeler identifies as a White male of European descent. Other members of the research team all identify as male and work at predominantly White academic institutions in the United States. Two of the research team identify as European American and one as a Marwari Bihari out of place. The research team therefore identifies as outsiders to the research context. Further, they did not engage local communities or decision-makers in the development of the model. The psychological distance and potential for (inequitable) bias are therefore both large.

The reflexive exercise in Appendix A provided an opportunity for the modelers to reflect on their positionality and how the research may have perpetuated existing inequities. Further, it helped to make certain assumptions inherent to the problem framing transparent, in particular that the adaptation strategies being examined were both top-down, external interventions that may not agree with local belief systems. The exercise was conducted after the modeling project had concluded and therefore was not able to affect the model development or application itself. As we have stated, such reflection should ideally be embedded within the iterative modeling cycle. However, this example demonstrates an important point: that it is unreasonable to expect all ABM research to immediately become "equitable" in every way. Realizing equity in ABM is a large (and perhaps impossible) endeavor. The reflexivity statement, even when a post-hoc consideration or only pertaining to a subset of the five themes, begins to push the needle toward this goal.

Table 2.3: Questions to assist a reflexive approach to recognize inequities in the modeling process. For all questions, the researcher should consider how their answers relate to (in)equity. Responses to these questions can either be synthesized into a standalone bias and positionality statement or integrated with other ABM communication. The right-most column provides suggestions for where to include the recognitional themes within standardized ABM documentation protocols.

Theme	Questions	Where to include
Positionality	• What are the racial and cultural backgrounds and identities of the modeler(s)? • How might these identities have influenced how the modeler(s) experience the world and approach research? • How do these identities, worldviews, and objectives relate to the participants and/or context of the research? (e.g., in what ways are the modelers insiders or outsiders?) • What is the social, institutional, and historical nature of inequity in the context of the research?	• Article: methods; acknowledgements • TRACE: Problem formulation (#1) • ODD: Purpose
Framing	• What **narratives** underlie the formulation of the challenge, problem, or research questions? • What kinds of **solutions** does the problem framing open itself to? • What entities are included/excluded? Who are the **actors** involved within the framing of the problem and solutions? • What **outputs** are included and prioritized? • What is the **scale** of focus within the problem framing? • What **theories** and/or relationships is the conceptual model predicated on? If relevant, are there alternative explanations?	• Article: introduction; methods; discussion • TRACE: Problem formulation (#1) • ODD: Purpose; Entities, state variables and scales; Design concepts (basic principles)
Data	• How could historical patterns of inequity exist within the data? • How could marginalized people or groups be misrepresented in or excluded from the data? • How could the process of data collection have perpetuated inequity?	• Article: data and methods; discussion • TRACE: Data evaluation (#3) • ODD: patterns; input data
Process quantification	• What subjectivity is involved in defining model variables and/or translating information from data sources into the model format? (e.g., are model variables latent constructs?) • Could the inclusion or exclusion of model processes misrepresent or lead to bias against certain groups?	• Article: methods • TRACE: Model description (#2); Conceptual model evaluation (#4) • ODD: process overview and scheduling; input data; submodels
Model interpretation	• How could calibration and validation procedures prioritize models that (dis)advantage certain modeled subgroups? • How could pre-conceived understandings or objectives affect which model structures and outputs are considered acceptable and subsequently communicated?	• Article: methods (model calibration); results • TRACE: Model output verification (#5); Model output corroboration (#8) • ODD: purpose and patterns; observation

2.4.3 Engage stakeholders and society

In the final action pathway, the modeler engages stakeholders throughout the modeling process. There exists a rich body of research on stakeholder engagement, participatory modeling, and co-production of knowledge. Much of the guidance within this existing work already touches on equity dimensions. Our discussion of this action pathway is therefore relatively brief, and we use this section to relate existing work to the equity language within our framework. For the purposes of this chapter, we consider a stakeholder to be anyone who exists within the modeled system, could use model outputs, or could be affected by decisions inspired by the model. Stakeholders comprise actors across a spectrum of scales of agency, from individual behavioral change (e.g., community members) to affecting institutional and governance structures (e.g., policymakers). This may preclude the relevance of stakeholder engagement in some contexts, such as highly theoretical ABM research.

Engaging stakeholders in ABM can improve equity through at least two mechanisms. The first is through empowerment. Stakeholders each bring their own perspectives and opinions (i.e., positionalities) that can contribute to reducing the effects of the modeler's individual bias. Including a mix of diverse identities (i.e., 'insiders' and 'outsiders' (Holmes, 2020)) thus reshapes the filter and lens through which knowledge is projected into the model (Figure 2.2) and has the potential to profoundly affect the ABM process and outcomes (Steger et al., 2021a). For example, stakeholders may identify actors omitted from the model (thereby improving recognitional equity) or offer alternative descriptions for model processes. Stakeholder engagement at early stages can affect the problem definition and framing (Voinov and Bousquet, 2010; Steger et al., 2021b) and thereby the scope of the model (for instance, if some elements *should not* be modeled due to problems with positivist interpretations of indigenous knowledge (Smith, 2013)). In all of the above, stakeholder engagement improves procedural equity because it increases stakeholders' influence over the decisions made throughout the ABM process.

Second, engaging stakeholders improves equity in the distribution of access to science and knowledge. Particularly when stakeholders are collaboratively involved in a process of co-production, models can act as boundary objects to facilitate knowledge generation and shared understanding (Voinov et al., 2016; Lemos et al., 2018). This two-way learning process increases the legitimacy of models and drives better decisions (Reed, 2008). For example, modeling projects with frequent exchanges between modelers and decision-makers have shown greater potential to influence policymaking (Will et al., 2021). Additionally, integrating collaboration into scientific funding can generate more effective and usable research outcomes (Arnott et al., 2020). The contributions to equity through this mechanism are thus twofold: equity in access to knowledge and increased propensity for ABM to drive equitable decision-making.

Yet, if stakeholders are not appropriately recognized and empowered, stakeholder engagement amounts to a co-opting or tokenizing of knowledge (Arnstein, 1969). Accordingly, researchers have developed a range of guidelines for stakeholder engagement. Critical insights include the importance of involving stakeholders early within the project (Reed, 2008; Steger et al., 2021b) and taking time to develop trust (modeler-stakeholder, stakeholder-stakeholder, and stakeholder-model) (Voinov and Bousquet, 2010). Attention should be paid to stakeholders' motivations for involvement (Voinov et al., 2016), the diversity of stakeholder interests and identities (Steger et al., 2021a), who is selected as a leader (Hämäläinen et al., 2020), and potential power dynamics between stakeholders (Barnaud and Van Paassen, 2013; Voinov et al., 2018). In all cases, stakeholder engagement takes significant time and effort (both for the stakeholders and the modeler) (Voinov et al., 2018), and modeling objectives should be accordingly adjusted from the model as an end product to participatory modeling as a *process* for shared learning (Reed, 2008). The process is not one-size-fits-all and clearly must be tailored to individual contexts. Yet, neither is there only one appropriate approach for a given context, and there are a range of levels of engagement possible (Pretty, 1995).

Finally, at a more meta-level, this action pathway also pertains to the questions of who is allowed to participate in science and who is given access to scientific knowledge. Some of these considerations are beyond the scope of agent-based modeling per se, such as who makes decisions on research funding priorities. However, modelers can improve access to scientific knowledge by publicly posting their code (e.g., at CoMSES.net), using open-source software (e.g., Netlogo, Python, R), and making research articles open access.[10] Such considerations transcend the spectrum of theoretical to applied models and play an important role in the democratization of science.

2.5 Discussion

Historical, current, and future societal inequities demand that greater attention be paid to equity within ABM research. This chapter contributes to this goal at two levels. First, it develops a conceptual framework for understanding the equity-ABM interface. Second, it presents three action-oriented pathways for achieving greater equity-ABM integration. These pathways, both individually and together, reduce the risk of ABM inadvertently perpetuating inequity and demonstrate the opportunities for using ABM to advance equity.

[10] Albeit noting the issues of predatory journals and how publication fees can be prohibitive for researchers with less financial means (ironically, the very populations open access systems are designed to benefit).

2.5.1 Benefits and opportunities in the action pathways

A core conceptual advance in our framework is the centrality of modeler recognition and reflexivity, which is rarely (if ever) addressed in ABM research. To attempt to narrow the gap between theory of reflexivity and its practice, we presented a set of reflection questions to assist modelers in a reflexive exercise (Table 2.3). Recent history in participatory modeling suggests that the uptake of modeler reflexivity is possible. Participatory engagement of stakeholders was once a peripheral consideration within the modeling field, yet is nowadays a mainstay in many modeling application contexts (Voinov and Bousquet, 2010). The rise of participatory modeling to the mainstream was, in part, precipitated by the increasing complexity of modern decision problems. These new problems, for instance sustainability, pushed the boundaries of hegemonic positivist methodologies (Pretty, 1995) and demanded integration of more diverse ways of knowing. In these contexts, stakeholder engagement was demonstrated to achieve better decision outcomes and greater learning (Reed, 2008), facilitating its uptake by scientists and funding organizations (Arnott et al., 2020). In a similar vein, recent societal events (e.g., the Black Lives Matter movement) have challenged society to reflect on its positionality, biases, and (hidden) forms of discrimination. This has already led to relatively radical changes in parts of the United States criminal justice system. We believe that a similar recognition is both necessary and possible within the ABM field. Our action pathway illustrates how these reflection questions are compatible with existing ABM frameworks (Table 2.3), thereby facilitating their integration into good modeling practice. Future research is needed to test whether modeler reflexivity contributes to better modeling outcomes in different contexts.

With respect to assessing equity in agent-based models, the results to our keyword search demonstrate a considerable existing research base. This suggests that representing equity is not limited by methodological constraints. From the existing research, we synthesized a set of generic approaches and specific examples of equity-oriented agent-based models (Table 2.1; Table 2.2), which we hope can inspire future research on this theme. In particular, ABM research could more deeply engage with issues of socio-cultural recognition, i.e., how society differentially treats different socio-cultural identities (e.g., race) and thereby contributes to disparities in outcomes. Such research could use agent-based models to formalize multiple understandings of reality (i.e., engage with nonpositivist methodologies (Lincoln et al., 2011)) to contrast a broader range of ontologies, narratives, and definitions (Pretty, 1995; Wilson, 2001; Leach et al., 2010). Beyond this, our results contained relatively few models dealing with cross-scale power dynamics or governance, which are important to many real-life contexts (Ostrom, 2009). Future research could therefore integrate such dynamics more deeply, as well as explore cross-scale thresholds and tipping points toward equitable system states (Radosavljevic et al., 2021). For example, there is considerable research on 'regime shifts' in ecology (Lade et al., 2013; Filatova et al., 2016; Horan et al., 2011), which could be extended in agent-based models to include social thresholds and systemic shifts in power balances and governance forms (Schlüter et al., 2021).

Yet, modeling advances are of limited value if agent-based models are not used to inform decision-making or meaningfully contribute to societal knowledge. The insights gained from the participatory modeling field (the 'engage' pathway within our framework) demonstrate the utility of stakeholder engagement and knowledge co-production (Reed, 2008; Steger et al., 2021a). This research underscores a need to shift focus from the model outputs to the modeling process itself (Verburg et al., 2019; Voinov and Bousquet, 2010; Sandker et al., 2010). Future ABM research could therefore more frequently use models as boundary objects to stimulate societal debates about equity and work toward forming consensus in diverse stakeholder groups (Voinov et al., 2014).

Although the three action pathways are conceptually distinct, they are in practice highly interdependent. For example, equity-oriented agent-based models must codify different socio-cultural identities and this codification may demand greater stakeholder engagement. Further, processes of stakeholder engagement need to reflect on the stakeholders' own biases, and it may be useful to develop (collaboratively with stakeholders) a statement of stakeholder bias and positionality. Thus, equity-oriented ABM research may need to bundle approaches from different action pathways.

2.5.2 Pragmatic constraints and generalizability

As with any call for more rigorous modeling practice, our recommendations have the potential to add considerable burden to modelers operating with limited time, financial, and social resources. Thus, it is important to consider the benefits and costs of improving equity through any action pathway. We have emphasized that reflection provides value in itself, meaning equity can be improved without substantive changes to the model or the modeling process. Moreover, given that many agent-based models already represent actor heterogeneity, transitioning from 'heterogeneity' to a focus on 'equity' may only require attaching a normative value to model output heterogeneity. This can be as simple as acknowledging that inequality within an agent population is undesirable, or that outcomes such as "the poor people get poorer" represent inequitable or unfair distributions. Additional effort is required to deconstruct and mitigate inequities (e.g., by broadening the problem framing or mitigating bias in data), but doing so early within the model design process is likely to yield the largest benefit to cost ratio.

Equity considerations may be more or less relevant in different modeling contexts. Agent-based model applications are incredibly diverse, both in terms of application area, level of realism, and level of knowledge integration (Schlüter et al., 2019b; O'Sullivan et al., 2016). Moreover, models and modeling projects have diverse objectives, such as system understanding, prediction, forecasting, informing decision-making, and social learning (Kelly et al., 2013). It might seem intuitive that equity considerations are most relevant for highly applied models that address current societal issues. However, this is not necessarily the case. For example, Sugarscape (Epstein and Axtell, 1996) presents a highly abstracted description of society, but sheds considerable insight into power dynamics and the emergence of inequality. Should Epstein and Axtell have engaged with "stakeholders" in developing their model? Perhaps not. However, their unique positionalities unquestionably shaped how they interpret the world and approached the modeling. Ultimately it is difficult to develop generalized guidelines for when equity could or should be considered, and we leave it as a task for modelers to reflect on in their individual contexts. However, reflecting on equity is particularly important when there is a large degree of physical or psychological distance between the researcher and the research context (Holmes, 2020).

There are several considerations that our framework and action pathways do not explicitly address. One is how they generalize to collaborative modeling projects with *multiple* principal modelers. In such cases, it may be necessary to describe the role that each modeler played within the modeling project and develop a joint bias and positionality statement. Just as modelers are themselves also stakeholders (Voinov et al., 2014), progress on equity can be made by ensuring that the modeling team represents diverse socio-cultural identities and perspectives. A second question not covered by our framework and pathways is how to determine whether ABM is an appropriate methodological paradigm for the given research context (Schlüter et al., 2019b). Agent-based models require extensive time and expertise to develop, validate, and apply (Voinov et al., 2018), and these resource demands can be prohibitive. Related to this, we do not address how the equity dimensions and action pathways generalize to other modeling approaches, for instance system dynamics, economic equilibrium, state-and-transition models, or bioeconomic simulations. The 'assess' pathway is relatively specific to ABM, as it is underpinned by agent-based models' abilities to represent heterogeneity and nuanced decision-making processes, which are frequently limiting factors in other modeling paradigms (Emmerling and Tavoni, 2021). Nevertheless, the 'engage' and 'recognize' action pathways are likely relevant in other process-based modeling domains, as they rest on a more general description of a modeler who is nested within a socio-political context.

2.6 Conclusions

In this chapter, we developed an operational, conceptual framework for the equity-ABM interface. The framework describes the modeler as a filter and a lens at the boundary of the model and society, and thereby a key locus for action on equity. We identified three action pathways through which agent-based modelers can improve equity, including engaging stakeholders, recognizing positionality, and modeling (in)equity.

The challenge of "equitable ABM" is by no means insurmountable. We found that ABM research already engages with equity, to some extent, in many respects. There is thus ample precedent and existing knowledge from which future modeling work can draw. We have also underscored that improving equity does not require drastically changing the direction or focus of the modeling project. Rather, it minimally requires some time reflecting on how the problem is being framed and how this is conditioned by researcher positionality. By adopting a more reflexive attitude, agent-based modelers can better engage with society and non-modelers to increase the usability of knowledge derived from agent-based models. Further, ABM is uniquely positioned to model recognitional, procedural, and distributional equity, so could be a useful tool for researchers interested in equity but not experienced with ABM. Taken together, we believe the equity-ABM interface is a fruitful area for future research and a mechanism through which ABM can achieve greater societal impact.

Chapter 3

Resilience and Equity[1]

Strategies aiming to increase the climate resilience of smallholder agricultural systems may not equally benefit all groups of the smallholder population. To reduce the potential for aggravating existing vulnerabilities, quantitative resilience analyses therefore need to acknowledge the possibility for inequities in the effects of proposed resilience-enhancing strategies (RESs). In this study, we develop, validate, and apply a household-level agent-based model to explore the equity of climate RESs in an Ethiopian smallholder farming system. Specifically, we study the potential effects of two RESs, involving access to seasonal climate forecasts and increases in non-farm job availability, on household food security under climate variability. We measure these effects in two distinct ways: "poverty-reduction," which describes food security improvements relative to existing conditions; and "shock-absorption," which isolates the strategies' effects on food security during and following a drought. Our results reveal that the different measures of resilience lead to divergent assessments of equity in policy effects. Relative to baseline levels of food security (poverty-reduction), both strategies disproportionately favor the most vulnerable households—i.e., they are equity-enhancing. Under this assessment, increases in job availability provide slightly stronger benefits to the most vulnerable households than climate forecasts. However, when isolating the effect of a drought (shock-absorption), both RESs benefit the moderately vulnerable households at the expense of the more vulnerable households—i.e., they are inequitable. These results demonstrate that a pure focus on poverty reduction may be insufficient to promote equitable development. Given the prevalence of climate shocks in smallholder systems, future studies of resilience should therefore jointly consider both poverty reduction and shock recovery, as well as the potential for inequity in the effects of RESs.

[1]Published as Williams T.G., Guikema S.D., Brown D.G, Agrawal A. "Resilience and equity: Quantifying the distributional effects of resilience-enhancing strategies in a smallholder agricultural system" *Agricultural Systems* http://dx.doi.org/10.1016/j.agsy.2020.102832

3.1 Introduction

Weather and climate directly affect the prosperity of smallholder agricultural communities. Vulnerability to climatic shocks often materializes as food insecurity and malnourishment, which currently affects over 800 million people around the world (FAO, 2008). In the pursuit of sustainable development, many strategies for increasing climate resilience have been proposed. However, the effects of such interventions are not necessarily shared equally throughout a population. This invokes questions regarding equity; is it possible that by increasing resilience at an aggregate level, vulnerable populations are left behind?

To evaluate and compare resilience-enhancing strategies (RESs) therefore requires an approach that can represent the distribution of potential effects over a population, as well as include the interactions and temporal dynamics inherent to socio-environmental systems. Agent-based modeling can address these challenges and has been extensively used to model agricultural systems (Matthews et al., 2007; An, 2012; Kremmydas et al., 2018), with common emphasis on the impacts of policy interventions (Berger et al., 2017; Ng et al., 2011; Kiesling et al., 2012; Amadou et al., 2018; Ziervogel et al., 2005). Application of agent-based models (ABMs) for resilience has included, among others, evolving community vulnerability under repeated natural disasters (Reilly et al., 2017; Tonn and Guikema, 2018), resilience of pastoral and common-pool resource systems (John et al., 2019; Dressler et al., 2019a; ten Broeke et al., 2019), resilience of agricultural systems (Tian et al., 2016; Bitterman and Bennett, 2018), and infrastructure system resilience (Baroud et al., 2014; Çağnan et al., 2006; Tabucchi et al., 2010).

However, distributional effects are seldom considered in quantitative resilience analysis. In the context of smallholder agriculture, households are highly diverse, so may respond differently towards policy initiatives (Kansiime et al., 2018) and agricultural coping strategies may lead to asymmetries in resilience between different groups (Béné et al., 2012) or even reinforce poverty (Miller et al., 2010). With an exclusive focus at the population level, these effects may be missed. This has prompted discussions about the distribution of resilience and questions such as "resilience for whom?" (Cutter, 2016). Such a perspective is also prevalent in "pro-poor" policy analyses (Dorward et al., 2004; Tarawali et al., 2011). ABMs can represent a heterogeneous population of interacting agents, so are well-suited to explore distributional effects (e.g., (Berger et al., 2017; Evans et al., 2011; Dressler et al., 2019b)), facilitating a focus on equity when evaluating and prioritizing potential RESs.

With a focus on equity, how we measure resilience is critical and should carefully consider the ways in which vulnerable groups are represented by resilience measures. Applications of ABMs to resilience are typically *specific* (i.e., resilience "of what, to what") (Carpenter et al., 2001; Urruty et al., 2016), focusing on the impact of shocks on some measure of system functionality, often desiring the return (or "bouncing back") to some pre-shock state. These types of approaches have been criticized, as such measures may overlook vulnerable populations (Miller et al., 2010; Cretney, 2014; Manyena et al., 2011), potentially masking or even inadvertently increasing inequalities.

In this study we develop, validate, and apply an ABM to explore the distribution of climate vulnerability in a stylized Ethiopian mixed crop-livestock smallholder agricultural system and the potential resilience-enhancing benefits of selected interventions. Due to the historical prevalence of food insecurity and poverty in Ethiopia, the diverse range of climatic zones and livelihoods, and its development potential, Ethiopia provides a useful context in which to explore the effects of climate and agricultural development on smallholder populations. The model is based on the notion that time and resource allocation decisions made by smallholder households — influenced by interactions with variable and uncertain climatic and economic conditions, as well as other households and communal grazing land — aggregate to monthly variations in food security. We measure household food security as the monthly occurrence of food shortages. We measure resilience as the extent to which food security might be compromised by a drought (vulnerability), in combination with the time taken to recover (recovery).

In our analysis we examine the resilience-enhancing potential of two common measures for supporting rural producers: increases in urban job availability, allowing households more steady streams of income throughout the year; and seasonal climate forecasts that provide information that enables better-informed agricultural decisions. Given the different mechanisms by which these operate, we are interested in exploring how the interventions might benefit different groups of people and how these ways might be different. We examine the effectiveness of these strategies at maintaining food security during and following a drought relative to the levels of food security in: (1) the absence of the drought ("shock-absorption") and (2) the absence of the intervention ("poverty-reduction"). In all cases, we interpret the results of this analysis in light of our interest in understanding equity in quantitative resilience assessment.

There are four main contributions of this chapter. First, our exploration of the distributional effects of drought and policy interventions is a much-needed advance for quantitative resilience analysis. Second, in doing so, we operationalize resilience in a smallholder agricultural setting and explore and discuss the implications of different resilience measures for policy assessment. Third, for a case study set in an idealized Ethiopian community, we demonstrate the utility of this approach and its potential to be used to assess resilience in socio-environmental systems. Finally, because we conceive and operationalize resilience of food security in a more encompassing way than other quantitative measures such as crop yield or agricultural production, we incorporate both environmental and social processes into a single outcome measure of relevance for decision-makers and sustainable development.

3.2 Methods

3.2.1 Approach

We use an ABM to model the livelihoods of smallholder households, including how these livelihoods may be differently affected by climate shocks and policy interventions. Agent-based modeling is a suitable tool for this purpose because it allows for the simulation of: (1) a population that is heterogeneous in its characteristics and spatial location, which is critical when examining distributional effects; (2) the interactions and feedbacks inherent to agricultural systems; (3) stochasticity and uncertainty, which are important to the notion of resilience; and (4) the effects of ex-ante experimental intervention.

3.2.2 Model description

Our model description generally follows the ODD+D format (Müller et al., 2013). See Appendix B.7 for further details.

3.2.2.1 Overview

Purpose The purpose of the model is to produce assessments of household-level and community-wide resilience to climate shocks in a stylized Ethiopian smallholder mixed crop-livestock farming setting. The model is not designed to make policy recommendations for a specific location. Rather, we designed the model as an experimental platform to evaluate and prioritize interventions with respect to their community-level and distributional resilience-enhancing potential more generally.

Entities and state variables Each agent represents a household that engages in farming and livestock rearing and can partake in off-farm income-generating activities (Figure 3.1). Each agent is defined by sets of (1) non-changing attributes such as household size, land ownership, and risk aversion, (2) evolving state variables such as food stores and livestock herd size, and (3) beliefs about variables such as crop prices and job availability. In addition, each agent has a non-evolving preference for either income or leisure, which influences decision-making.

Each agent owns a livestock herd, which produces milk on a monthly basis and is grazed on a combination of on-farm crop residues and communal rangeland. Livestock can be bought and sold at the beginning of each year as well as sold as a coping mechanism throughout the year in order to meet food and cash requirements.

The agents populate a spatially explicit environment, which consists of cropland and rangeland. Each agent occupies a number of cropland pixels, which are each characterized by a static measure

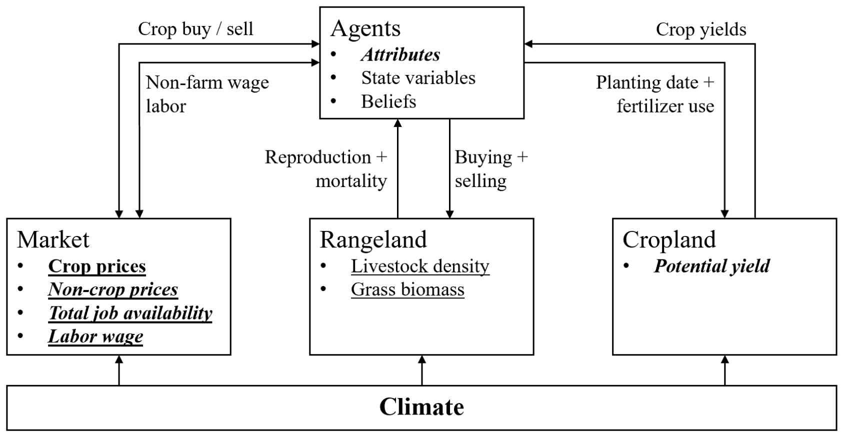

Figure 3.1: The primary entities in the ABM. **Bold: exogenous**. *Italics: constant over time*. Underline: region-level.

of potential crop yield (i.e., agricultural suitability). A single crop type (maize) is modeled[2], which is grown in a single cropping season (Block et al., 2008). Agents cannot buy or sell cropland or reclaim rangeland for agriculture. Rangeland pixels represent commonland used for grazing of livestock. Rangeland is characterized at the regional level by an evolving livestock density and grass biomass (Janssen et al., 2000).

Agents interact with their neighbors through shared beliefs, with the rangeland and its limited capacity, and through the market, which has a limited supply of non-farm jobs and unlimited supply and demand for crop and livestock products. Non-farm employment, for which agents earn a fixed wage, could represent either employment in an urban center or on another farm.

Setting and spatio-temporal representation The model region is located in Amhara in the Ethiopian highlands (Figure 3.2), primarily representing agro-ecosystem 1 of the Choke Mountain area (Simane et al., 2013). This region is characterized by extensive production of sorghum, tef, and maize, which are grown primarily for subsistence. The spatial resolution and extent of the model landscape are 0.15 ha (39x39 m) and 11,600 ha (8.9x13.1 km) respectively. There are approximately 4,250 agents. Spatially explicit climate and landcover information was used to parameterize the landscape, and agent population parameters were drawn from the World Bank's 2015 Living Standards Measurement Study (LSMS), subsetted to Amhara. Because these demographic data come from a wider region, the agents are not representative of the exact conditions in

[2]This is a simplification of reality as crop diversification is an important aspect of smallholder agriculture (Tekle-wold et al., 2013; Barrett et al., 2001). However, in the modeled region, soil type and quality pose major constraints to crop production (Simane et al., 2013) and it is rare for farmers to switch crop types. Additionally, maize is fairly sensitive to the effects of rainfall (Eggen, 2018), so it provides a useful backdrop against which to assess agricultural resilience. This simplification is hence deemed to be reasonable for the purposes of this chapter, but we acknowledge that it likely leads to an underestimation of resilience.

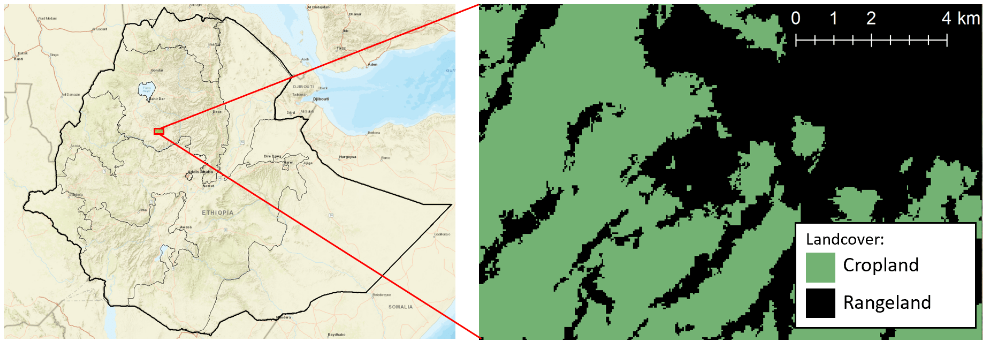

Figure 3.2: The ABM is designed to be representative of an Ethiopian mixed crop-livestock smallholder farming system. Climate and landcover data are drawn from Choke Mountain, Amhara, Ethiopia.

Choke, but this region serves as an exemplar of the types of places represented in the model and its geographic conditions serve as the environmental backdrop for our analysis.

The model is run for a period of 30 years, with agents making annual agricultural management decisions and monthly food consumption decisions. This time period enables the observation of inter- and intra-annual recovery from climate shocks, but we do not model environmental degradation, population growth, or climate change.

System boundaries First, we assume that the modeled region is of insufficient size to influence regional market dynamics and that the agents cannot influence climate conditions through their decisions. Hence, prices, labor wages, and climate are exogenous (see Figure 3.1). Second, the effect of climate on yields is exogenous to the system; agents can influence their experienced yields by shifting their planting date or applying fertilizer to their fields, but we do not model environmental feedbacks in the farming system (e.g., soil degradation). Third, regional livestock reproduction/-mortality is influenced by both the exogenous climate conditions and the regional population of livestock, which is an emergent outcome of the agents' decision-making process. Fourth, although regional job availability is exogenous, the supply is limited and agents' experiences are influenced by individual labor allocations, which collectively determine the regional demand for wage labor.

Process overview and scheduling Income and food security are modeled at a monthly time step, but the primary model processes operate on an annual basis (Figure 3.3). First, given their beliefs, agents make their livelihood decisions. These decisions and the actual climate are then used to calculate crop yields, prices, and job allocation. Monthly food consumption and income are then simulated under these realized conditions, with agents having the ability to self-consume their own yields as well as buy/sell from the market. Here, agents can engage in coping measures (stress selling of livestock and food stores) if they cannot meet their food requirements. At the end of the year, livestock reproduction and mortality are simulated (rangeland dynamics). Finally, informed

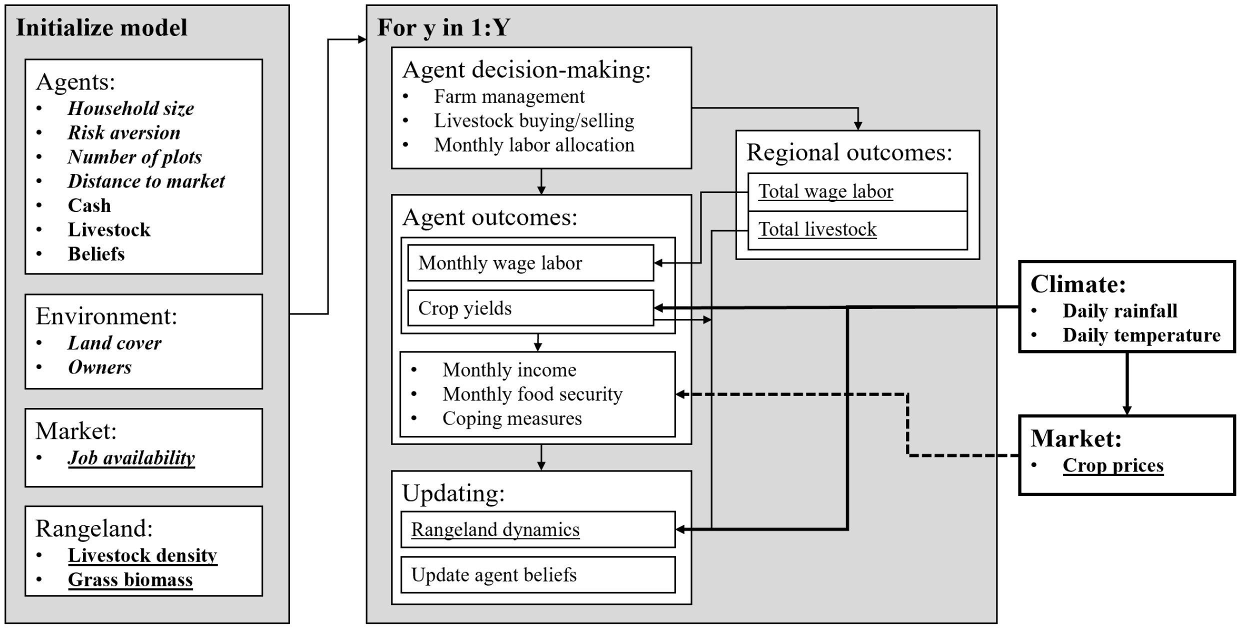

Figure 3.3: Overall flow of the ABM. **Bold: exogenous**. *Italics: constant over time*. Underline: regional level.

by their and their neighbors' experiences, agents update their beliefs.

3.2.2.2 Design concepts

Theoretical and empirical background The livelihood processes included in the model represent the main livelihood activities in Ethiopian mixed crop-livestock systems and are consistent with previous agent-based modeling efforts (Berger et al., 2017; Hailegiorgis et al., 2018; Hajjar et al., 2019). Rangeland dynamics are modeled using a simple system dynamics model (Janssen et al., 2000) and the effects of water stress on crop yields are modeled using established process-based methods (FAO, 1984, 1998; Block et al., 2008). We assume that the agents are boundedly rational and pursue a method of satisficing inspired by Kaufman (1990). This acknowledges the hierarchical objectives of smallholder farmers, in which poorer households have been observed to be primarily motivated by risk minimization and income stabilization, while wealthier households pursue maximization strategies (Demissie and Legesse, 2013). We assume that agents have a preference for either income or leisure. The preference of some agents for leisure is in accordance with Chayanovian subsistence or "full belly" theory (Kaimowitz and Angelsen, 1998; Meyfroidt et al., 2018; Chayanov, 1986), in which households seek to satisfy a consumption target with minimal labor input.

Individual decision-making: annual livelihoods Each year, the agents choose between a finite set of livelihood options (Figure B.1). These consist of agricultural management decisions (planting date and fertilizer application), livestock management (buying or selling from herd), and labor allocation (water and firewood collection, farming, livestock, and off-farm). We conceptualize the agents with preference for leisure as "traditional" households that always engage in farming. Under the satisficing framework, agents have two levels of objective; they first pursue the option that leads to the lowest levels of food insecurity (i.e., fundamentally they wish to be food secure), and if two options are equal in this respect, they choose the option with the highest expected utility (defined either by wealth or leisure).

Evaluations of decision options are made under imperfect knowledge, which is represented through subjective agent-level beliefs. These beliefs form as a result of agent experiences and interaction with their direct neighbors, both of which are spatially dependent. Beliefs are updated each year using Bayesian updating.

We explicitly incorporate uncertainty in the decision-making process. The agents' beliefs about crop yields are formulated using a set of several forecasting models (similar to Magliocca et al. (2013)), each of which will give different predictions and consequently different estimated levels of food security. Agents evaluate each decision option using all of these forecasting models separately and then calculate expected food security and an expected risk-averse utility on wealth to resolve the uncertainty. Sections B.7.5 and B.7.4 give more details on beliefs and decision-making.

Individual decision-making: monthly food consumption In addition to this annual decision-making process, we model monthly food consumption using a heuristic (Figure B.2). This heuristic is employed using subjective beliefs in the livelihood decision-making stage and using realized values (e.g., actual crop yields) after the decisions have been made. First, agents attempt to satisfy their consumption requirements (section B.7.6) through their own crop stores and production. Next, they purchase food from the market. If, after this, sustenance has not been satisfied (because of inadequate food stores and cash availability), agents reduce their consumption for the month. However, if the sustenance deficit is above some threshold and the agent has also experienced food insecurity in the previous month, they sell livestock (if possible) as a coping mechanism to fill this deficit. This threshold-based approach reflects the fact that selling livestock can be an important coping mechanism (Dercon and Christiaensen, 2011; Bellemare and Barrett, 2006; Demeke et al., 2011), but that people may prefer to reduce their consumption before selling livestock assets (Mogues, 2006; Little et al., 2006).

Each month, agents that cannot satisfy their food requirements or that are forced to sell livestock are classified as food insecure. This is a binary measure of food insecurity representing the incidence of food shortages at the household level. Although calorie deprivation is an imperfect measure of food insecurity (Headey and Ecker, 2012), our model does incorporate the FAO's pillars of availability, access, and stability (FAO, 2008) at a household level. Additionally, it enables model calibration using questions from the LSMS.

Learning We do not model learning as defined by Müller et al. (2013) (i.e., the decision rules themselves do not change).

Individual sensing Each year, agents observe their own crop yields, success at receiving off-farm employment, livestock reproduction/mortality, and regional crop prices. Agents also observe their immediate neighbors' beliefs about these values. They do not make mistakes in these perceptions.

Individual prediction Agents predict future conditions using their beliefs, which are represented using probability distributions. In all cases except crop yields, the expected value of the distribution is used (see section B.7.4).

Interaction The model contains both direct and indirect interactions. First, agents interact directly with their immediate neighbors through the sharing of beliefs. This influences decision-making. Second, agents interact indirectly through the rangeland; regional livestock density influences agent-level reproduction and mortality, so the decisions of one agent (e.g., increasing herd size) affect all others. Additionally, crop residues can be used as fodder for livestock, so each agent's crop production influences the amount that their livestock must be grazed on the rangeland. Third, agents interact indirectly through the job market; a finite number of jobs are distributed between the agents seeking employment, so a job taken by one agent is not available for another.

Collectives Agents do not form collectives.

Heterogeneity Agents are heterogeneous with respect to household size, land endowment, location, initial number of livestock, risk aversion, time required for collecting water and fuelwood, and their preference for wealth or leisure.

Stochasticity There is stochasticity in the initialization and assignment of agents to the landscape as well as in the calculation of crop yields and the allocation of regional livestock reproduction/-mortality and non-farm jobs between the agents.

Observation The primary output of interest is the monthly incidence of agent-level food insecurity, which emerges as a result of the various interacting components of the model. We also observe agent-level livestock herd sizes as a measure of wealth, as well as the agents' livelihood decisions.

3.2.2.3 Details

Implementation details The model is implemented in python 3. Source code can be made available upon request.

Inputs, calibration, and initialization See Appendix B.1 for input data sources. Agent parameters were drawn collectively from the LSMS to account for dependencies (i.e., from the joint rather than marginal distributions) (Berger and Schreinemachers, 2006). Due to the sample size of the LSMS (484 households in Amhara), surveyed households are duplicated in the model, but the overall population-level characteristics are preserved.

A number of parameters could not be determined directly from available data sources. We selected the values for these parameters using pattern-oriented modeling (Grimm et al., 2005); first, we combined several qualitative desired model characteristics with a variety of empirical distributions created from LSMS data (labor allocations, food security, livestock ownership, and subsistence levels). Second, we used a genetic algorithm to identify parameter sets for which the ABM outputs most closely reflected these empirical distributions and qualitative characteristics (section B.7.7). We chose to match the ABM with distributions rather than single values or ratios as we are interested in the distributional nature of outcomes.

Crop yields We model crop yields to be dependent on climate, planting date, soil properties, and fertilizer application. Each agent has the ability to directly influence their own yields by changing their planting date and their application of fertilizer. To account for the effects of precipitation, temperature, solar radiation, and planting date, we use a climate yield factor (CYF). A CYF is a process-based representation of the cumulative effect of water stress on crop yields. Developed using methodologies from the FAO's Irrigation and Drainage Papers Nos. 33 and 56 (FAO, 1984, 1998), CYFs have been used in previous ABMs (Wossen et al., 2014) and Ethiopia-based modeling exercises (Block et al., 2008). The calculation of the CYF simulates the infiltration of precipitation through the soil, soil moisture change in the soil column, and actual evapotranspiration at a daily time step. Critical values are taken for each crop growth stage to give an overall seasonal CYF. CYFs range between 0 and 1, where 1 represents no water-induced yield reductions and 0 complete crop failure, i.e., $Y^a_{i,f} = CYF_i * Y^p_{i,f} + \epsilon_i$, where $Y^a_{i,f}$ is the actual water-constrained yield at location i under management condition f, $Y^p_{i,f}$ is the theoretical maximum attainable yield with no water constraints, and ϵ_i is a random perturbation that we add when calculating yields each year. To honor the yield variability within a single year between different fields in the region, we model the CYFs and Y^p heterogeneously (hence the dependence on location, i). Section B.7.1 gives more details.

Off-farm labor The different agent preferences imply qualitatively different behavior with respect to off-farm labor; agents with preference for leisure will only engage in non-farm wage labor if they expect to be food insecure, while agents with preference for wealth will allocate all remaining labor to this activity. Given this, we assume that the "wealth maximizer" agents have access to a different kind of labor market that has guaranteed employment options. For the leisure maximizer agents, we assume that there is a constant, finite amount of non-farm work available each month. Jobs are allocated to these agents on a daily basis, with a probability dictated by the ratio of the total job availability to the total demand for off-farm work in a given month.

Crop prices To provide a proxy for temporal crop price dynamics, we use a regression model that predicts monthly crop prices in Amhara using regional climate variables (section B.7.2). In using this statistical model, we must assume that the relationships discerned from the historical data hold into the future and under drought conditions. An economic equilibrium model could be used to calculate prices under hypothetical drought conditions and reduce these assumptions, but this is beyond the scope of this study. In any case, historical prices alone would not represent drought conditions and our approach is more thorough than assuming a constant price throughout the simulation as is often done in agricultural ABMs.

Rangeland dynamics Livestock are an asset that can accumulate through reproduction. However, issues of overpopulation and drought can lead to mortality and population declines (Mogues, 2006; Crépin and Lindahl, 2009; Desta and Coppock, 2002). We model livestock population using a simple system dynamics model inspired by Janssen et al. (2000, 2004). This model operates at the regional level, with all livestock considered as a single object that shares the entire region's grazing land. It would be unnecessarily complicated for the purposes of this chapter to incorporate spatially explicit herding behavior. Each year, based on current population and climate conditions, grass growth is simulated and rainfall and overcrowding effects on livestock are calculated, giving a net regional livestock growth or decline. This is apportioned randomly between the agents. This simple representation provides a first approximation of these dynamics. Section B.7.3 gives the model equations.

3.2.3 Representing droughts

The effect of a drought on food security will depend on both the antecedent and succeeding climatic conditions. To account for this variability, we expose the system to droughts under multiple background "climate time series," which we created by repeatedly randomly sampling years from the 2000-2015 observational climate record.

We represent droughts through reductions in rainfall. For example, a 50% shock represents a year in which the rainfall at every pixel location in the modeled region is reduced by 50%. The rainfall reduction affects crop yield, grass growth, livestock mortality, and market prices. The effects on yield, due to the CYF representation, are both spatially explicit and non-linear in rainfall. To give a single metric to represent each climate shock, we do not vary the temperature in our drought scenarios. In all cases, we run the model for 30 years and test the effect of a single-year 50% drought in the 5$^{\text{th}}$ year of the simulation.

3.2.4 Conceptualizing climate resilience

We conceptualize a resilient system as one that has the capacity to maintain high levels of food security throughout and following a drought. At a regional level, we measure food security as the fraction of the population in each month that is food secure. To assess distributional effects, we examine how food security differs based on household characteristics and vulnerability. Food security is an appropriate outcome metric for our study for several reasons (Ansah et al., 2019). First, it is of relevance to development (e.g., refer to the United Nations' Sustainable Development Goals). Second, it is the outcome of a variety of interconnected social and environmental processes (e.g., agricultural production, market access, and employment opportunities). Third, while it has occasionally been questioned whether resilience is a desirable trait (e.g., resilient poverty systems) (Miller et al., 2010; Cretney, 2014), we argue that higher levels of food security are inherently desirable. Finally, with respect to equity, our measure of food insecurity identifies the most vulnerable households. Thus, higher levels of regional food security constitute improvements in food security for the more vulnerable (i.e., food insecure) households.

The temporal evolution of regional food security in the wake of a drought can be split into different dimensions (Figure 3.4). *Vulnerability* measures the maximum magnitude of the damage to food security. *Recovery* describes the time taken for the system to recover to some state, which we take here to be 90% recovery from the maximum damage. Additionally, one possible measure of *overall resilience* can be formed by calculating the area over the curve. Figure 3.4 is a stylized representation of a resilience curve; in reality (and in the model), food security will fluctuate both within and between years.

Although we consider resilience with respect to a state of functionality (food security) to which we desire the system to return, resilience is not conceptualized with respect to discrete equilibria and our modeling approach does not necessitate that the system must achieve this functionality through the same means. It is possible that the drought transitions the system into a new structure with either inferior, on-par, or better functionality. For example, large livestock losses due to drought-induced mortality may lead households to spend more time in off-farm labor, qualitatively

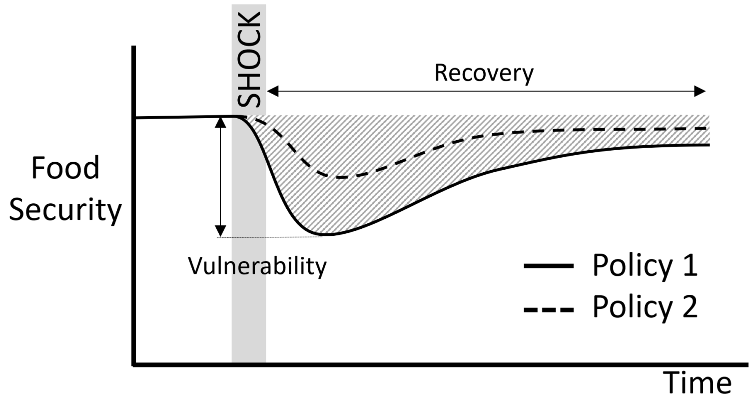

Figure 3.4: Schematic representation of the dimensions of resilience.

changing their livelihoods and patterns of food security.

We emphasize that the pre-shock level of food security does not represent universal food security; there still exists food insecure population in the absence of a shock. Hence, it is indeed possible and desirable for food security to increase above the "zero line". This highlights the potential for resilience assessments of this nature to overlook vulnerable populations that are always food insecure, motivating the consideration of distributional impacts.

3.2.5 Scenario analysis

We explore the resilience-enhancing potential of increased availability of non-farm jobs and provision of seasonal climate forecasts. Both of these are common measures for supporting rural producers and are in line with several of the options recently proposed by the Government of Ethiopia (Federal Democratic Republic of Ethiopia, 2019).

Non-farm job availability is already increasing in rural Ethiopia through infrastructure and development (Bachewe et al., 2016). Additionally, Ethiopia's Growth and Transformation Plan II calls for "industrial-led development," which can be expected to result in future increases. Shifts to non-farm employment can decrease smallholders' dependence on agriculture and, through extension, climate, potentially increasing climate resilience (Headey et al., 2014). Non-farm employment can also act as a source of income in the lean period of the year (Bachewe et al., 2016), potentially generating pro-poor effects (Van den Broeck et al., 2017; Herrmann, 2017). We represent increased non-farm job availability in our model by increasing the supply of non-farm jobs by 20%.

The ability for seasonal climate forecasts to have positive effects on smallholder communities has been extensively debated. Previous studies have highlighted, for example, the importance of implementation, communication, skill, and trust in the forecasts (Ziervogel et al., 2005; Narisi et al., 2018; Marshall et al., 2011; Luseno et al., 2003). However, if such barriers can be overcome, forecasts can allow smallholders to make better-informed agricultural management decisions (Risbey et al., 1999; Hansen et al., 2011). The provision of seasonal climate forecasts in our model gives agents perfect information about the time at which it is optimal to plant. Also, they provide information about the climate conditions for the upcoming year, which is of relevance for calculating crop yields, crop prices, and livestock dynamics. This information is not perfect, however, as it does not resolve the random perturbation that is added when calculating agent-level crop yields and does not give information about other agents' livestock decisions and the implications of these for rangeland dynamics. Agents have an evolving *trust* in the seasonal climate forecast, which is influenced by its accuracy relative to their other methods for predicting yields (section B.7.4). The extent to which the forecast affects agent decision-making is influenced by this trust as well as their risk aversion.

3.2.6 Output analysis

We analyze the effect of each intervention on resilience both at the regional- and household-levels. We assess the effect of each intervention (int) in two distinct ways. The first, which we refer to as the "shock-absorbing" effect, isolates the effect of a drought in the presence of the intervention:

$$R_{int,1} = FI(int, shock) - FI(int, no\_shock) \tag{3.1}$$

where R denotes resilience and FI denotes the sum of food insecurity over the population at a point in time. The second way, which we refer to as the "poverty-reducing" effect, combines the effect of the drought with the effect of the intervention relative to the baseline:

$$R_{int,2} = FI(int, shock) - FI(baseline, no\_shock) \tag{3.2}$$

By isolating the effect of the shock, the first measure offers a more accurate assessment of resilience as we have conceptualized it. However, it is more difficult to compare interventions, since the $FI(int, no\_shock)$ value for each intervention will be different. The second measure gives a more accurate idea of the overall effect of the intervention on the system under drought conditions. For example, an increase in job availability enables food insecure households to more regularly engage in non-farm wage labor, potentially reducing their chance of experiencing food insecurity both in regular, non-drought years and in the wake of a drought. Our first measure isolates the drought-related impact, while the second incorporates both. We present both measures of resilience in our results.

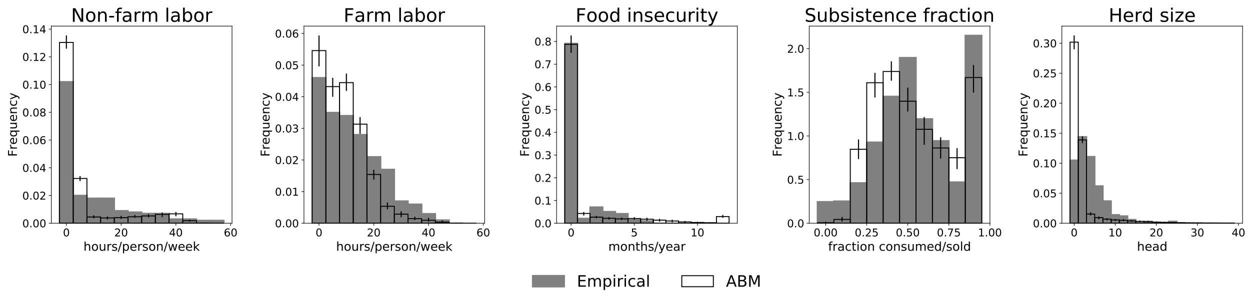

Figure 3.5: Comparison of empirically-derived and ABM-generated distributions. Whiskers show the 1% and 99% values over the ABM replications.

3.2.7 Simulation procedures

The two sources of variability relevant to our resilience assessment are (1) the climate time series and (2) stochasticity in the ABM itself. We used 50 different climate time series as it enables robust averaging yet is not overly computationally intensive. To adequately account for the stochasticity of the ABM we conducted a convergence analysis, which determined that 50 replications were sufficient to achieve a relative error of 0.1 with a confidence level of 90% (section B.7.8). This therefore results in a total of 2,500 simulation replications for each experimental setting. We use common random numbers to reduce the variability in the outputs when comparing between policies. Finally, to reduce the impact of initialization (e.g., of agent beliefs) on our analysis, we conducted a 13-year (2003-15) burn-in period before beginning each model run.

3.3 Results and discussion

After presenting the calibration results (section 3.3.1), we begin by assessing the effect of a drought under baseline conditions (3.3.2). Then we evaluate the effects of the RESs (3.3.3), explore the sensitivity of our results (3.3.4), and discuss limitations (3.3.5). In all cases, we focus primarily on food security as our outcome of interest. Although this is a purely social outcome, it emerges as a result of interactions within the different components of the ABM.

3.3.1 Calibration and cross-validation

Overall, the parameterized ABM reproduces the empirical distributions of several household-level variables (Figure 3.5). However, the ABM overestimates the proportion of households without any livestock. Potential reasons for this discrepancy and a cross-validation of the fitting metrics used for the calibration procedure are discussed in Appendix B.4.

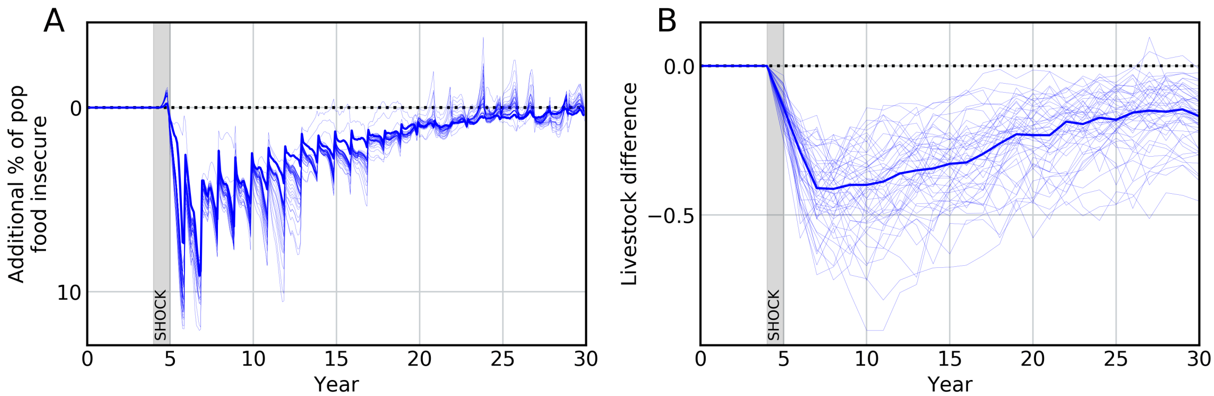

Figure 3.6: Effects of a 50% drought on: (A) regional food security, measured as the *additional* percentage of the population that is food insecure; and (B) agent-level livestock holdings (head). Thin lines represent the average effect in each of the 50 climate time series. Bold lines represent the overall median effect.

3.3.2 The effect of a drought

3.3.2.1 Regional effects

The largest effect of the drought on regional food security is experienced in the year following the drought (Figure 3.6A); this is because the agents that have received only small crop yields at the end of the drought year begin the following year with insufficient food stored to last them until the next harvest. Because the agents and the effects of the drought are heterogeneous, different agents fall into food insecurity at different times, and hence the line in Figure 3.6A is sloping downwards throughout year 5-6. This seasonality also exists in the subsequent years, with progressively more agents over time able to replenish their assets sufficiently to restore levels of food security.

Drought-related livestock mortality occurs during the drought year, but the largest effect on livestock is not felt until several years later (Figure 3.6B). This is a result of the food shortages described above; many agents must destock their livestock herds as a coping measure in the years following the drought. Livestock herds slowly recover (through both natural reproduction and purchase from the market), but over the simulation period analyzed do not on average fully recover to the levels of the drought-free counterfactual.

Together, these results demonstrate that the effects of drought on smallholder livelihoods persist for many years. Agents may be able to cope with the immediate effects, but in doing so they deplete their capital, causing increased levels of vulnerability even under regular conditions in the following years.

3.3.2.2 Distributional effects

Household characteristics play an important role in shaping the drought experience of the agents. The most vulnerable agents (i.e., both most likely to be affected and slow to recover) include those with high subsistence fractions, small land holdings, and small livestock herd sizes (Figure 3.7). These results are generally consistent with other discussions of food insecurity in smallholder populations (Devereux and Sussex, 2000; Bogale et al., 2005; Beyene and Muche, 2010).

However, our results show a few counter-intuitive trends. First, although agents with large land holdings are faster to recover, they are just as likely to be affected by the drought as those with small land holdings (Figure 3.7). In part, this is because of the dramatic effects of the drought on crop yield; agents with large land holdings rely heavily on agricultural production and hence are strongly affected by the drought. However, this trend also can be explained as an artefact of the model structure; in the model, it is impossible for the agents to farm only a fraction of their land, so agents with large land endowments have large farming-based labor requirements. As the model is parameterized, these labor requirements are sufficiently high that these agents in some cases cannot farm their land. Their outcomes are consequently similar to those with very little land; that is, high levels of food insecurity.

Another non-linear pattern is seen in household size; both small and large households are most affected by the drought (Figure 3.7A). As household size increases in the model, so do both labor capacity and consumption requirements. Our results suggest that small households are constrained by their small labor capacity (because, for example, they are less able to engage in non-farm labor to help mitigate the effects of the drought) and large households are constrained by their higher consumption requirements, both of which can increase drought vulnerability. Larger households, however, generally can recover more quickly from the drought (Figure 3.7B), suggesting that a larger labor endowment is beneficial during the recovery period.

Together, these results have three implications. First, the strongly distributional nature of the effect of drought on food security suggests that measures to increase agricultural productivity (through crop production and livestock holdings) could offer promise to substantially reduce smallholder drought vulnerability and enhance recovery. From an assessment perspective, the relationships between household characteristics and food insecurity are not necessarily linear. Although this is in part an artefact of our model's structure, model-based and empirical studies alike should nonetheless acknowledge the potential for non-linear relationships in the data. Finally, a high propensity for being affected (Figure 3.7A) need not imply a long recovery time (Figure 3.7B). Different priorities may prevail in different contexts, and this result demonstrates that a single metric alone cannot fully represent resilience.

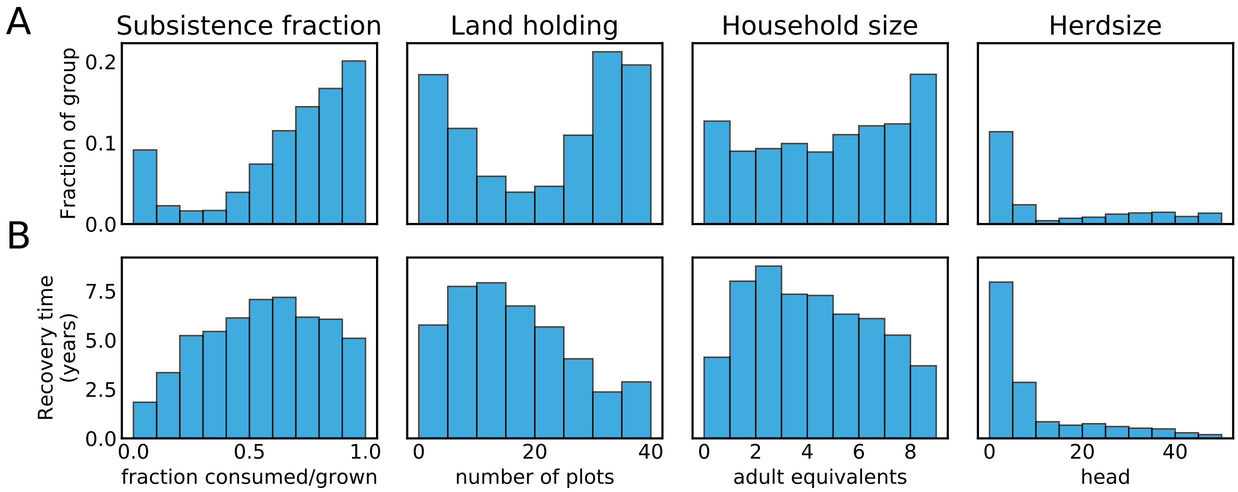

Figure 3.7: Fraction of each population group that's food insecurity is higher at the time of the maximum drought impact (A) and the time taken for 90% of each group to recover (B).

3.3.3 Assessing and comparing interventions

3.3.3.1 Regional resilience: measurement method matters

When isolating the "shock-absorbing" effect of each RES (i.e., excluding their potential for food insecurity reduction relative to baseline conditions; Equation 3.1), both strategies provide mixed levels of benefit relative to the baseline (Figure 3.8A); both strategies do little to improve food security in the year following the drought, but provide a mild benefit in the subsequent years. Increased job availability even slightly reduces food security in the long term[3]. When both strategies are implemented together (Figure 3.8A, right), there is also little benefit in the year following the drought, but almost complete recovery occurs within three years, demonstrating an enhanced benefit to having the strategies in conjunction.

A different picture emerges when the overall "poverty-reducing" effect relative to the baseline is included in the assessment (Figure 3.8B; Equation 3.2). Here, both RESs reduce the vulnerability (i.e., the maximum impact of the drought), with the climate forecasts providing a slightly greater benefit. This is because climate forecasts enable agents to engage in *ex-ante* coping measures; in the year of the drought, a portion of the agents anticipate the effects and choose not to farm (Figure B.4), thus separating themselves from the drought-induced yield effects. Despite smaller vulnerability-reduction benefits, the increased job availability scenario provides a faster return to the counterfactual levels of food insecurity (in year 8 of the simulation); this is because the higher availability of non-farm wage labor, which facilitates benefits throughout the year, reduces

[3]This effect is because in the job availability scenario, the shock induces a permanent reduction in the number of agents farming (Figure B.4). Because of the uncertainty in the non-farm wage labor market, higher reliance on this contributes to increased levels of food insecurity for these agents.

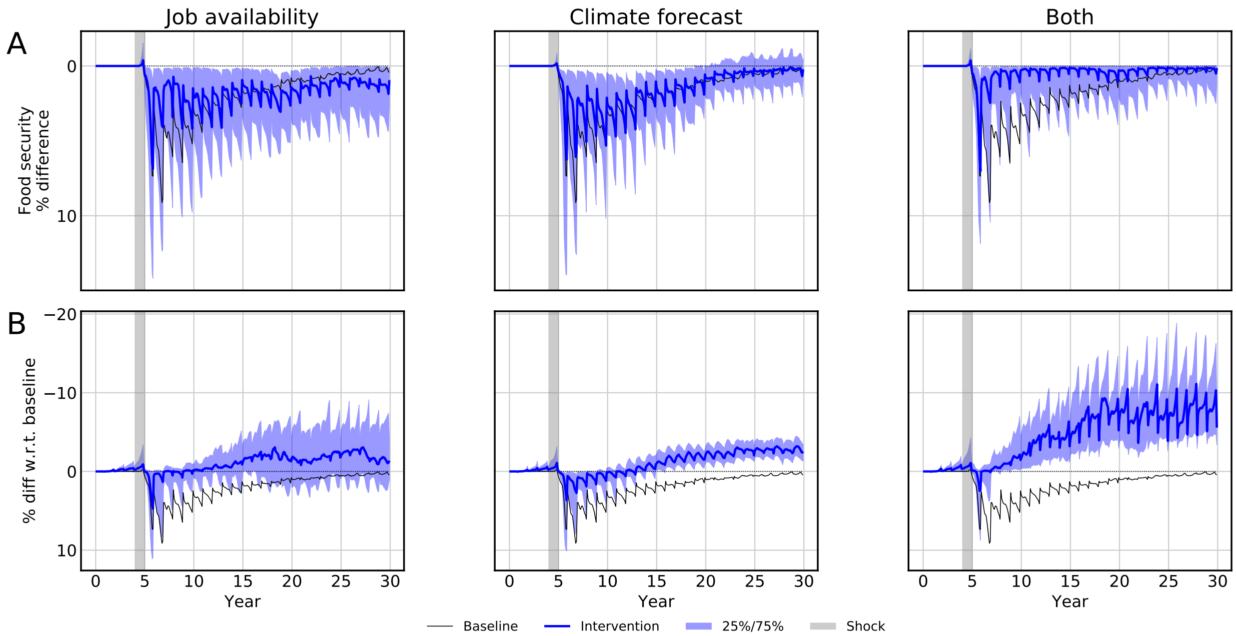

Figure 3.8: Effects of the interventions on the resilience to a 50% reduction in rainfall in comparison to the baseline. "A" isolates the "shock-absorbing" effect of each scenario (Equation 3.1) and "B" includes the food security improvements ("poverty-reducing benefits") relative to the baseline (Equation 3.2).

the extent to which agents must sell livestock as a coping mechanism (Figure B.5, Appendix B.3). When both RESs are included, the drought only leads to increased food insecurity for one year, beyond which food insecurity is lower than the $\{baseline, no\_shock\}$ simulation. Additional analysis revealed that having both strategies together generally provided synergistic benefits to overall resilience — i.e., larger benefits than the sum of the two RESs in isolation (Appendix B.5). In general, the benefits provided by the strategies to the different components of resilience are similar for different magnitudes of drought (Figure B.3, Appendix B.3).

Contrasting the resilience assessment of Figure 3.8B with that in 3.8A demonstrates that the interventions enable a transformation of the state of food security in the modeled system; even when affected by a 50% drought, the interventions enable better food security than the no-drought baseline scenario (Figure 3.8B). In Figure 3.8A (Equation 3.1) this development transition is not measured.

3.3.3.2 Effects on inequality

Using our measure of food security, we calculated the evolution of the Gini coefficient, a common measure of inequality (Figure 3.9). Our results show that, in all cases, the drought increases inequality; by affecting only a portion of the households (Figure 3.7), the drought causes food security to be less equally shared between the population, thus increasing the Gini coefficient. Ad-

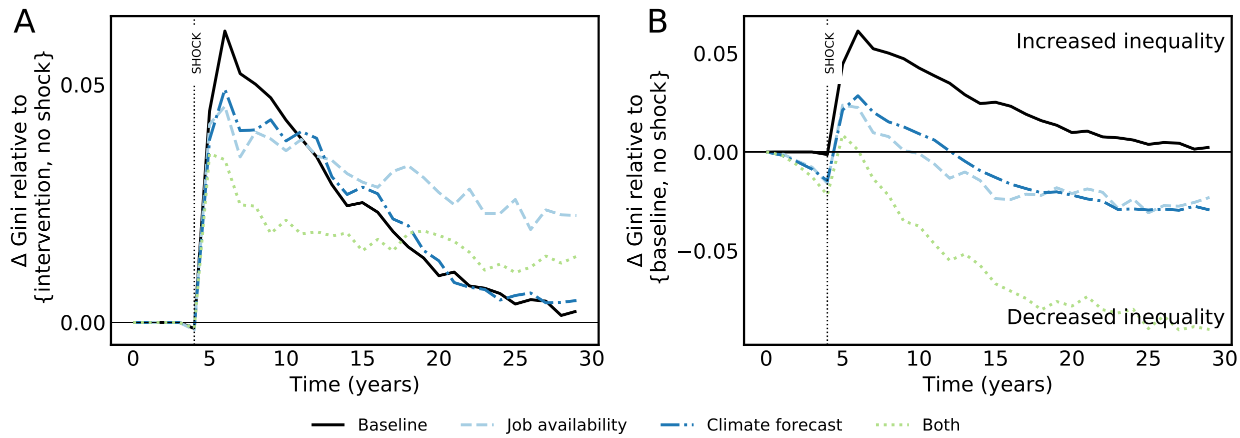

Figure 3.9: Effect of drought and interventions on the Gini coefficient, calculated with respect to food security, over the simulation. "A" isolates the "shock-absorbing" effect (Equation 3.1) and "B" includes the "poverty-reduction" relative to the baseline (Equation 3.2).

ditionally, in all cases, the interventions reduce the impact of the drought on inequality; an overall increase in population-level food security necessitates that the cumulative food security is more equally shared between the population, thus decreasing the Gini coefficient. When assessing the overall effects relative to the baseline (Figure 3.9B), all interventions result in long-term inequality reductions. These results demonstrate that enhancements to resilience as conceptualized here go hand in hand with reductions in population-level inequality.

The form of resilience measurement influences the assessment of the relative benefits of the two interventions on longer-term inequality; when isolating the shock-absorbing effect (Figure 3.9A), the job availability scenario results in a long-term *increase* in inequality relative to the climate forecast scenario, whereas it leads to similar long-term results when the poverty-reduction effect is included (Figure 3.9B). This is largely because, as already mentioned, food security never recovers following the shock in the job availability scenario (Figure 3.8A).

3.3.3.3 Distribution of effects based on vulnerability

The population-level analyses presented above do not tell us whether all groups receive equal benefit. For example, even with a decrease in overall inequality, it is possible that the benefit for some comes as the expense of others. Additionally, recall that the "zero" line in Figures 3.6 and 3.8 does not indicate that there is no food insecurity, rather that food insecurity is no worse than it would have been had a shock not occurred. By this metric, those who are food insecure both with and without the climate shock are overlooked. The distributional analysis we present here speaks to some of these issues.

Assessing the effect of the drought on agent-level food insecurity under each intervention relative to each agent's baseline state reveals that the largest benefits under all scenarios are experienced by those that were moderately vulnerable under baseline conditions (i.e., 20-40% chance of being affected by the drought with no intervention) (Figure 3.10A; Equation 3.1). The benefits these agents experience are slightly stronger under the climate forecast scenario. Under all scenarios, however, the more vulnerable agents (60% chance) are *worse* off. This suggests that, when assessing solely the shock-absorbing effects of the interventions, a resilience benefit for some comes at the expense of the more vulnerable.

Similar to before, including the poverty-reducing effects relative to the baseline gives rise to a substantial difference in results (Figure 3.10B); under this measure, food security is improved for everyone, with the agents that are most vulnerable under the baseline conditions receiving the greatest overall benefit (i.e., the greatest vertical deviation from the 1:1 line). Although both strategies improve outcomes for all agents, climate forecasts more strongly benefit the moderately vulnerable agents than increased job availability, whereas increased job availability more strongly benefits the most vulnerable agents.

Together, these results suggest that there may be differences between interventions that provide pro-poor benefits relative to existing vulnerability (Figure 3.10B) in comparison to pro-poor benefits during and following a drought (Figure 3.10A). Given the prevalence of climate shocks in smallholder systems, assessing drought recovery is critically important. Thus, a focus purely on poverty reduction may be insufficient. To mitigate the potential for inequitable development, future studies should therefore consider both poverty reduction and drought recovery in their assessments.

3.3.3.4 Distribution of effects based on household characteristics

Disaggregating the effects of the interventions by household-level characteristics (Figure 3.11) reveals a similar story, which we summarize in Table 3.1; both interventions *increase* the effect of drought on food security in some of the most vulnerable population groups when the effect of the shock is isolated (Figure 3.11B), but when food security improvements relative to the baseline are considered, both interventions provide the largest benefits to the most vulnerable populations — i.e., they are equitable (Figure 3.11C). Increased job availability provides slightly larger poverty-reduction benefits to the most vulnerable groups — particularly those with few land and livestock resources[4] (Figure 3.11C). This is likely because an increased job availability provides an additional source of income, whereas climate forecasts enable better management of land and livestock resources — resources that the more vulnerable agents do not possess.

In any case, it is noteworthy that both interventions provide similar levels of benefit to the agent population. We do not model the mechanisms through which these system changes could be achieved, but note that increased non-farm job availability in a country like Ethiopia could only be

[4]Though note that the outcomes for the agents with a large number of plots in the model are similar to those with few plots. See section 3.3.2.2.

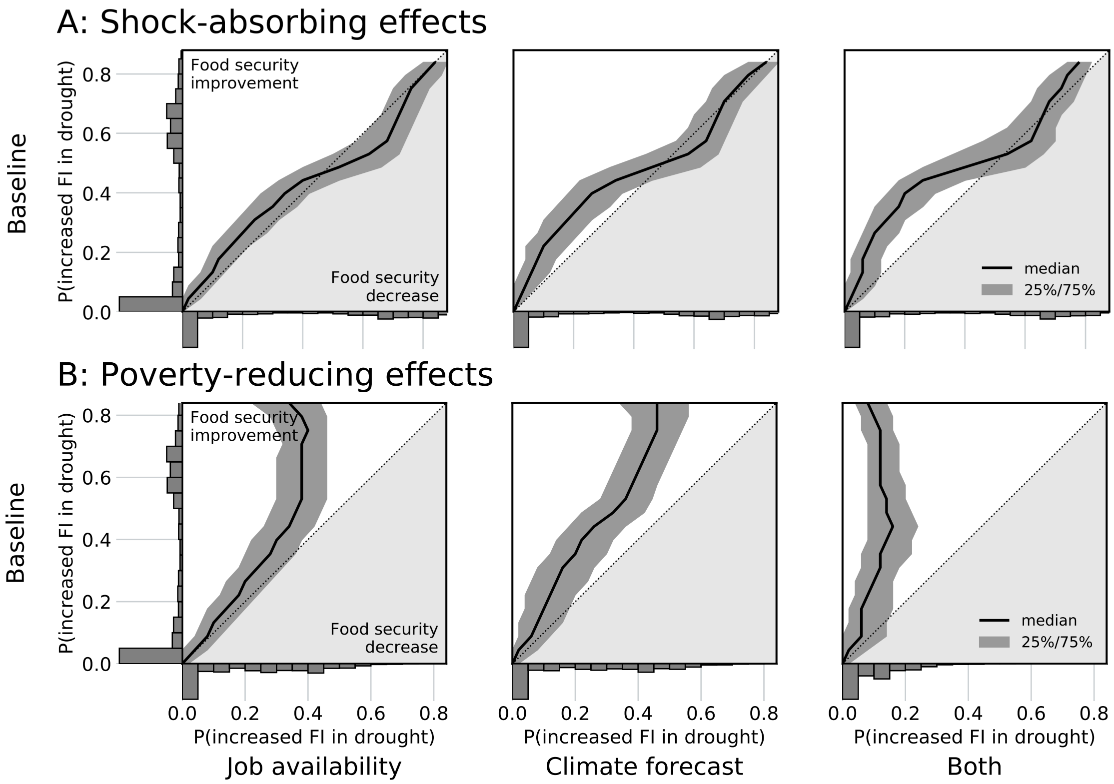

Figure 3.10: Distributional effects of interventions on agent-level food insecurity (FI). The axes represent the probability that an agent's food insecurity is higher in the drought simulation, calculated over all replications of the ABM. The interventions are assessed with respect to their shock-absorbing effect (A; Equation 3.1) and including their poverty-reducing effects (B; Equation 3.2). Values to the left of the 1:1 line indicate food security improvements brought about by the intervention. For example, a point $(x, y) = (0.1, 0.2)$ represents an agent with a 20% chance of being affected by the drought under baseline conditions and 10% with the intervention.

Table 3.1: Summary of distributional results in Figures 3.10 and 3.11. Differences between the interventions under each resilience measurement approach are italicized.

	Shock-absorbing	**Poverty-reduction**
Job availability	Inequitable • Benefit moderately vulnerable • Increase vulnerability for most vulnerable	Equitable • Benefit all households • Stronger benefit for *most* vulnerable
Climate forecast	Inequitable • *More strongly* benefit moderately vulnerable • Increase vulnerability for most vulnerable	Equitable • Benefit all households • Stronger benefit for *moderately* vulnerable

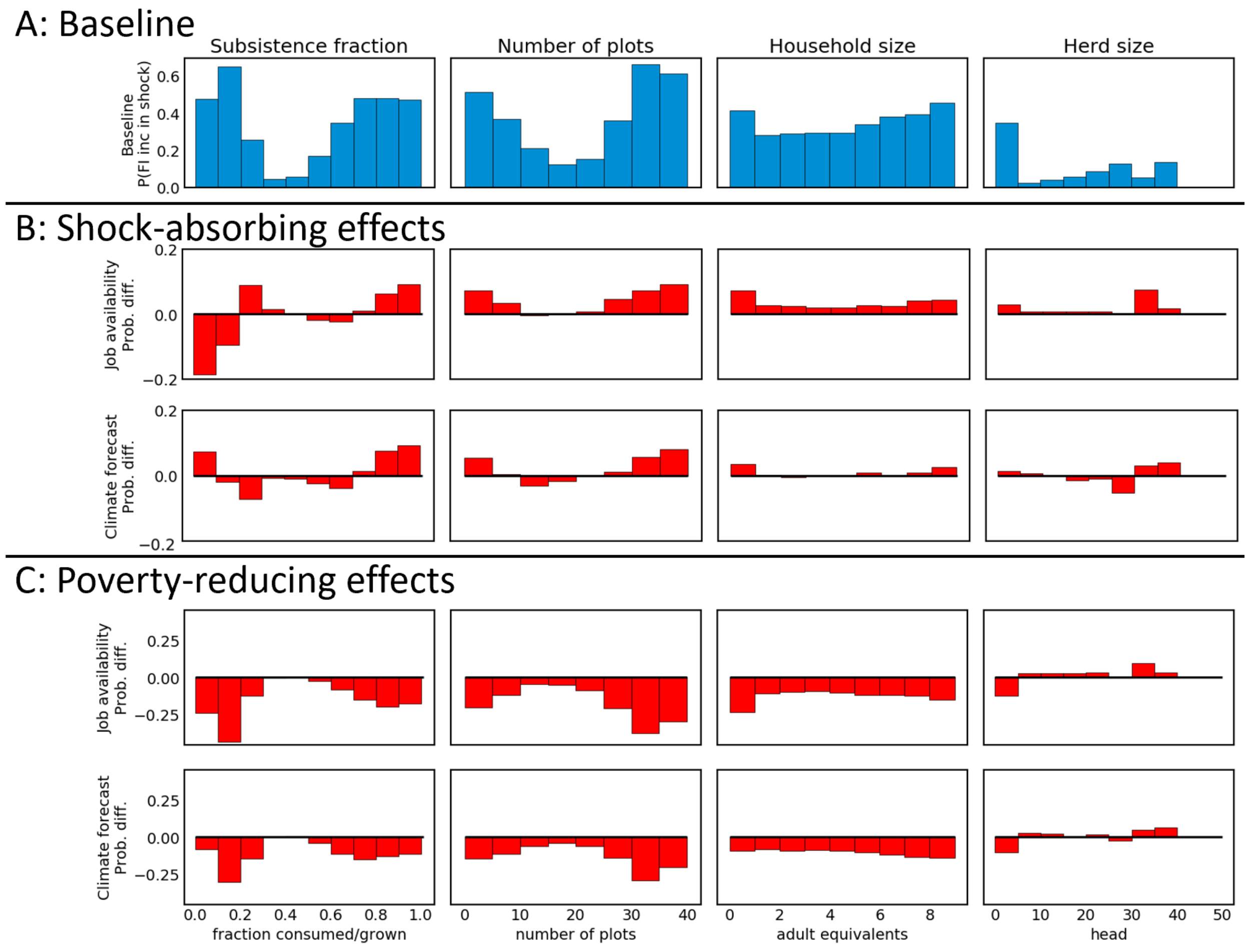

Figure 3.11: Probability of the drought increasing total food insecurity as a function of agent characteristics (A) and the impact of interventions on this when isolating the shock effects under each intervention (B) and when including food security improvements relative to the baseline (C).

achieved through significant infrastructural development, which would entail large costs. Climate forecasts, conversely, are limited both by their practical ranges of accuracy, information dissemination, and farmer trust (Hansen et al., 2011). In different contexts, different barriers may be more or less important, favoring one type of strategy over another. Given that our results suggest that the strategies offer similar resilience-enhancing potentials, practicality may be the most critical factor in selecting between these options. In either case, however, our results underscore the importance of ensuring that vulnerable populations benefit in the wake of a drought.

3.3.4 Sensitivity of results

We conducted a one-way sensitivity analysis on six uncertain model parameters. In most cases, the parameter changes had little influence on the intervention comparison and the distribution of food insecurity throughout the population. However, there were some interesting effects. For consideration of space, we present and discuss these results in Appendix B.6.

3.3.5 Limitations and extensions

These results emerged from our ABM. As such, they should be considered in light of the model representation and may change upon inclusion or exclusion of processes, or under alternative formulations of the modeled processes. Our sensitivity analysis does not assess the influence of model structural uncertainty on our results. Here we mention several important considerations that may affect the generalizability of our results to different contexts or real smallholder systems.

With no data with which to calibrate nutrient effects on crop yields, we did not model soil parameters (e.g., organic matter or nutrient levels). In reality, soil erosion and quality are limiting factors in Choke Mountain (Simane et al., 2013) and in many smallholder agricultural systems (Sanchez et al., 1997). These exclusions may present trade-offs with resilience as defined in this study and future research could incorporate environmental feedbacks in the farming system.

The decision-making framework and methodology is a critical component of any agent-based model. Although our approach was informed by generalized empirical evidence, context-specific information gathered through surveys or interviews would provide further knowledge about how households make decisions (e.g., how households might incorporate forecast information into their decision-making). Alternatively, if no such information is available, it would be possible to explore the sensitivity of the conclusions to changes in the assumed decision-making framework (Schlüter et al., 2017).

Future efforts could work to incorporate measures of dietary diversity into agent-based models or to model issues of nutrient utilization (Nicholson et al., 2019). Representation of multiple crops would also allow agents to diversify their production, which is a common strategy of smallholder farmers (Teklewold et al., 2013; Barrett et al., 2001), and would most likely lead to increased resilience under the same conditions.

We assumed that job availability is constant over time, livestock can always be bought and sold from the market, and crop prices are influenced solely by climate. Especially in drought years, these assumptions may not be appropriate. Additionally, we did not incorporate temperature variations in drought years, which would further affect our calculated crop prices. Efforts to endogenize prices by integrating spatially-explicit ABMs with partial or general equilibrium models are welcomed and should be prioritized in future research agendas (van Wijk, 2014; Berger and Troost, 2014).

Finally, our parameterized ABM failed to accurately recreate the empirical distribution for livestock herd size. Livestock is a critical component of smallholder farming systems and important for food security and the specific policy implications of our analysis could be biased by this important limitation. First, it is possible that our results overestimate the drought vulnerability of the smallholder system; we overestimate the proportion of households with no livestock relative to the LSMS data, therefore underestimating the coping capacity. In the policy comparison, it is possible that this leads to an underestimation in the benefit of the climate forecast scenario; the climate information provided to the agents can be used to inform livestock selling decisions, which more households would benefit from in reality than do in the model. Future work can improve the process description in the model to increase the policy-relevance of results.

3.4 Conclusions

We developed and analyzed results from an ABM to quantify drought resilience in a stylized Ethiopian smallholder farming system and explore the effects of selected policy interventions, including the provision of seasonal climate forecasts and increases in non-farm job availability. We quantify resilience as the extent to which a measure of household food security is affected by a drought, combined with the time taken to recover. We pay particular attention to equity in the distribution of the effects of both shocks and interventions.

Our analysis has produced the following insights. First, considering the effect of drought on the system, many households could cope with the immediate effects, but eroded their assets in doing so, making them vulnerable even under normal climate variability in the following years. The households most affected and taking the longest to recover included those with small land holdings, small household sizes, and small livestock herds. Methods that increase the absorptive and adaptive capacities of these households could reduce the extent to which assets are lost and improve the speed of recovery.

Second, both of the strategies that we assessed generated similar effects to resilience at both regional and household levels. This is noteworthy given the qualitatively different mechanisms through which each strategy acts. These results suggest that, in the case of these interventions, practical constraints with respect to implementation may be more important considerations than the different levels of resilience enhancement provided. In particular, if reliable, trustworthy forecasts are available, climate resilience could be increased without significant capital or infrastructure investment. Further, the potential for synergistic benefits emphasizes the utility of considering agricultural interventions as interacting, rather than in isolation.

Third, and most importantly for this study, when isolating the shock-absorbing effects of the strategies, both strategies led to *increased* levels of vulnerability for the more vulnerable households — i.e., they were inequitable (Figures 3.10A and 3.11B). However, if the strategies' poverty-reducing benefits relative to the baseline were also included, both strategies benefited all households (Figures 3.10B and 3.11C). Here, increased job availability provided stronger poverty reduction benefits for the most vulnerable households, while climate forecasts provided larger benefits for the moderately vulnerable households (Figure 3.10B). These results emphasize the need to jointly consider both poverty reduction and shock vulnerability when assessing the effects of development interventions; in our case, this distinction led to discordant results. Thus, it is not necessarily a question of *if* an intervention produces inequitable effects, rather *in which ways* an intervention might be inequitable. Future studies should therefore clearly articulate their resilience measurement approach and consider the implications that this might have.

In conclusion, the approach we have demonstrated in this chapter enables disparate development strategies to be analytically compared under a common framework with a focus on the potential for inequities in their effects. In the pursuit of sustainable development, we argue that quantitative policy analyses should consider how benefits are shared across the population and prioritize development that fosters resilience not just on average, but for everyone.

Acknowledgments

This project was supported by a U.S. National Science Foundation (NSF) INFEWS Grant (1639214) and an NSF CNH Grant (DEB-1617364). In addition to the helpful comments from the anonymous reviewers, the authors would like to thank Allison Kelly for her help with the LSMS data, Uriah Israel for work on the genetic algorithm, Ying Zhang for the climate yield factors, and Ben Zaitchik for providing the climate data and general climate-related guidance.

Chapter 4

Assessing Model Equifinality[1]

Equifinality—a situation in which multiple plausible explanations exist for a single outcome—presents a challenge for socio-environmental systems modeling. When equifinality is ignored in model calibration, subsequent policy analyses may mis-estimate the range of potential policy effects. In this chapter, we present and demonstrate an approach, called DMC-RPA, for generating a set of diverse model calibrations (DMC) to enable more robust policy analysis (RPA). The optimization-based approach maximizes diversity in the model parameters and/or structural configurations to efficiently represent any equifinality in the model set. We demonstrate the approach for an agent-based model that is used to compare resilience-enhancing strategies in a smallholder farming system. Results over the set of diverse model calibrations demonstrate consistent policy effects, enabling stronger conclusions than a single model analysis. Going forward, this approach can be applied in the development of socio-environmental systems models to facilitate more robust policy analysis and inference.

4.1 Introduction

Process-based models are regularly used for ex-ante evaluation of policy interventions (Verburg et al., 2016; Schulze et al., 2017; Kremmydas et al., 2018; Schmolke et al., 2010). However, model outputs—and hence any policy recommendations derived from model analysis—are dependent on both the chosen model structure and parameterization (van Vliet et al., 2016). In particular, the complexity of socio-environmental systems (SES) and the corresponding uncertainty inherent in modeling them (Hornberger and Spear, 1981; Liu et al., 2007; Ligmann-Zielinska et al., 2014) means that there may exist *multiple* plausible model configurations that reasonably fit observed

[1]Published as Williams T.G., Guikema S.D., Brown D.G, Agrawal A. “Assessing model equifinality for robust policy analysis in complex socio-environmental systems” *Environmental Modelling & Software* `http://dx.doi.org/10.1016/j.envsoft.2020.104831`

outcomes (Axtell and Epstein, 1994; Oreskes et al., 1994; Beven, 2006). This is known as "equifinality." Under this condition, it is possible that a single model's behavior does not represent the full range of plausible outcomes, leading to biased policy recommendations. This can have large implications when adaptation actions are limited in their reversibility or are costly to counteract, potentially locking systems into maladaptive pathways (Leclère et al., 2014).

However, the vast majority of SES modeling studies employ single, "best-fit" models to conduct policy analysis (Parker et al., 2003; Huber et al., 2018; Brown et al., 2013; O'Sullivan et al., 2016). To address equifinality in the development of SES models and prioritize robust policies for sustainable development therefore requires that two questions be answered:

1. Do multiple plausible structural and/or parameter representations exist for a given SES model? If so, how do they vary?

2. When applying a set of plausible models to new conditions (e.g., a policy analysis), do they lead to qualitatively consistent results? If not, what can we learn from this?

To attend to these questions, we present an optimization-based approach for identifying equifinal models and exploring their implications for policy analysis. We name the approach DMC-RPA: Diverse Model Calibration for Robust Policy Analysis. First, we systematically identify multiple parameter sets and/or model structural characteristics that each match calibration data within a specified level of fitness, yet are as diverse as possible (DMC) (Brill et al., 1982; Zechman and Ranjithan, 2004). Next, we conduct a policy analysis and explore the consistency of policy effects over the equifinal model set (RPA). There are two main contributions in this approach. First, by explicitly maximizing diversity within the model set, our approach enables an efficient representation of equifinality in a small number of models. This assists in the communication and understanding of equifinal models, as well as reduces the computational complexity of subsequent model experimentation. Second, we focus on the *implications of equifinality* for policy analysis, which allows for more robust policy assessments and inference in SES modeling.

We demonstrate our approach using a case study, in which an agent-based model (ABM) is used to compare strategies for enhancing climate resilience of smallholder farmers in an Ethiopian context. We measure climate resilience with respect to the effect of drought on household food security and examine which of two policy interventions provides the greatest resilience-enhancing benefit. Using the DMC-RPA approach, we identify a set of plausible, diverse model configurations and explore the implications for policy recommendations. We seek to examine what identifying equifinal models can mean for inferences drawn throughout the modeling process, and what implications it might have for model-based policy studies.

4.2 Background: Equifinality in socio-environmental systems models

Equifinality describes a situation in which a given set of observed patterns or outcomes can be produced by multiple distinct explanations (Beven and Freer, 2001). This is equivalent to the terms "nonidentifiability," "nonuniqueness," "multi-realizability," and the "parameter identification problem," in which multiple parameterizations or generative process descriptions are observationally equivalent (Oreskes et al., 1994; Conte and Paolucci, 2014; Walter, 2014; Guillaume et al., 2019). A particularly strong case of equifinality is known as structural nonidentifiability, which is most easily imagined in the case of a fully parametric equation-based model; if the model has more free parameters than the number of data points used to calibrate it, an infinite number of observationally equivalent parameterizations may exist (Schmidt et al., 2020).

Equifinality is relevant to SES modeling. Many SES models are highly complicated (Sun et al., 2016; Lee et al., 2015); that is, they contain a large number of parameters and structural assumptions. In many cases, there is limited knowledge of the processes driving the modeled system (Ligmann-Zielinska et al., 2014), as well as limited empirical data against which to compare model outputs (Augusiak et al., 2014). Given this potential mismatch between the dimensionality of the model and the data, multiple structures and/or parameterizations may exist that generate outputs consistent with the data. In other words, "the mapping from micro-rules to macro-structures may be many-to-one" (Axtell and Epstein, 1994). Due to stochasticity, feedbacks, and the non-analytical nature (i.e., no fixed, structural form) of many SES models (Windrum et al., 2007), such equifinality may be difficult to identify. Yet, due to this very nature, SES models can exhibit high degrees of path dependence and non-linear dynamics, which may lead to significant implications if the potential for equifinality is not acknowledged.

Equifinality is most pertinent to the calibration stage in the iterative model development cycle (Grimm and Railsback, 2005). The purpose of model calibration is to improve a model's fit to real-world conditions by adjusting its parameter values and/or structural representations (van Vliet et al., 2016; National Research Council, 2012). Typically, this involves finding the single, "best-fit" model to utilize for subsequent analysis. Proponents of the equifinality thesis argue that the possibility for *multiple* acceptable models should not be rejected and the model calibration process should instead constitute "a mapping of the landscape into a space of feasible models" (Beven, 2006).

There exist various approaches for model calibration and analysis that either explicitly or implicitly acknowledge equifinality. Table 4.1 gives a non-exhaustive overview. Approaches explicitly dealing with equifinality have been most extensively discussed and developed in the field of hydrology (e.g., (Beven and Freer, 2001; Blazkova and Beven, 2009; Smith et al., 2008; Efstratiadis and Koutsoyiannis, 2010; Yen et al., 2014; Vrugt et al., 2008)). In these contexts, the objective of allowing for multiple model configurations is typically to produce a wider uncertainty band on model predictions that is more likely to contain the true value. In contrast, acknowledgement of equifinality is surprisingly absent in process-based model evaluations of policy impacts in SES. This literature could therefore benefit from an approach that builds from frequently used methods in SES modeling to (1) identify equifinality and (2) explore its implications for policy analysis.

4.3 Approach: Diverse Model Calibration for Robust Policy Analysis (DMC-RPA)

4.3.1 Overview

We propose an approach for incorporating equifinality into the model development cycle to enable more robust policy analysis in socio-environmental systems (Figure 4.1; Figure 4.2; (Grimm and Railsback, 2005; Latombe et al., 2011; Schmolke et al., 2010)). The approach consists of two steps, which together we refer to as DMC-RPA: Diverse Model Calibration for Robust Policy Analysis. First, an optimization-based model calibration procedure seeks to identify the maximally diverse set of model configurations that can explain the observed data (diverse model calibration; DMC). Second, this small set of maximally different models is applied to a policy analysis (robust policy analysis; RPA). If policy recommendations are qualitatively different over the set of diverse, plausible model configurations, this is evidence to suggest that these policies may not be robust in reality or that further information is needed to reduce equifinality. Conversely, if results are consistent, this provides strength to any model-driven inferences and policy recommendations beyond a best-fit model analysis. In either case, our approach enables a more robust policy analysis[2].

The motivation underlying our approach is that, given the complexity of SES and the paucity of empirical data in many situations, we cannot claim to have all potentially relevant data for model calibration—i.e., there are inherently objectives that are *unmeasured* in the calibration process. Thus, it is not instructive to focus only on the single, "optimal" calibration to this imperfect set of data. Rather, it is more useful to generate a number of calibrated solutions that each "perform well with respect to modeled issues, and are significantly different with respect to the decisions they specify" (Brill et al., 1982). These diverse solutions may consequently behave considerably differently under conditions not modeled in the calibration process, such as a policy analysis.

[2] We note the distinction in the use of the word "robust" between our approach and Robust Decision Making (RDM); in RDM, a *robust policy* is one that is beneficial over a wide range of uncertain exogenous conditions (e.g., input data, future trajectories) (Kasprzyk et al., 2013), whereas we refer to a *robust policy* as one that is beneficial over a range of equifinal model configurations. The more general phrase *robust policy analysis (RPA)* refers to our overall approach, whether the policy itself is robust or not.

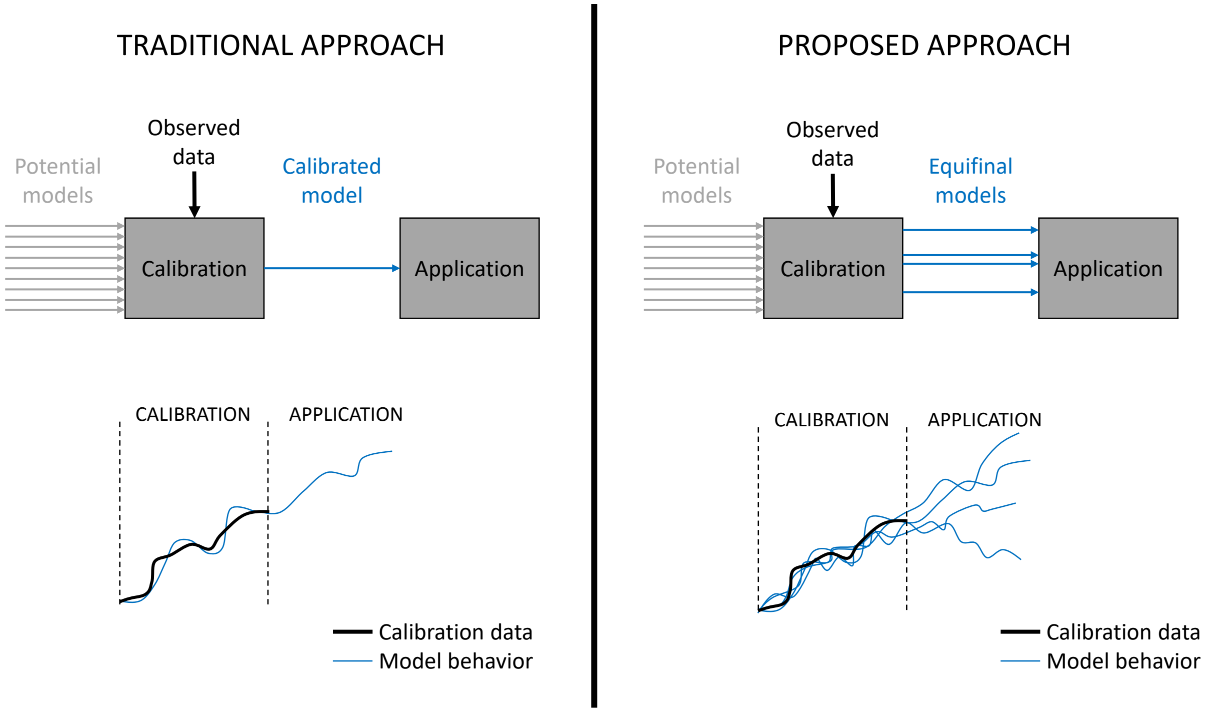

Figure 4.1: Our approach identifies a set of models that each similarly match the calibration data, yet are as different from each other as possible with respect to the parameters and/or structures that they specify. As a result, the models may produce qualitatively different behavior when applied to new conditions (e.g., in a policy analysis).

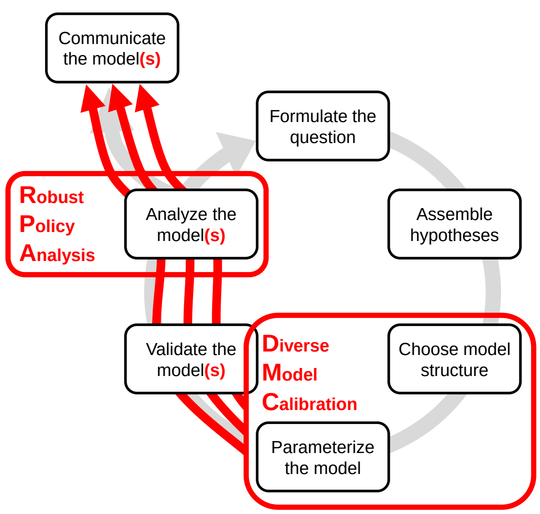

Figure 4.2: DMC-RPA in the model development cycle. A set of diverse model calibrations are identified and then used to conduct a policy analysis.

Table 4.1: Other approaches relevant to assessing equifinality in process-based modeling

Name	Stage in modeling cycle	Description	Relation to equifinality
Pattern-oriented modeling (POM)	Calibration	Identify the model configurations that can simultaneously produce several desirable "patterns" that are consistent with reality (Wiegand et al., 2003; Grimm et al., 2005). POM is similar to the "limits of acceptability" approach (Beven, 2006), which seeks to find "behavioral" models that provide predictions within specified limits.	Multiple model configurations may generate the patterns or be behavioral, thus allowing for equifinality.
Bayesian model calibration	Calibration	Use Markov chain Monte Carlo (MCMC) simulation to derive distributional estimates of model parameters and/or structures. Each potential model configuration is assessed using a measure of "likelihood," which may represent a true likelihood function (i.e., $P(D_{obs}\|\phi)$ where D_{obs} denotes observed data and ϕ is a model calibration) (Hartig et al., 2011; Vrugt, 2016) or an informal likelihood measure (Beven and Freer, 2001).	By probabilistically representing parameter values and/or structures (which may be multimodal), Bayesian methods explicitly allow for equifinality. In some cases, Bayesian model averaging (BMA) is used to aggregate over multiple models and produce a single distribution of estimates, where each model is weighted by its likelihood (Touhidul Mustafa et al., 2020; Ajami et al., 2007).
Robust decision-making (RDM)	Analysis	Evaluate policy alternatives on their ability to succeed over a wide range of deeply uncertain conditions (Lempert, 2003), as well as identify the sets of conditions under which policies may fail (see "scenario discovery" (Bryant and Lempert, 2010) and "robust regions" (Lempert, 2002)).	RDM incorporates model-level uncertainty into the policy analysis. However, this uncertainty is not restricted to that which produces results consistent with available data. Thus, it does not explicitly incorporate equifinality.
Sensitivity analysis	Analysis	Describe and explain the factors contributing to the variability in model outputs (Abreu and Ralha, 2018; Ligmann-Zielinska et al., 2014; Lee et al., 2015). Part of uncertainty quantification (National Research Council, 2012).	Sensitivity analysis approaches the problem of model-level uncertainty from the perspective of a *single* model and explores the sensitivity around this single point. It also does not consider the extent to which model-level uncertainties contribute to outcomes that are consistent with available empirical data. As such, sensitivity analysis does not address equifinality in models directly.
Model intercomparison	Post-analysis	Compare common scenarios over a range of independently developed models (Robinson et al., 2014; Duan et al., 2019). See also model "docking" (Axtell et al., 1996).	This incorporates a large amount of model uncertainty and even allows for comparison of different model types (e.g., simulation, economic optimization), thus directly allowing for equifinality.

By focusing on a small number of diverse, plausible solutions, this approach efficiently encompasses any equifinality, enabling each calibration to be individually examined and reducing the computational burden of subsequent policy experiments. This kind of approach—with an explicit focus on solution diversity—was first proposed as a method for generating a set of diverse decision options in the context of land use planning (Brill et al., 1982) and has since been applied as a decision support tool in other areas (DeCarolis et al., 2017; Ligmann-Zielinska et al., 2008; Harrison et al., 2001). In this chapter, we extend this work to the context of model calibration and its implications for policy analysis.

We note that the approach, as presented in this chapter, focuses on ABMs, but is also applicable to any process-driven model that requires calibration (e.g., cellular automata, partial or general equilibrium models, system dynamic models, or other simulation-based biophysical models). We also note that the DMC-RPA approach focuses only on model calibration and application. It does not attend to other stages of the model development cycle, such as model validation. Rather, our approach simply suggests that the modeling process be modified to allow for the possibility of multiple plausible model configurations through the model analysis stage (Figure 4.2). The equifinal calibrated models could go through a further validation refinement before being applied to assess policy; we discuss this in section 4.6.

4.3.2 Diverse Model Calibration (DMC)

We implemented the DMC approach in Python 3. Pseudocode is given in Online Appendix B for the published article[3] and the code for the case study application is hosted on comses.net.[4]

Our approach for diverse model calibration uses EAGA (evolutionary algorithm to generate alternatives) (Zechman and Ranjithan, 2004), which is an extension of a conventional genetic algorithm (GA)[5]. Using this approach, we are solving a multimodal, multiobjective optimization problem (Efstratiadis and Koutsoyiannis, 2010; Singh and Deb, 2006) in which the uncertain model parameters and/or structures constitute the decision variables. The objective function has two components, the first representing the degree to which a model's outputs match empirical data and the second representing the degree to which a model's configuration is different from other candidate models. For ABMs of SES, which are stochastic and can exhibit nonlinear dynamics, evolutionary approaches are a useful heuristic method for searching parameter spaces (Reed et al., 2013; Thiele et al., 2014), so are appropriate in this context. Our operational definitions for various terms are given in Table 4.2.

[3] `http://dx.doi.org/10.1016/j.envsoft.2020.104831`

[4] `https://www.comses.net/codebases/5c7710b4-f9c1-47cc-ad61-8734febdb2f0/releases/1.1.0/`

[5] GAs are a form of optimization inspired by Darwin's theory of evolution. A population of "individuals" is evolved toward better solutions. Each individual is characterized by a vector of parameters (i.e., their genes). Selection and reproduction occur within the population and are mediated by each individual's "fitness," which in this case represents the degree to which an individual's output matches calibration data.

Table 4.2: Operational definitions and symbols for the DMC.

Term	Symbol	Definition
Model	-	A process-based representation of a socio-environmental system (e.g., an ABM)
Structure	-	A process-based abstraction of reality within a model
Parameter	-	A continuous or discrete value representing a level or state within a modeled structure
Configuration	S	The set of structures and/or parameters that make up a model
Pattern	r	A qualitative or quantitative stylization of reality
Individual (/solution)	k	A single model within the genetic algorithm
Population	p	A set of individuals
Loss	L	A measure of the discrepancy between a model-generated and empirical pattern
Hyperparameter	-	A value or setting for the genetic algorithm

4.3.2.1 Optimization procedure: Genetic algorithm

In an extension to a regular GA, which consists of a single population of individuals, the EAGA consists of multiple *subpopulations* (SPs) of individuals (i.e., potential model configurations) that coexist in the decision space (Figure 4.3). Each SP is analogous to the population of individuals in a regular GA, but each individual evolves to both increase its fitness to the empirical data (i.e., reduce the "loss" in Equation 4.1, below) and increase its difference from the models in the other SPs (Equation 4.4, below). Evolution and genetic selection occur within each SP (i.e., there is no genetic spillover between SPs) and utilize standard GA evolution procedures.

One SP is defined a priori as the *master* SP. The master SP seeks to find the globally optimal solution and each individual in the master SP evolves solely based on fitness (i.e., diversity is not important). Solutions in other SPs are considered *feasible* when their fitness, given by Equation 4.1 below, lies within some tolerance of the best solution in the master SP (e.g., up to 20% larger). The feasibility of solutions affects the selection of "parents" in the EAGA; the binary tournament selection process selects individuals to act as parents by randomly pairing two individuals and selecting one of these to pass on its genetic material using the following heuristic: if both potential parents are feasible, select the more diverse of the two (i.e., higher $D^{q,k}$ in Equation 4.4 below); if only one potential parent is feasible, select this one; and if both potential parents are infeasible, select the one with the better fitness (Equation 4.1 below).

4.3.2.2 Decision variables: Uncertain parameters and structures

Each continuous uncertain model parameter is defined within some specified bounds (i.e., constraints) and potential process descriptions and categorical parameters are represented using categorical decision variables. In this sense, uncertain structural characteristics are treated no differently than model-level parameters, with the specification of both falling under the general term "configuration."

Initialize: Create $P + 1$ SPs, each with population size K.
for *each generation* **do**
  Evaluate the fitness of all solutions (i.e., the calculate loss, L).
  Feasibility:
  - In SP_0 (master), identify the fittest solution.
  - In $SP_1, ..., SP_P$, determine the feasibility of each solution (i.e., whether each solution is within the tolerance of SP_0's best solution).

  Diversity:
  - Identify the centroid of each SP in the decision space using a fitness-weighted average over each decision variable.
  - For each solution k in $SP_q(q \neq 0)$, calculate the distance $D^{q,k}$ to the closest $SP_p(p \notin \{0, q\})$ centroid.

  Selection: Apply binary tournament selection
  - In SP_0, select parents using fitness only.
  - In $SP_q(q \neq 0)$, select the feasible solution with the higher $D^{q,k}$. If both solutions are infeasible, select using fitness.

  Evolve: Apply recombination operators (e.g., uniform crossover) to combine the genetic material and create offspring from the parents selected in the above step.

end
Choose the feasible solution from each $SP_q(q \neq 0)$ that is most diverse (i.e., highest $D^{q,k}$).

Figure 4.3: DMC procedure

Specifically, each candidate model k is characterized by a set of S configuration elements ($\mathbf{x_k} = x_1, ..., x_S$), which for these calculations are each normalized to the $[0, 1]$ unit interval. In the case of a categorical configuration element, it can be assumed that all categories are equidistant from each other by specifying a single distance (e.g., 1) for models that are different with respect to this element.

4.3.2.3 Objectives: Matching patterns and increasing disparity

The first objective is to minimize the discrepancy between a set of model-generated and empirically-observed patterns. As such, our approach is a form of pattern-oriented modeling. For each individual k, discrepancies are weighted and combined over the R patterns to give a single measure of fit:

$$L_k = \sum_{r=1}^{R} weight_r * discrepancy(model_{k,r}, data_r) \tag{4.1}$$

The discrepancy measure (or "loss," L) is an informal measure of likelihood (Hartig et al., 2011; Smith et al., 2008) similar to those used in other studies (Calvez and Hutzler, 2006; Stonedahl and Wilensky, 2010; Chica et al., 2017). The discrepancy measure could take a number of forms, depending on the type of pattern to be fit (Table 4.3).

The second objective is a measure of difference between a given model configuration and the other candidate models, which is measured as a distance in the configuration space. Because the goal is to evolve increasingly disparate SPs, this difference is calculated between each individual k in SP q and the *centroid* of each of the other SPs ($SP_p(p \neq q)$). The centroid of the p^{th} SP ($\mathbf{C_p}$) is calculated as a fitness-weighted average over its individuals' configuration elements (i.e., parameters and/or structures):

Table 4.3: Potential discrepancy measures for different fitting pattern types.

Pattern type		Discrepancy measure
Model	**Data**	
binary	binary	$\mathbb{1}(model == data)$ where $\mathbb{1}$ is an indicator function
categorical	categorical	$\mathbb{1}(model == data)$ where $\mathbb{1}$ is an indicator function
numeric	numeric	$(model - data)^2$
histogram	histogram	$\sum_{b=1}^{B}(model_b - data_b)^2$ where b are the histogram bins
distribution	distribution	$\int_{-\infty}^{\infty}(F_{model}(y) - F_{data}(y))^2 dy$ where $F()$ is the cumulative distribution function
distribution	numeric	$CRPS = \int_{-\infty}^{\infty}(F_{model}(y) - \mathbb{1}(y - data \geq 0))^2 dy$ where $CRPS$ represents the cumulative ranked probability score
relational	relational	$(\beta_{model} - \beta_{data})^2$ where β is some relation between two quantities, such as a regression coefficient or correlation
spatial	spatial	$sim(model, data)$ where $sim()$ is some measure of similarity between two mapped distributions or features

$$\mathbf{C_p} = \frac{1}{K}\sum_{k=1}^{K} weight_k * \mathbf{x_{p,k}} \tag{4.2}$$

where the best-fitting individual in each SP receives a weight of one, the worst-fitting individual receives a weight of zero, and there is a linear scaling in between based on fitness.

The distance, d, between model k in SP q and SP_p's centroid is calculated as the sum of the absolute differences (Manhattan distance) in the normalized configuration space:

$$d_{q,k\to p} = \sum_{s=1}^{S} weight_s * |x_{q,k,s} - C_{p,s}| \tag{4.3}$$

Different weighting schemes or distance calculations may be chosen, if desired. Finally, for each model k in SP q, the second component of the objective function, $D^{q,k}$, is then evaluated as the distance to the closest SP centroid:

$$D^{q,k} = \min_{p:p\neq q} d_{q,k\to p} \tag{4.4}$$

4.3.2.4 Selecting models

Given the stochasticity in both the model and the GA, the objectives (Equations 4.1 and 4.4) are likely to be non-monotonic and will fluctuate as each SP evolves. To assess convergence, the modeler can visually assess the two components of the objective function and discern whether they have stabilized. Because the DMC evolution is likely highly influenced by the dynamics of the particular socio-environmental model, we do not present a quantitative measure for assessing its convergence.

Once convergence has been reached, a single solution is chosen from each SP (not including the master SP). To prioritize diversity among the selected models, select the most diverse feasible solution within each SP.

4.3.2.5 Selecting DMC hyperparameters

The DMC procedure contains several hyperparameters that need to be specified by the modeler. For example, as the number of SPs (N_{SP}) is increased, the well-fitting regions of the parameter space will become more crowded and, in turn, the solutions less diverse. There is hence a tradeoff, whereby if N_{SP} is too low, plausible regions of the parameter space may not be discovered, and if N_{SP} is too high, the solutions lose their diversity and begin to collapse on top of each other. Applications of this approach should therefore explore the effect of varying N_{SP} on the solutions reached by the DMC procedure. Alternatively, practical considerations may drive the choice of N_{SP}. For example, if the individual models are to be presented to decision makers or if subsequent policy-related computational requirements are high, the number has to be manageable. The example presented in the original description of the EAGA included four SPs (Zechman and Ranjithan, 2004). Other hyperparameters such as the number of generations and the population size within each SP will also affect the solutions reached. Again, the effect of these values on the results should be assessed to encourage appropriate choices. We present an example assessment of hyperparameter values for our case study application in Appendix C.1.

4.3.3 Robust Policy Analysis (RPA)

Following the identification of a diverse set of model calibrations, the second step is to assess the implications for system behavior under policy intervention. This simply involves conducting the same policy assessment for each selected model. There are two general possible classes of outcome: (1) if results are consistent over the set of models, we can have greater confidence in any policy recommendations; and (2) if results are qualitatively different between the selected models, our approach has exposed sensitivities that may have been missed in an analysis using a single, best-fit model. In this case, the equifinal models can be explored to suggest the potential socio-environmental conditions or mechanisms that may give rise to the success or failure of a policy intervention, or additional data can be included to attempt to restrict equifinality (see the discussion in section 4.6). In either case, we achieve a more robust policy analysis.

4.4 Case study description: Smallholder climate resilience

Using an ABM of smallholder farmer resilience, we apply the DMC-RPA approach to generate a diverse set of models and explore whether these lead to policy-related assessments that are qualitatively consistent. We give a brief overview of the ABM here. An ODD+D protocol (Müller et al., 2013) is provided in Online Appendix A for the published article and in Williams et al. (2020a).

4.4.1 ABM description

Smallholder agricultural systems are highly vulnerable to climatic variability (Vermeulen et al., 2012). It is therefore important to identify ways through which their climate resilience can be supported (Hansen et al., 2019). They also exhibit key properties of complex adaptive systems (de Vos et al., 2019); smallholder populations are highly heterogeneous in their attributes and access to capital, and household-level mechanisms to cope with shocks (e.g., selling of livestock or assets) can lead to path dependencies and poverty traps (Haider et al., 2018). Additionally, interactions between smallholder households and agroecosystems can give rise to dynamically evolving system trajectories (Giller et al., 2011; Tittonell, 2014). Thus, agent-based modeling is an appropriate tool through which to assess these resilience dynamics (Bitterman and Bennett, 2018; Schlüter et al., 2019b).

The purpose of the ABM is to provide temporal and distributional assessments of smallholder drought vulnerability and the potential household- and community-level effects of selected resilience-enhancing strategies. The ABM is designed to represent an Ethiopian smallholder mixed crop-livestock farming system. It draws from several sources of empirical data to represent the conditions of Amhara in the Ethiopian highlands. However, the model is not intended to produce policy recommendations for a specific location; rather, it serves as an experimental platform to evaluate the potential effects of resilience-enhancing strategies in smallholder systems more generally.

Each agent represents a single smallholder household. The modeled livelihood activities include farming, livestock rearing, and wage-based employment (Figure 4.4). Agents are heterogeneous with respect to their household size, land holding, and risk aversion. Additionally, each agent has preference for either maximizing wealth or leisure. Livestock are grazed on a combination of on-farm crop residues and a communal rangeland system, which each provide amounts of fodder that vary over time based on both climate and endogenously-driven rangeland demand. The availability of wage-based employment is exogenous and does not vary over time. Climate affects the following model components: crop yields, which are calculated at an agent-level on an annual basis; rangeland dynamics, which is simulated at the regional level on an annual basis; and crop prices, which vary each month at the regional level, but are exogenous to the modeled system.

At the beginning of each year, agents make decisions about how to manage their farmland (fertilizer application, planting date), whether to buy/sell livestock from their herd, and how much labor to allocate to non-farm wage-based employment. These options are represented as a finite

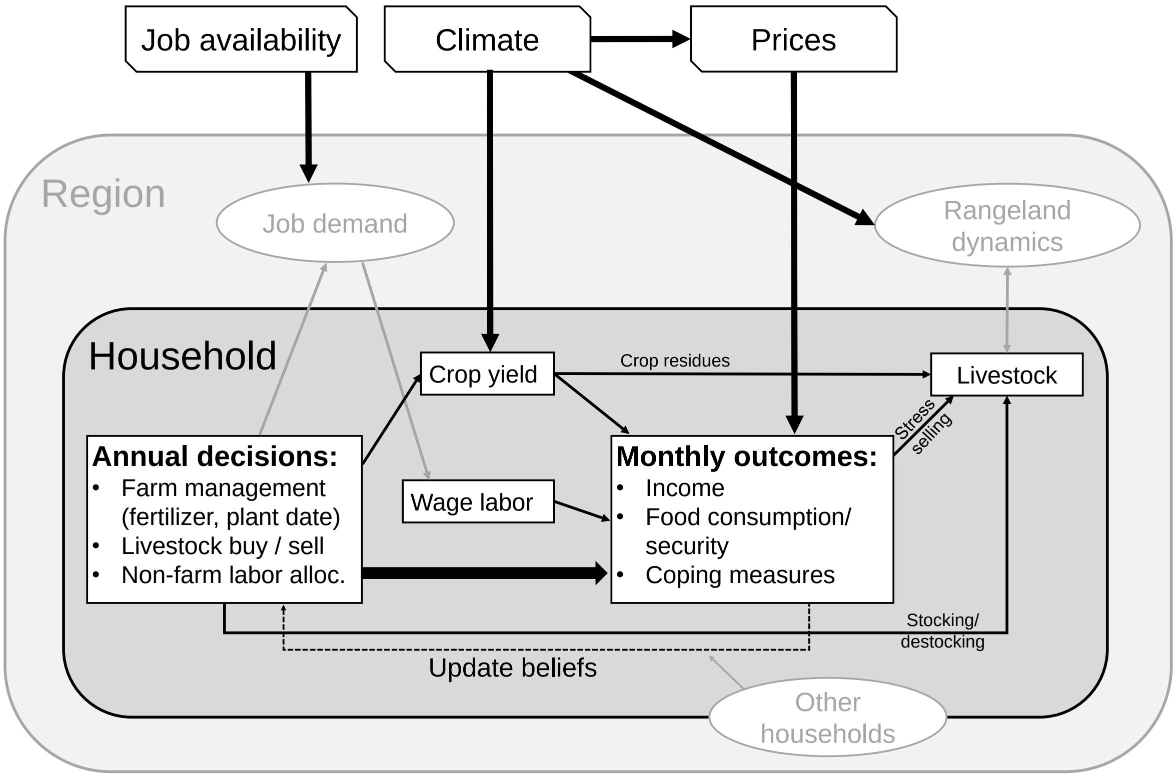

Figure 4.4: Schematic overview of the smallholder resilience ABM. Light grey circles represent processes through which a household directly or indirectly interacts with other households. Objects outside the large box are exogenous.

set of livelihood options. A single crop type (maize) is modeled. These start-of-year decisions are made under a degree of uncertainty about the future climate and market conditions. Agent-level beliefs about these conditions are formed from their previous experiences as well as interaction with neighboring agents. Following these decisions, crop yields are calculated and monthly wage employment allocations are made. Crop yields are influenced by both agent decisions and the exogenous climate, and wage employment allocations depend on the regional demand for labor (Figure 4.4).

At each month throughout the year, agents attempt to satisfy their food and cash consumption requirements through their own crop production, food stores, and cash holdings. Agents can buy and sell crops from the market each month as well as sell livestock as a coping measure to smooth cash and food consumption. Each month that an agent cannot satisfy their food requirements they are classified as "food insecure." The primary output of the model is this binary, monthly, household-level representation of food security, which emerges as a result of interactions between the different modeled components of the agricultural system.

4.4.2 Calibrating diverse models

4.4.2.1 Uncertain model characteristics

The ABM contains a variety of uncertain parameters and structures (Table B.4). Although the differences in alternative model structures are relatively minor, they are sufficient to demonstrate how the DMC procedure works to evaluate potential model structures. Other applications could integrate DMC more thoroughly into model structure development, in addition to identifying model parameter values.

Critically, what Table B.4 demonstrates is that uncertain model structural representations—although fundamentally different from parameters in how they affect the ABM—are treated no differently by the DMC approach, and are simply coded as binary or categorical switches. For example, we allow the possibility for two alternative decision-making representations: expected utility maximization and satisficing (exp_util_DM in Table B.4). In doing so, we apply our approach to contrast alternative theories (Grimm et al., 2005). Under expected utility maximization, each agent chooses the livelihood option each year that maximizes either their wealth or their leisure time, depending on their preference. Under satisficing, all agents have two levels of hierarchical objectives (Kaufman, 1990): (1) to choose the option that leads to the lowest food insecurity; and (2) to choose the option that maximizes expected wealth or leisure time, depending on their preference. The second-level objective only activates if multiple options tie with respect to the first objective. These two alternatives represent different functional representations in the ABM, yet their reduction to a binary switch is more an issue of model design than a qualitative difference in the DMC procedure.

4.4.2.2 Patterns

Using data from the 2015 World Bank's Living Standards Measurement Study (LSMS), we identified eight emergent outcomes (i.e., patterns) that we wish the model to match. To generate these patterns, the ABM was run from 2003-2015 and outputs from the final year of the simulation were compared against the calibration data, which represent the empirical conditions in Amhara in 2015.

Five of the patterns represent agent livelihood characteristics. We represent these using histograms, thus combining information from the household level into a regional pattern. These distributions include: non-farm labor allocation, agricultural labor allocation (farming and livestock), months of food insecurity, subsistence fraction (i.e., percent of production consumed), and large livestock holdings.

In addition to these distributions, we include three binary indicators representing desirable qualitative model-level characteristics. The first specifies that at least 70% of the agents choose to farm their land on average. This indicator is included to encourage the generation of models

Table 4.4: Uncertain parameters and structures in the ABM. All continuous parameters were initialized uniformly over the specified ranges, which were defined as conservative bounds of expected realistic values.

Name	Unit	Range	Description
Structures			
exp_util_DM	-	$\{F, T\}$	Agents make decisions to maximize expected utility (T) or using a satisficing process (F)
$labor\_hiring$	-	$\{F, T\}$	Ability for agents to hire others to work on their own fields
$job\_mkt\_participants$	-	$\{all, leisure\}$	The agents that experience limited non-farm job availability (all=all agents, $leisure$=leisure-maximizing agents only)
Continuous parameters			
min_sust	%	$20, 70$	Sustenance deficit threshold for "severe" food insecurity
$risk\_aversion\_mult$	-	$1, 1000$	Multiplier on risk aversion coefficient
$frac\_income\_max$	-	$0, 1$	Fraction of income-maximizing agents
$job\_availability$	hours/month	$0, 5$	Non-farm wage labor job availability per agent
$labor\_wage$	Birr/day †	$8, 100$	Non-farm labor wage per day
per_cap_labor	hours/day	$5, 13$	Total labor availability per adult equivalent
$farm\_labor\_mag$	-	$100, 1600$	Multiplier on farm labor requirements
$ls\_per\_head\_herding$	hours/head	$0, 50$	Monthly livestock labor requirement
$planting\_fraction$	fraction	$0.5, 1$	Fraction of land that can be planted
$grass\_regen$	-	$1, 5$	Rangeland grass regeneration rate
$rf\_intercept$	-	$-1, 1$	Rainfall effect on rangeland grass with zero rainfall
rf_slope	-	$0.001, 0.004$	Sensitivity of rangeland grass to rainfall
ls_max_growth	-	$0.2, 0.5$	Maximum livestock reproduction rate
g_max	kg/ha	$400, 10000$	Maximum rangeland grass biomass
Discrete parameters			
$months\_b4\_ls\_coping$	months	$\{0, 1, 2, 3\}$	Months of food insecurity before agent considers selling livestock

† Birr is the Ethiopian currency

in which farming is the dominant livelihood activity, which is consistent with empirical data for the modeled region (CSA, 2017)[6]. The second indicator specifies that the grass biomass in the communal rangeland does not decrease to zero at any time to discourage unrealistic rangeland dynamics. The third indicator specifies that no agent should ever have more than 80 head of livestock, which further encourages livestock holdings to be consistent with the empirical LSMS data.

To measure the discrepancy between the model outputs and the histograms, we convert each histogram into its empirical cumulative distribution function (ECDF) and calculate the average squared difference between each ABM-generated and empirical ECDF step. This represents a discretized version of the "distribution-distribution" pattern type in Table 4.3. We chose this as it bounds the maximum possible loss for each distributional pattern between zero and one, meaning that each distribution exerts comparable influence on the total loss. We weight all distributions equally, as each distribution was chosen to represent a relevant, independent livelihood characteristic. An additional value of one is added to the total loss for each qualitative pattern that the model does not generate. Thus, the overall loss is bounded between zero and eight.

4.4.2.3 Genetic algorithm

We conducted an experiment with four SPs (five total, including the master), each comprised of 30 individuals, run for 300 generations (Table 4.5). We chose to present the results for four SPs in this chapter for visual clarity, but we also experimented with alternative numbers of SPs and population sizes: see Appendix C.1. To select models, we selected the feasible individual from each SP that is most distant from any other SP's centroid.

Although the ABM is stochastic, for the calibration we ran a single simulation replication for each model configuration and calculated the loss from this single set of outputs. We found the variability in model outputs to be much larger *between* model configurations than *within* each model configuration, so opted for this approach due to computational feasibility.

To assess the sensitivity of each parameterization, we conducted a local sensitivity analysis. We systematically perturbed each parameter from its calibrated value and assessed the effect on the fit to the empirical data (see Appendix C.5).

4.4.3 Policy analysis

The objective of the case study is to examine which of two resilience-enhancing strategies provides the greatest benefit to climate resilience in the smallholder agricultural system. The strategies in-

[6] We note that the LSMS reports higher farming percentages than this (89% in Amhara in 2015 (CSA, 2017)). We use a more permissive value because we do not model land rental dynamics, which is a common practice in the modeled region but would unnecessarily complicate the model.

Table 4.5: EAGA settings for the smallholder resilience case study.

Name	Value	Description
Number of SPs	4	Not including the master SP
Number of generations	300	-
Population size	30	-
Probability of mutation	2%	-
Type of crossover	Uniform	Each of the offspring's genes is chosen randomly from their parents
Type of selection	Binary tournament	Objective function value is used as the selection criterion
Feasibility threshold	30%	Individuals within 30% of the master SP's best solution are considered feasible

clude the provision of seasonal climate forecasts and a 20% increase in the availability of non-farm employment opportunities. These are not necessarily explicit *policies* at a government or institutional level, but represent potential policy-relevant interventions to the system (Federal Democratic Republic of Ethiopia, 2019). The climate forecasts give the agents information at the start of each year about the upcoming climate conditions. This information is not perfect, but enables the agents to make better-informed agricultural decisions (e.g., shifting their planting date or choosing to not apply fertilizer to their fields), potentially increasing their climate resilience. Increased job availability, in contrast, provides an opportunity for agents to diversify their livelihoods and can act as an important source of income for agents that do not otherwise have access to land- or livestock-based capital. Given the different mechanisms through which these strategies operate, they may differently affect the households' and, in turn, the system's collective ability to respond to drought under alternative socio-environmental conditions.

We quantify specific resilience (Carpenter et al., 2001) using a measure of the effect of drought on the agents' food security. We represent droughts by imposing reductions in rainfall; a 50% drought represents a year in which the rainfall is reduced by 50%. This affects the crop prices, rangeland dynamics, and crop yields. The effects on prices and rangeland dynamics are at the regional level, while the effects on crop yields are both nonlinear in rainfall and spatially explicit. To isolate the overall effect of a 50% drought, we run two simulations: one for 30 years under regular climatic variability, and one for 30 years under the same climatic variability but with a 50% drought imposed in the fifth year of simulation. In every month of each simulation, we record the percent of agents that are food insecure (i.e., are unable to satisfy their food consumption requirements). We then take the difference between food insecurity in these two simulations, giving a monthly measure of the *additional* percentage of the ABM agents that are food insecure as a result of the drought. When analyzed over time, this incorporates both the initial impact of the drought and the long-term recovery of food security in the years following. Because food security varies throughout the year, this will exhibit an annual cycle.

To isolate and compare the strategies' resilience-enhancing benefits, we evaluate this measure of resilience both under baseline conditions and with each of the two strategies in place. We assess a strategy's overall benefit as the cumulative amount of food insecurity that it avoids in the wake of the drought—i.e., the total number of agent-insecurity months avoided in the 25 years following the drought. Finally, we calculate the difference between the two strategies' resilience benefits (Δ_{res}) as:

$$\Delta_{res} = \sum_{a=1}^{N} \sum_{y=5}^{30} \sum_{m=1}^{12} \left(FI_{a,y,m}^{climate\ forecast} - FI_{a,y,m}^{job\ availability} \right) \tag{4.5}$$

where a indexes the N agents, y indexes the years, m indexes the months, and FI denotes the incidence of household-level food insecurity.

There are two levels of stochasticity relevant for this policy assessment, both of which will affect the quantity in Equation 4.5. The first level represents within-model stochasticity introduced in the assignment of the agents to the landscape, crop yield calculation, and allocation of regional wage labor and livestock reproduction/mortality. To account for this, we replicate each simulation 50 times[7]. The second level represents uncertainty associated with the drought; as already described, we define our droughts using single-year reductions in rainfall. However, the ultimate effects of this on the smallholder system will depend on both the preceding and succeeding climatic conditions[8]. To account for this, we generated 40 "climate timeseries," each 30 years long, by repeatedly randomly sampling years from the 2000-2015 observational climate record. We impose a drought in the fifth year of each timeseries.

Finally, to assess sensitivity in the policy comparison, we conduct two additional experiments: one with nine SPs (ten, including the master) and one assessing the effect of a 20% drought.

[7] A convergence analysis ((Law, 2008), pg. 502) determined that 50 replications were sufficient to achieve a relative error of 0.1 (i.e., $|\bar{X}(n) - \mu|/|\mu| < 0.1$ if $\bar{X}(n)$ is the estimate based on n replications and $\mu = E[X]$) with a confidence level of 90%.

[8] For example, the effects of the recent Ethiopian drought—in which some parts of the country only experienced 50-75% of the regular rainfall—were in part exacerbated by continued dry conditions in the following year (Singh et al., 2016).

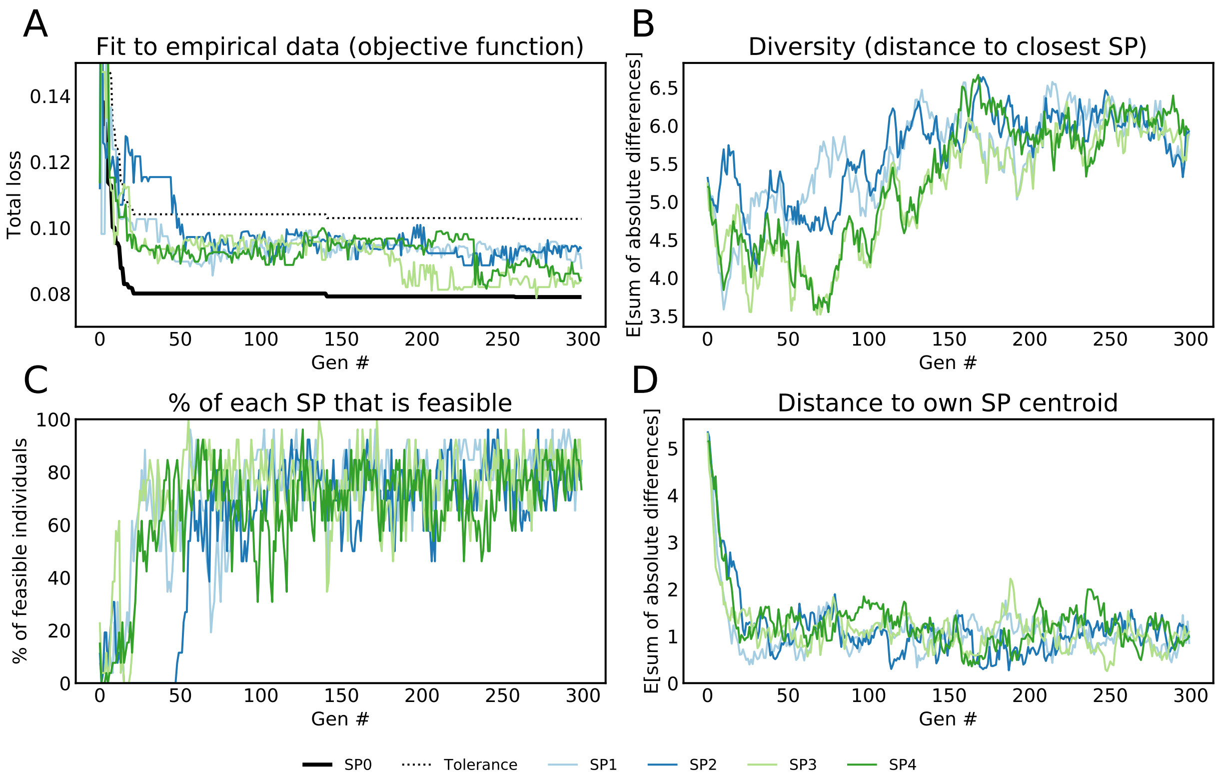

Figure 4.5: Convergence of the genetic algorithm with four SPs. (A) Overall convergence of the objective function (i.e., the loss; Equation 4.1). (B) The diversity of the feasible solutions (Equation 4.4 evaluated over the SP centroids). (C) The percent of individuals in each SP that are within the specified tolerance of the master SP's objective function. (D) The separation within each SP. Note that the master SP ($SP0$) is only displayed in A, as measures of diversity are not calculated for the master SP.

4.5 Case study results

4.5.1 Diverse model calibration

The master SP's solutions converge in their fit to the empirical data after approximately 25 generations of the genetic algorithm (Figure 4.5A). After approximately 50 generations, the solutions in all other SPs begin to become feasible (Figure 4.5A and C)—i.e., within 30% of the master SP's best solution. Once the solutions are feasible, their diversity begins to increase (Figure 4.5B) and stabilizes after approximately 175 generations.

Overall, the ABM does well at matching the empirical patterns. The total summed loss in the master SP is approximately 0.08 (Figure 4.5A) out of a maximum possible value of 8 and all other SPs are within 30% of this, indicating that all qualitative patterns are matched and the levels of fit to the empirical histograms are high. However, all calibrated models overestimate the percentage of households with no livestock herds (Figure 4.6 and Appendix C.2). Livestock represent an

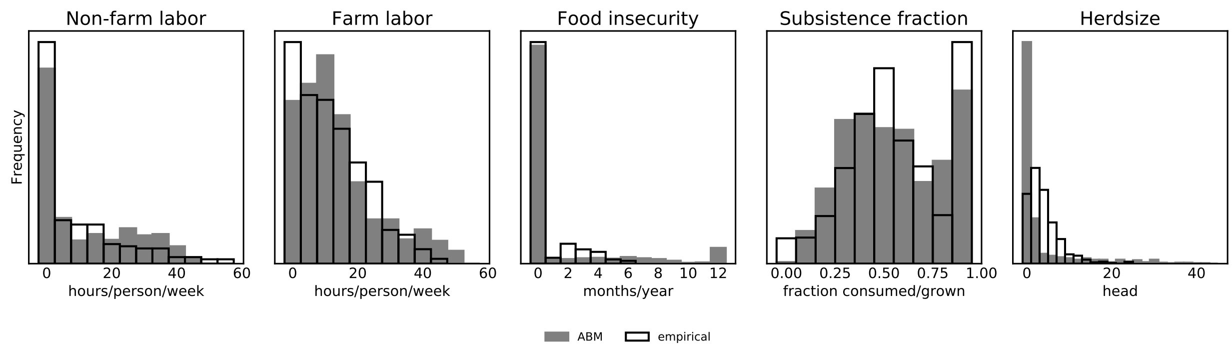

Figure 4.6: Comparison of empirical and ABM-generated patterns in SP 1. The patterns include: non-farm labor (hours spent on non-farm activities); farm labor (hours spent on farming and livestock activities); household-level food insecurity (number of months experiencing food shortages within the past year); subsistence fraction (the fraction of crop production consumed by the household); and herd size (number of large livestock). Comparisons for the other SPs are shown in Appendix C.2.

important coping mechanism both in reality and in the ABM; given this, our models may underestimate resilience or lead to biased policy assessments. For example, climate forecasts provide the agents with information that can aid their livestock stocking/destocking decisions; because our model underestimates the number of households with livestock, it may underestimate the benefits of climate forecasts through this mechanism. In addition, the calibrated models exhibit different levels of fit to the non-farm labor distribution (Figure 4.6 and Appendix C.2); SP2 underestimates and SP4 overestimates the proportion of agents engaging in non-farm labor.

Some variation is to be expected in the performance of any set of equifinal models, and that variation is important to provide structure to the variation in any policy-relevant conclusions. It is therefore important to incorporate understanding of the variations in level of fit into policy analysis based on the calibrated models, as different levels of fit may imply different degrees of credibility over the model set. Future work could expand the range of structural representations included in the model calibration to improve the level of fit to the livestock distribution, as well as include a validation step to filter out models from the calibrated set that less adequately represent reality in ways that might significantly affect the policy analysis and, in turn, bias conclusions.

4.5.2 Equifinality in the calibrated models

The selected model configurations are diverse (Figure 4.7), suggesting that the complexity in the model allows for very different parameter sets to produce similar levels of fit to the data. Specifically, the distance between each of the SPs in the normalized parameter space is approximately 6 (Figure 4.5B). Given that there are 18 uncertain model elements (Table B.4), the maximum possible distance (Equation 4.3) between any two models is 18.

Different parameters are more or less consistent over the model set. For example, the risk aver-

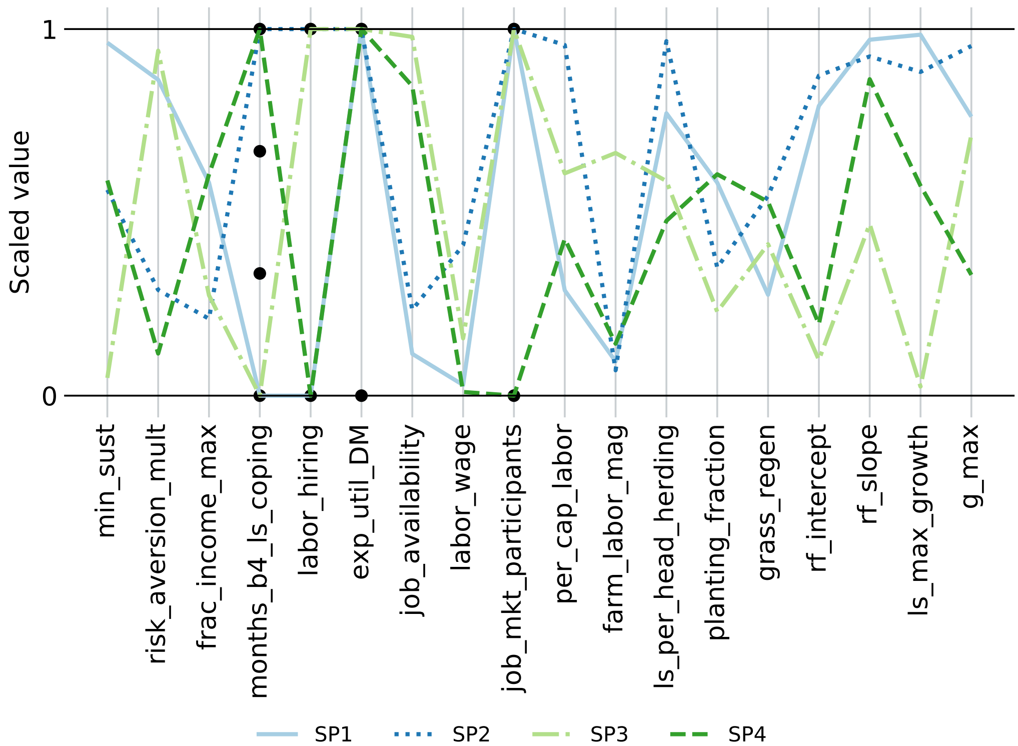

Figure 4.7: Parallel coordinate plot showing the resultant parameterizations. All parameters are normalized to a [0,1] interval given by their initial lower and upper bounds (Table B.4). Black points represent the possible values for model structural elements and discrete parameters. Variable descriptions are provided in Table B.4.

sion parameter ($risk\_aversion\_mult$), which represents a dimensionless multiplier on the agents' risk aversion coefficient, is present over almost its entire range in the four models (Figure 4.7). This implies that the model is insensitive to changes in this parameter. In contrast, parameters such as $planting\_fraction$ are only present over a smaller range in the calibrated model set (Figure 4.7). This implies that the model is more sensitive to changes in these parameters and that only a narrow range of values produce plausible model outputs. Our analysis of the sensitivity of the model calibrations (Appendix C.5) confirms these observations and additionally shows that the SP3 and SP4 calibrations are less stable; slight perturbations in their parameters result in some of the qualitative fitting patterns not being matched, thus degrading the calculated fit.

We observe an interesting result in the uncertain model structural representations; all four models contain agents that utilize expected utility maximization to make their livelihood decisions (exp_util_DM in Figure 4.7)—i.e., the satisficing decision-making representation is never selected by the calibration procedure.

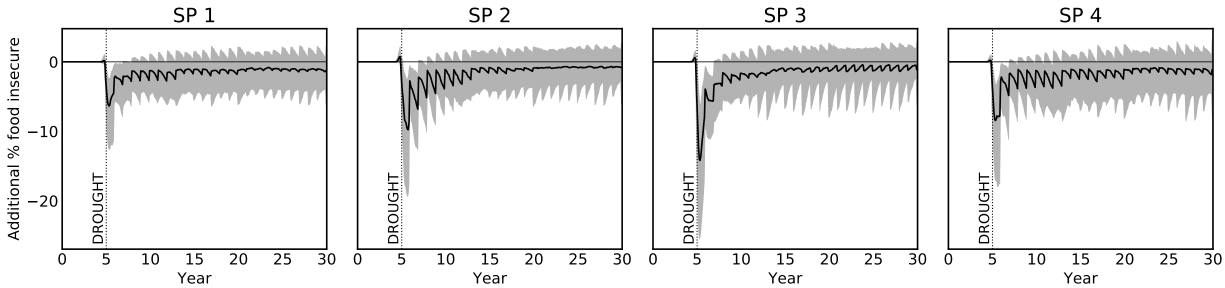

Figure 4.8: The additional percentage of agents that become food insecure as a result of a 50% reduction in rainfall in year 5 of the simulation. Lines plot the median response over all ABM replications. Shaded regions show 5/95% intervals over all model replications.

4.5.3 Policy analysis: enhancing climate resilience

With no policy in place, the effect of the drought on household food security differs over the selected models (Figure 4.8); for example, the model selected from SP1 exhibits the smallest drought vulnerability, with only a maximum of 6% of households affected by the drought (at the median simulation output). In comparison, in SP3, food insecurity is at the median 14% higher in the year following the drought. SP3 also exhibits the highest levels of food insecurity under baseline conditions (Appendix C.2). In terms of recovery, in no model does food security recover completely to its level in the no-drought counterfactual (Figure 4.8), showing that, in all cases, the drought permanently alters the livelihood trajectory of some households. The differences between the SPs suggest that in this case there are implications of equifinality, as each of the equifinal models exhibits different behavior when applied to a situation (drought event) not used in the calibration.

When comparing the effects of the two interventions on the system's resilience, all four models yield the same directional result: climate forecasts offer larger potential benefits to food security in the wake of a drought than an increase in job availability (Figures 4.9 and 4.10). Thus, in spite of the differing model configurations (Figure 4.7) and baseline levels of vulnerability (Figure 4.8), there are no large implications of equifinality for our policy analysis. In this case, the DMC-RPA approach has therefore yielded a conclusion that is likely similar to that utilizing a single model calibration, but the consistency of this result over the diverse calibrated models increases the robustness of this conclusion.

4.5.4 Sensitivity analyses

Reducing the magnitude of the drought to 20% does not affect these conclusions; the 20% drought has a smaller, yet similar effect on household-level food security, and the climate forecast still consistently provides larger benefits under these drought conditions (Appendix C.3).

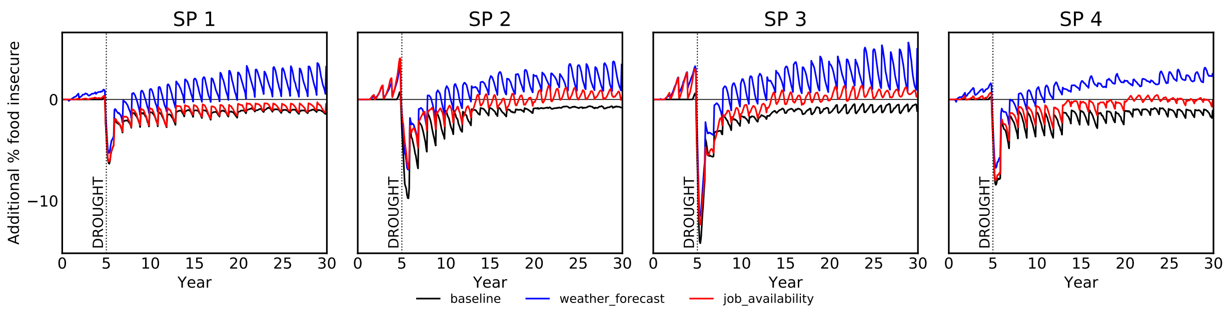

Figure 4.9: Effects of resilience-enhancing strategies on food insecurity. Lines plot the median response over all ABM replications, relative to the no-drought conditions.

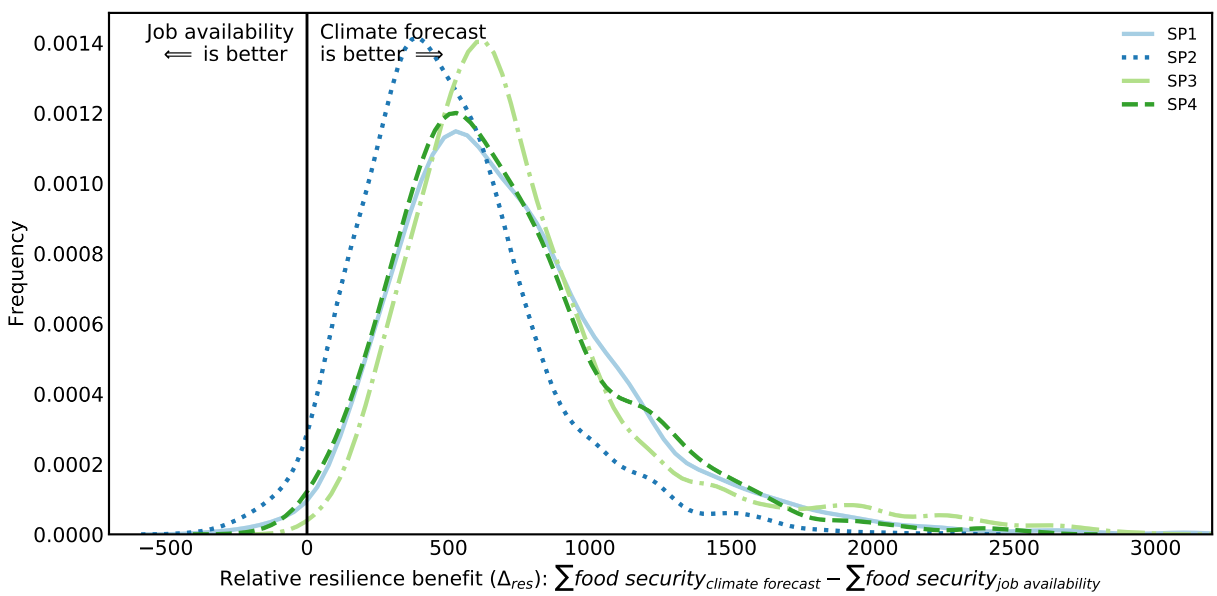

Figure 4.10: Comparison of the resilience-enhancing benefits of the two interventions. Positive values indicate that climate forecasts provide greater cumulative benefits to food security than an increased availability of non-farm jobs (Equation 4.5). Distributions represent the variability over the stochastic simulations.

In an experiment with nine SPs, not all SPs become feasible after 300 generations and, due to the larger number of solutions, are more closely located within the uncertain parameter space (Appendix C.4). However, the overall diversity of the feasible solutions is higher; for example, in contrast to the results above in which all models specified expected utility maximization, two SPs in this experiment specify satisficing as the decision-making framework. Additionally, the overall range of many continuous variables over these models is higher. This higher diversity between parameterizations contributes to a higher diversity in drought responses and policy comparisons (Appendix C.4). However, the result from above—that climate forecasts provide greater resilience benefits—is demonstrated in seven of the eight retained parameterizations. Thus, we conclude that this result is a robust one.

4.6 Discussion

4.6.1 DMC-RPA for model development and inference

The DMC-RPA approach allows the researcher to retain multiple hypotheses, represented by multiple disparate model configurations, through the analysis stage of the model development cycle (Figure 4.2). Doing so is in accordance with the notion of strong inference (Platt, 1964). Although we have focused on the implications for policy analysis, the DMC approach also presents opportunities for model development, model-driven theory development, and systems inference.

First, the parameters and structures of the calibrated models themselves—i.e., the conditions *measured* in the calibration—may suggest different socio-demographic and environmental contexts under which similar outcomes may be possible. Parameters that are especially variable or consistent may represent critical or sensitive factors in driving the empirical outcomes. Alternatively, if one process representation is consistently selected over another, this constitutes a form of combined model-based and empirical evidence supporting its appropriateness for describing reality in the modeled context. This could be compared against existing theory and evidence to aid in developing more generalized knowledge (Magliocca et al., 2018; Schlüter et al., 2019b,c).

Second, the calibrated models could be compared along some *unmeasured* axes—i.e., information not assessed in the calibration process, yet contained within the calibration simulations (Wiegand et al., 2003; Khatami et al., 2019). If these unmeasured axes represent factors that are *unobservable* or *unobserved* in reality, they could provide new information about the diverse mechanisms that may give rise to the observed outcomes, acting as a complement to studies that use statistical methods to explore causal mechanisms in empirical data (Ferraro and Hanauer, 2014). Alternatively, unobserved factors that differ over the set of models could be identified to prioritize the collection of new empirical observations.

A key context within which model equifinality is salient is when model complexity outstrips the availability of calibration data (Schmidt et al., 2020). In these contexts, a single calibrated model runs the risk of being "overfit" to the data—i.e., it could fit the available data well but have low ability to recreate patterns not used in its calibration (van Vliet et al., 2016; Sun et al., 2016). By using multiple, information-rich patterns, POM attempts to reduce this concern (Latombe et al., 2011; Grimm and Railsback, 2012). Sensitivity analysis, such as in Appendix C.5, can indicate the stability of each parameterization, but does not indicate the degree to which the calibrated model may be overfit to the data. If the model outputs exhibit an extremely high variability under conditions not used for the calibration (e.g., a policy analysis), this may indicate that overfitting is occurring (Calvez and Hutzler, 2006). However, it is indeed possible that there are multiple plausible, discordant model representations. In any case, acknowledging that there need not be a single, "optimal" solution reduces the risk that incorrect inferences are made, even if each individual solution is overfit.

It is possible that including new tests of emergent model characteristics—i.e., testing against data that are *observable* in reality, yet *unmeasured* in the original calibration procedure—would reduce equifinality (Platt, 1964; Latombe et al., 2011; Guillaume et al., 2019). This could involve, for example, a "model output corroboration" (Augusiak et al., 2014) or cross-validation procedure that employs patterns or tests unmeasured in the calibration process itself (Wang et al., 2018; Wiegand et al., 2003). If the model is spatial, it could be validated by applying it or comparing its outputs to those from a different region (Magliocca et al., 2015; Brown et al., 2005). Alternatively, consultation with domain experts could be used to filter out unreasonable model configurations (van Vliet et al., 2016).

With respect to model development, our approach is similar to the virtual laboratory and "building block" approaches (Magliocca and Ellis, 2016), which use pattern-oriented modeling to systematically evaluate hypothesized model structures—potentially representing contrasting theories or differing levels of complicatedness (Sun et al., 2016)—against empirical data. Our approach builds on this by identifying multiple, diverse model configurations. Systematic procedures for model structure and parameter specification that allow for equifinality have been more extensively developed in the field of hydrology (Khatami et al., 2019; Touhidul Mustafa et al., 2020) and could be turned to for future inspiration in the SES modeling community. In principle, DMCA-RPA is a formalization of the pattern-oriented modeling approach and we situate it within this body of research.

Finally, as we have previously stated, the policy assessment can lead to two general classes of outcome; either the conclusions are similar over the set of models, or they are not. In our case study example, we observed remarkably consistent effects over the four model configurations. This implies a greater level of robustness of our conclusions than if we had used a single model parameterization. Had we instead observed inconsistent effects, these results could be used to shed light on the socio-environmental conditions under which different interventions may be more or less effective, helping to inform the targeting of policy interventions (Giller et al., 2011).

4.6.2 Comparison to alternative calibration approaches

There are three main features that set our approach apart from alternative model calibration methods: (1) identifying a small set of N models, (2) retaining these models as separate, and (3) maximizing model diversity. These features have both philosophical and practical implications in interpreting the model configurations and policy analysis results.

First, our approach identifies a small set of N models (in the case study application, we chose N=4). As we have discussed, this allows for equifinality and is more appropriate than any "best-fit" model calibration procedure. By calibrating a *small* set of distinct models, each model configuration can be individually examined, enabling enhanced inference and communication with decision-makers (Schwartz, 2012). Having a small number of models is also particularly advantageous in situations when subsequent policy-related experiments are computationally expensive. In these situations, it is desirable to not only have efficient sampling over the prior distributions for the parameters (Vrugt and Beven, 2018; Yen et al., 2014), but also to efficiently encompass equifinality in a small set of models. MCMC-based calibration approaches require repeated sampling from the posterior distribution, and in high-dimensional cases a large number of samples may be necessary. For example, the application in an initial presentation of the MCMC-based DREAM algorithm used 2,500 draws from the posterior distribution (Vrugt et al., 2009). Other implementations of multimodal evolutionary algorithms for model calibration use diversity-based filtering to reduce the number of solutions (Chica et al., 2017; Moya et al., 2019), so are comparable to our approach in this respect.

Second, and related to the first, we do not assign a probability or relative likelihood to each model, but present each model configuration and its policy assessment separately. This differs from Bayesian model calibration methods that estimate a posterior distribution for the model parameters/structures and use this to produce a single predictive distribution (Ajami et al., 2007; Vrugt et al., 2008; Hartig et al., 2011; Touhidul Mustafa et al., 2020). Similarly, Monte Carlo calibration approaches typically aggregate outputs over the entire set of behavioral models (Beven and Freer, 2001) (although an interesting exception exists in (Khatami et al., 2019)). Conceptually, our approach is motivated by the deep uncertainty in modeling socio-environmental systems, which makes it problematic to assign a probability to each model (Polasky et al., 2011). In this regard, the DMC-RPA approach is similar to robust decision-making (Lempert, 2003) and we employ it to identify policies that are beneficial over a wide range of potential states, without assigning a probability to these states. If a single predictive distribution is desired, however, our approach does not preclude using the model fits (Equation 4.1) or, preferably, the success of each model in some independent validation exercise to develop (informal) weights for each model.

That we do not assign a probability to each model configuration also demonstrates some similarities (and differences) to both the "limits of acceptability" (LOA) (Beven, 2006) and the POM concepts (Grimm et al., 2006). In LOA, models that satisfy a number of predetermined acceptable limits (e.g., consistently produce outputs within 20% of an observed value) are characterized as "behavioral." No behavioral model is considered to be more or less behavioral than another. POM also employs an equivalent behavioral notion. Similarly, in our approach, the final model configurations are each equally considered. However, our approach differs in that model configurations are defined as feasible based on their fit relative to the master SP's best solution. This feasibility criterion is not defined a priori, but evolves with the algorithm (e.g., dashed line in Figure 4.5A). This helps to maintain the balance between fit and diversity during the genetic algorithm's evolution, but is not as strongly based on theory or domain expertise as is required for LOA or POM (Beven, 2006).

Third, our approach is unique in how it *maximizes* diversity within the feasible model set while staying within a specified tolerance of the "optimal" model. Other set-theoretic and MCMC-based calibration approaches generally aim to *maintain* model diversity, which is accomplished by removing or penalizing solutions that are similar in the configuration space (i.e., the parameters and/or structures that they specify) (Singh and Deb, 2006; Olalotiti-Lawal and Datta-Gupta, 2015). Our approach, because of the dual objectives—i.e., model fit and diversity—is technically a bi-objective optimization. These are not combined into a single objective, nor do we seek to find a Pareto-optimal set of solutions. Our approach does not maintain that diverse models are more desirable, rather that the diverse set of plausible models most efficiently encompasses equifinality. The diversity component enters the objective function only in models that are feasible (e.g., with an error within 30% of the master SP's best solution). Due to this treatment of the two objectives, combined with the fact that we do not seek to identify a posterior distribution, our approach could not to the best of our knowledge be directly integrated into an MCMC procedure, at least without modification. Thus, the DMC procedure constitutes a distinct approach to model calibration.

While providing important computational and inferential advantages, these features of the DMC-RPA approach may also present tradeoffs in certain situations. Importantly, because the number of models (N_{SP}) must be specified a priori, the algorithm is limited in its ability to characterize the structure of unknown objective spaces. Other evolutionary approaches, such as niching genetic algorithms (Goldberg and Richardson, 1987; Miller and Shaw, 1996; Deb et al., 2002), can more flexibly identify an unknown number of local optima. These can be filtered to a smaller number ex-post, if desired (Chica et al., 2017). Additionally, although Monte Carlo methods are inefficient in exploring the parameter space (Vrugt and Beven, 2018), they also do not place any restrictions on the number of desired solutions. In our case study example, we conducted a sensitivity analysis to the N_{SP} hyperparameter (Appendix C.1) and found that increasing the number of models did not significantly affect the fitness and diversity of the solutions. This showed that, in this case, there are many possible equifinal model representations and that this can have implications for the policy analysis (Appendix C.4). Thus, applications of DMC-RPA should carefully consider the choice of N_{SP}.

Additionally, our approach does not consider the sensitivity of each parameterization. This sensitivity can exist on at least two levels. First, model stochasticity and uncertainty in the input data can both lead to drastically different system behavior under a single model configuration, and our approach does not allow for this variability when evaluating a model configuration, in contrast to probabilistic approaches (Olalotiti-Lawal and Datta-Gupta, 2015). Second, due to the complexity of socio-environmental systems, small parameter or structural changes can massively affect system behavior (Lempert, 2002; Liu et al., 2007); by identifying N discrete models, each defined by a fixed set of parameter/structure values, our approach does not give information about this sensitivity. In our case study example, we conducted a local sensitivity analysis to the resultant model configurations, and found two of them to be highly sensitive to small changes in some of the parameters (Appendix C.5). Thus, we recommend that sensitivity analyses are integrated with future applications of DMC-RPA.

4.6.3 Potential extensions

Our intention has been to demonstrate diverse model calibration for robust policy analysis. We did not conduct computational experiments to compare EAGA to alternative optimization procedures, so we do not claim that EAGA represents the most computationally efficient or appropriate method for diverse model calibration in all contexts. Niching genetic algorithms are an alternative evolutionary approach that model a single population of solutions that evolve based on a combination of feasibility and the density of other solutions in the parametric neighborhood (Miller and Shaw, 1996; Goldberg and Richardson, 1987; Singh and Deb, 2006). Multiobjective genetic algorithms have also been developed in other contexts to identify sets of Pareto-optimal solutions (Deb et al., 2002; Park et al., 2013; Komuro et al., 2006; Turley and Ford, 2009). Alternatively, the characteristics of this approach could potentially be integrated into MCMC-based methods, for example by including diversity requirements as a penalty in MCMC likelihood functions (Olalotiti-Lawal and Datta-Gupta, 2015), prioritizing diversity when sampling from MCMC-generated posterior distributions (e.g., some form of latin hypercube sampling), or through multi-objective MCMC algorithms. Although DMC-RPA is not directly comparable to these other optimization approaches, in particular due to the way it treats the fit and diversity objectives, future work could (1) compare the effectiveness of the EAGA with other adaptive sampling procedures in identifying a set of plausible, diverse model configurations and (2) integrate diversity objectives more explicitly into other calibration methods.

Within the DMC-RPA approach, there also exist promising avenues for future extension. Some of these relate to the evaluation of model fit. For example, for our case study, we combined the loss calculated for each histogram and system-level pattern into a single measure of fit (Equation 4.1). However, it has been demonstrated that such aggregation can inefficiently explore the entire parameter space (Park et al., 2013; Deb et al., 2002). Thus, a multiobjective loss measure may help to identify better calibrations. Additionally, to prioritize the generation of parsimonious model structures, a "complicatedness"-based penalty could be integrated into the objective function (Magliocca and Ellis, 2016).

When experimenting with the algorithm, we noticed a tendency for non-influential parameters to diverge to the extreme ends of the prescribed bounds (i.e., 0 and 1 in Figure 4.7). This is entirely a result of the diversity objective; the algorithm exploits non-influential parameters to increase the assessed diversity of the model configurations without measurably affecting the models' fit to the data. To reduce this effect and focus on diversity where it matters most, we encourage iterative model development (Figure 4.2) integrated with sensitivity analysis (as in Appendix C.5) to sequentially refine the parameters/structures included in the genetic algorithm (Ligmann-Zielinska et al., 2014).

To more comprehensively evaluate the policies' robustness using the equifinal models, the DMC-RPA approach could benefit from tools developed in the RDM literature. For example, the policy analysis could also consider uncertainty related to model inputs, future exogenous conditions, or elements of model configuration that cannot be fixed during the calibration process (Kasprzyk et al., 2013). Efforts could also be made to more explicitly map the calibrated parameter values and structural states to the scenario performance to identify "robust regions" within the configuration space (Lempert, 2002; Bryant and Lempert, 2010).

Finally, in this chapter we have applied DMC-RPA to calibrate an agent-based model using distributional data from a single point in time. However, the approach could be applied to any calibrated process-based model. In the most general sense, it requires simply a model (M) that produces an output (Y) that is dependent on some input parameters and/or structure X (i.e., $Y = M(X)$). Many other types of models (e.g., system dynamics, economic equilibrium, biophysical simulation) in many different fields (e.g., land system science, ecological economics, natural resource management) fit this description. Additionally, the loss function (Equation 4.1) is very flexible; it need not satisfy any statistical properties and only requires that a set of model configurations can be cardinally evaluated according to their level of acceptability. Thus, it would be possible to integrate, for example: (1) timeseries data by either summing or multiplying the discrepancy measure (in Equation 4.1) at each time point (Vrugt and Beven, 2018); (2) spatial features using a measure of landscape pattern similarity (Parker and Meretsky, 2004; Brown et al., 2005); or (3) other levels of uncertainty by either averaging losses over stochastic simulations or assessing robustness over multiple types of uncertainty (Lempert, 2003).

4.6.4 Case study results: Smallholder resilience

Our results to the case study suggest that—under the conditions of the modeled system, and noting the inaccuracy in the models' abilities to recreate the empirical livestock herd size distribution—climate forecasts may provide superior direct benefits to smallholder drought resilience than an increase in non-farm job availability. This result was consistent over all four model parameterizations (Figure 4.10) and over seven of the eight parameterizations in the sensitivity analysis (Appendix C.4). This suggests that dramatic improvements to resilience could be realized without the significant infrastructural investment that would be required to increase non-farm employment opportunities. Thus, "informational" forms of support like climate forecasting could play an important role in supporting smallholder resilience under a changing climate (Vermeulen et al., 2012; Hansen et al., 2019). However, appropriate communication of forecasts and integration into farmer decision-making would be necessary to achieve these benefits in reality (Hansen et al., 2011), in conjunction with adequate accuracy of the climate forecasts themselves (Ziervogel et al., 2005). Future modeling work could more thoroughly represent these elements and/or extend the structural breadth and empirical grounding of the ABM to increase the policy relevance of these results.

Expected utility maximization and the "rational actor" constitute a common approach for representing human decision-making in agent-based models (Kremmydas et al., 2018; Schlüter et al., 2017; Klabunde and Willekens, 2016; Groeneveld et al., 2017). However, this type of approach has long been criticized as not realistically representing how individuals actually make decisions (Simon, 1955). Interestingly, our results (Figure 4.7) showed that models in which agents maximize their expected utility produced better levels of fit to the data than models in which agents behave as "satisficers" that first attempt to ensure their food security, then maximize their utility beyond this (Kaufman, 1990). Because some agents' utility in the ABM is represented by leisure time, even under utility maximization these agents do not behave as maximizers in the traditional economic sense, potentially explaining this result. The DMC approach could be used in future work to evaluate and compare alternative approaches for modeling decision-making based on the degree to which they generate empirically-consistent behavior (Schlüter et al., 2017).

4.7 Conclusions

We have argued in this chapter that, given the prevalence of complex process-based models in socio-environmental policy analysis and the paucity of empirical data with which to calibrate these models for their intended purposes, equifinality is an issue of general concern to this community. We do not claim that process-based models are too sensitive to be useful. Rather, we advocate that modelers seriously consider the implications that model structure and parameterization may have on any model-generated inferences. The DMC-RPA approach that we outline and demonstrate in this chapter can be used to identify policies that perform well over an entire set of equifinal models, thus supporting robust decision making. Alternatively, the approach also can expose inconsistencies that more completely represent uncertainty in the relative benefits of policy interventions. In this case, divergent results may give information about the socio-environmental conditions under which certain policies may be more or less beneficial. In either case, the DMC-RPA approach facilitates more robust model development, policy analysis, and inference.

Acknowledgements

This project was supported by a U.S. National Science Foundation (NSF) INFEWS Grant (1639214) and an NSF CNH Grant (DEB-1617364). The authors would like to thank Allison Kelly and Uriah Israel for their help with the LSMS data and genetic algorithm respectively, as well as the helpful comments from five anonymous reviewers that greatly improved the chapter's quality and clarity.

Chapter 5

Ecological and Financial Strategies[1]

Researchers and development organizations regularly grapple with competing ecological and financial strategies for building climate resilience in smallholder agricultural systems, but rarely are such approaches considered in tandem. Using a social-ecological simulation model, we explore how different combinations of legume cover cropping—an "ecological insurance"—and index-based crop insurance—a "financial insurance"—affect the climate resilience of mixed crop-livestock smallholder farmers over time. The model simulates interactions between soil nutrient dynamics, crop yields, and household wealth, which is carried solely in the form of livestock. We assume cover cropping increases soil quality and productivity over time through biological nitrogen fixation, whereas microinsurance gives payouts in drought years that provide ex-post coping benefits. Our model results indicate that the benefits of cover cropping to mean household income strongly complement the shock-absorbing benefits of microinsurance. Specifically, we find: (1) insurance always provides larger benefits during and in the wake of a drought, while cover cropping progressively reduces poverty in the medium- to long-term; (2) the use of crop insurance solely as an ex-post coping strategy may not reduce the incidence of poverty; and (3) legume cover cropping offers larger relative benefits in more degraded environments and for poor farmers. These results underscore the complementary roles that ecological and financial strategies could play in building resilience in smallholder agricultural systems. The stylized model constitutes an important social-ecological foundation for future empirical research to inform agricultural innovation and sustainable development priorities.

[1]Published as Williams T.G., Dreßler G., Stratton A.E., and Müller B. 2021 "Ecological and Financial Strategies Provide Complementary Benefits for Smallholder Climate Resilience: Insights from a Simulation Model." *Williams, T. G., G. Dressler, A. Stratton, and B. Müller. 2021. Ecological and financial strategies provide complementary benefits for smallholder climate resilience: insights from a simulation model. Ecology and Society 26(2):14. https://doi.org/10.5751/ES-12207-260214*

5.1 Introduction

How to support climate resilience in smallholder agricultural systems remains a topic of uncertainty and debate among researchers and development organizations (Hansen et al., 2019; Tomich et al., 2019a). Institutional interventions such as microinsurance schemes have recently gained traction as tools for agricultural development and poverty reduction in the Global South (Hazell et al., 2010a; SwissRe, 2013; Kramer et al., 2019). Simultaneously, there is an increasing drive for ecological intensification to sustain or enhance both livelihoods and natural resources (Bommarco et al., 2013; FAO, 2018; HLPE, 2019). Such financial and ecological strategies both act as a form of "insurance" by reducing risk in agricultural production, yet they function through distinct mechanisms: "ecological insurance" improves ecological functioning to stabilize and increase production over time, whereas "financial insurance" stabilizes agricultural income on a seasonal basis against climate shocks. Given these distinct mechanisms, ecological and financial strategies may provide benefits for smallholder systems that are heterogeneous both throughout the population and over time. Thus, when considered together, these disparate strategies may be complementary. To make progress toward sustainable development therefore requires an integrated perspective on the benefits of ecological and financial development strategies. In this chapter, we aim to provide a valuable contribution toward this goal by conducting a rigorous comparative assessment of how two particular ecological and financial strategies may affect smallholder climate resilience.

Microinsurance is a form of low-sum financial insurance specifically targeted at low-income households. In recent decades it has gained traction in the international agricultural community as a resilience-enhancing strategy (SwissRe, 2013; Müller et al., 2017; Kramer et al., 2019). By providing financial compensation during droughts, microinsurance directly builds the ex-post coping capacity (i.e., following the occurrence of a shock event) of smallholder households. Additionally, by reducing production risk, microinsurance can provide ex-ante benefits that enable risk-averse households to engage in different production activities and escape poverty traps (Barrett et al., 2007; Carter et al., 2018). Index-based insurance, which gives payouts based on a predetermined climate index (e.g., rainfall) has been advocated as a tool for sustainable development, as it helps to overcome some of the "moral hazard" issues associated with conventional indemnity-based insurance (i.e., the tendency for insured households to reduce their own risk management and increase costs for insurers) (Hazell et al., 2010a).

Farm management practices based on ecological principles take a different approach to smallholder climate resilience. By increasing ecosystem functions and diversity, they provide farmers a form of "natural insurance" (Finger and Buchmann, 2015; Valente et al., 2019; Schaub et al., 2020). In particular, planting of nitrogen (N)-fixing leguminous cover crops to be incorporated into the soil as "green manure" builds resilience by increasing soil organic matter (SOM) and nutrient availability, which help to maintain or increase crop yields over time without other external inputs (Drinkwater et al., 1998; Snapp et al., 2005; Blanco-Canqui et al., 2012; Bommarco et al., 2013). Use of legume cover crops as green manures is receiving increasing attention in the academic literature, from governments, and from non-profit and development organizations advocating for conservation agriculture, regenerative agriculture, and agroecological approaches to smallholder resilience (Florentin et al., 2011; Kaye and Quemada, 2017; Wittwer et al., 2017; FAO, 2018; HLPE, 2019).

Despite their benefits, both microinsurance and legume cover cropping exhibit potential trade-offs that may affect their relative performance. For example, insurance often does not incentivize sustainable management practices (O'Connor, 2013) and may even lead to maladaptive outcomes in socio-environmental systems (Müller et al., 2017). In contrast, adopting legume cover cropping may lead to short-term losses in labor or yields as farmers transition to new management practices and build soil fertility (Martini et al., 2004). The structure of the "payouts" that these strategies provide may also contribute to divergent effects; although both entail annual costs, the ex-post benefits of index-based microinsurance are only experienced during shock years in which the index is triggered, whereas cover cropping provides a more consistent, though likely smaller, economic benefit (Rosa-Schleich et al., 2019). When considered together, it is therefore possible that microinsurance and cover cropping provide complementary benefits (Hansen et al., 2019).

However, it remains a challenge to understand the conditions—–when, where, and for whom—–under which each of these strategies may be most appropriate or beneficial to smallholder climate resilience. A deeper understanding of their benefits can help to inform and target agricultural research and development and contribute to the debate on the relative merits of financial- and ecological-centric development approaches (Tomich et al., 2019b). Given the nascence of research on the impacts of both microinsurance and legume cover cropping on the global agricultural stage, long-term observational datasets do not exist with which to systematically compare their relative or complementary short-term, long-term, and distributional effects. In addition, both strategies involve interactions and feedbacks between household assets and underlying ecological systems, necessitating an integrated social-ecological perspective.

Process-based simulation models are powerful tools for extending the understanding of these relationships and feedbacks beyond existing empirical datasets, as well as exploring changes in conditions and processes that would be impossible to control for in the field (Magliocca et al., 2013). Simulation models that combine social and ecological processes—henceforth "social-ecological simulation models"—–have been extensively used to explore questions related to resilience and smallholder agricultural livelihoods (Kremmydas et al., 2018; Egli et al., 2018; Dressler et al., 2019b). In the context of microinsurance, an agent-based model (ABM) was used to show that there can exist long-term maladaptive feedbacks related to livestock insurance in pastoral systems (John et al., 2019). Models incorporating soil nutrient dynamics have shown that access to credit, fertilizer, and improved seeds can help to reduce poverty but does not guarantee long-term social-ecological sustainability (Schreinemachers et al., 2007). Process-based models have been used to explore the effects of different policies to mitigate N losses (Kaye-Blake et al., 2019) and to assess the emergence of poverty traps (Stephens et al., 2012). Yet, despite the suitability of social-ecological simulation models to investigate short- and long-term tradeoffs and to compare disparate resilience-enhancing strategies across a population, such temporal and distributional effects are rarely studied (Williams et al., 2020a).

For this study, we developed a household-level social-ecological simulation model of a mixed crop-livestock smallholder agricultural system. Rather than being calibrated to a specific location, the model is purposely stylized and represents the general characteristics of many mixed crop-livestock systems in the Global South. As such, the model is intended as a tool for generating hypotheses to be empirically tested by researchers in specific contexts, as well as for illustrating key social-ecological dynamics relevant for informing future interventions, programs, or public policy directed at poverty alleviation.

Using the model, we address the following questions:

1. What are the relative effects of planting legume cover crops as green manure and index-based crop insurance on smallholder households' climate resilience?

2. Are there short- and long-term complementarities in these effects?

3. How do these strategies differentially affect rich and poor households?

In answering these questions, we operationalize the concept of resilience using measures of household wealth and income. In the model, these economic measures are mediated by ecological capital (i.e., soil nutrients). Our perspective is therefore an ecological-economic one. We hypothesize that financial insurance provides greater benefits to resilience in the short-term, but that over time the benefits of cover cropping for SOM will provide equal or superior resilience benefits. Thus, when applied together, the strategies will demonstrate complementarity over time. Additionally, because cover cropping constitutes a progressive ecological adaptation of the agroecosystem, we expect its benefit to be strongest for poor households with degraded soil fertility.

5.2 Methods

Our model description generally follows the Overview, Design Concepts, Details, and Decisions (ODD+D) format (Müller et al., 2013). We provide the full protocol in Appendix D.7. The model was implemented in Python and code is available at CoMSES.net[2].

5.2.1 Model purpose

The social-ecological simulation model was developed to investigate climate resilience in smallholder mixed crop-livestock farming systems, which are prevalent in many drylands regions in the Global South, where crop growth is limited by rainfall (Powell et al., 2004; Thornton and Herrero, 2015). To more easily disentangle the key social-ecological dynamics, we sought to limit model "complicatedness" (Sun et al., 2016). As such, the model does not draw from extensive empirical data to represent a specific location, but we draw several parameters from Ethiopian data sources to define the relative scales of model elements (e.g., crop and livestock prices). We affectionately name the model SMASH: Stylized Model of Agricultural Smallholder Households.

Our model analysis examines the general mechanisms through which selected household-level adaptation strategies affect climate resilience. Due to the model's stylized nature, we do not seek to directly generate policy-relevant recommendations through the model analysis. Rather, our assessment intends to (1) generate hypotheses that can be tested by researchers in future empirical studies and (2) provide theoretical grounding for future agricultural development and poverty reduction programs to integrate ecological and economic adaptation strategies.

5.2.2 Entities, state variables, and scales

The model (Figure 5.1) represents a population of smallholder households that engage in agriculture and carry wealth solely in the form of livestock. Each household is defined by static land holdings and consumption requirements and has dynamic income and wealth. Livestock are grazed on a combination of on-farm crop residues and an external rangeland, which is not explicitly modeled. Each household's land——or "field"—has an evolving level of organic and inorganic nutrients, the dynamics of which influence crop yields. The model is spatially implicit, no environmental feedbacks beyond the household scale are represented, and households do not interact.

5.2.3 Process descriptions

The model operates at an annual time step. Each year of the simulation involves calculation of (1) soil nutrient flows, (2) crop yields, and (3) household income and wealth.

[2] https://www.comses.net/codebases/ee47544a-7eb0-4482-8967-42d6b0c05060/releases/1.0.0/

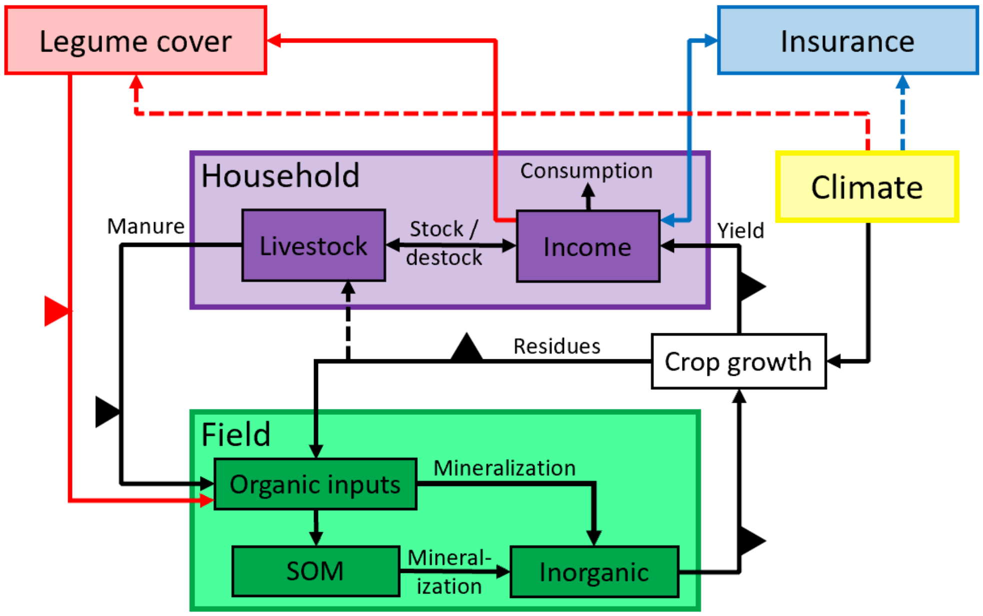

Figure 5.1: Conceptual diagram of the SMASH model showing the main interactions. Triangles pointing inwards (/outwards) indicate points at which nutrients are added to (/lost from) the system.

5.2.3.1 Nutrient dynamics

The model represents two pools of soil nutrients: organic and inorganic. The organic pool represents SOM and soil organic N together in a stylized manner, with fluxes primarily corresponding to the organic N portion of SOM. Although crop yields are also limited by other nutrients, we focus on N because it is generally the most limiting nutrient for crop growth (Robertson and Vitousek, 2009). We henceforth refer to this pool as SOM, though we note that we quantify it using kg N/ha rather than as a percentage of bulk soil. Each year, inorganic nutrients are mineralized from both added organic matter and from the SOM pool (Figure 5.1). These inorganic nutrients are available to that year's food crop.

There are several points at which nutrients enter and leave the system (Figure 1). First, a fraction of the mineralized nutrients is lost through leaching. This fraction is higher with lower levels of SOM (Drinkwater et al., 1998; Bommarco et al., 2013). Second, all nutrients contained in the harvested component of the crop are exported from the modeled system. Third, 10% of the crop residues are assumed to be lost or removed (Assefa et al., 2013). Nutrients enter the system through livestock manure, which qualitatively represents nutrient import from external grazing land. Hence, households with larger livestock herds have larger SOM additions and consistent cropping with no replenishment of SOM will slowly degrade soil fertility over time (Reeves, 1997).

In many mixed crop-livestock systems, households apply inorganic fertilizers to supplement in-soil nutrients for crop growth. However, inorganic fertilizer is not included in this version of the model. Including fertilizer would require additional assumptions about household decision-making related to fertilizer use and livestock nutrient management, as well as complicate the model dynamics. We interpret our results in the light of this assumption.

5.2.3.2 Climate and crop yields

We model crop yields using the yield gap concept, in which yields are reduced from a maximum potential value through water and/or nutrient limitations (Tittonell and Giller, 2013). We first simulate the regional climate condition, which is the same over all households and is independently sampled each year from a normal distribution. Using this, we calculate field-level water reduction factors. Here, field-level SOM helps to reduce drought sensitivity (Bommarco et al., 2013). Next, if the available field-level inorganic N is insufficient to produce this water-constrained yield, production is limited by the available inorganic N. Finally, the resulting value is perturbed by a field-level, normally distributed stochastic error term. This term conceptually represents all uncontrollable factors affecting crop yields and other positive or negative household-level shocks, as well as local variability in the observed climate conditions within a region containing a population of smallholder households.

5.2.3.3 Household income and wealth

The model makes several assumptions with respect to household income and wealth. First, households do not have access to financial savings and instead use livestock as a "bank account". Hence, "wealth" and "livestock" are equivalent in the model. Second, we do not consider non-farm employment markets. Third, households cannot purchase fodder for their livestock under baseline conditions, making livestock a risky wealth stock. These conditions are characteristic of many mixed crop-livestock systems in the Global South (Powell et al., 2004; Thornton and Herrero, 2015), in which livestock are the primary savings mechanism. We interpret our results in the light of these assumptions.

Households have a fixed annual consumption requirement. They earn income solely from harvested crops, which are sold each year at a constant price. If net income is in surplus, households add to their wealth stores by purchasing livestock. If net income is in deficit, households sell the required amount of livestock as a coping measure (Bellemare and Barrett, 2006; Moyo and Swanepoel, 2010). If income is in deficit and the household has no available wealth stores, we assume that they can perfectly reduce their consumption (i.e., wealth cannot be negative, and households do not exit the modeled system). Finally, we do not model livestock reproduction or mortality.

The ability for households to accumulate wealth is constrained by fodder availability for livestock (Valbuena et al., 2012; Assefa et al., 2013); we assume that a fixed percentage of livestock feed requirements must come from on-farm crop residues and that households cannot keep livestock that they cannot feed. Hence, households with larger land area (i.e., producing a greater quantity of crop residue) have larger wealth capacities. Additionally, this implies that in a year of complete crop failure, households lose all livestock that were dependent on crop residues.

5.2.4 Feedback loops

The structure of the model implies the existence of a feedback loop; surplus income enables accumulation of livestock, providing additional organic matter, which both decreases drought sensitivity and increases future crop yields and income. A household's ability to experience this positive feedback cycle is mediated by a combination of random and non-random factors; households' attributes such as land endowment and SOM determine their wealth-generating ability and hence predispose them to certain trajectories. In addition, stochasticity through household-level random yield effects introduces a degree of path dependence into the model; a household that is unlucky one year (i.e., has a large, negative random effect in their crop yields) may be pushed into a poverty trap (Tittonell, 2014; Haider et al., 2018), with decreasing livestock herds, SOM, crop yields, and income.

5.2.5 Calibration and specification of household types

Given our interest in exploring the distributional effects of resilience strategies, we specify the model with three types of household that differ exclusively in their land endowment. We refer to these types as: land-poor, middle, and land-rich. We used pattern-oriented modeling (POM) (Grimm et al., 2005) to estimate values for unknown model parameters that lead to a set of desired emergent model behaviors. To qualitatively represent both chronic and transitory poverty dynamics (Barrett, 2005), we selected baseline parameters such that the land-poor households are "always poor" (i.e., never maintain positive levels of wealth throughout the simulation), the middle households are "sometimes poor," and the land-rich households are "never poor". Additionally, we required that SOM never increases to a maximum value under baseline conditions and that the middle households can recover from shocks. See further details in Appendix D.7.

5.2.6 Resilience-enhancing strategies

We represent both microinsurance and legume cover cropping in the model as scenarios, rather than as an outcome of an explicit decision-making process. Thus, we do not focus on the question of *how* to expand the use of these strategies. Instead, we explore what the potential benefits might be *if* each strategy is taken up, when these benefits may be experienced, and by whom. We therefore assume that households always engage in a given strategy, regardless of their previous experiences or wealth.

We include a representation of index-based crop insurance. A household with insurance must pay an annual premium to participate and receives a payout in any year that the climate condition is below a pre-specified threshold (e.g., the 10^{th} percentile). The payout rate is the same for all households and is equivalent to the crop yield under average climate conditions, assuming a nutrient limitation factor of 0.5. Insurance payouts supplement the households' income and, in contrast to regular income, can be used to buy fodder for livestock. Thus, the insurance de-risks the wealth stock and represents a form of asset protection rather than replacement (Carter et al., 2018). Because we do not model fertilizer or other agricultural production investments, we consider only the ex-post coping effects of microinsurance and not its ex-ante risk-reducing benefits.

Legume cover crops are grown in the fallow season and incorporated into the soil as green manures. Through biological N_2 fixation and production of high-N biomass, green manures provide additional organic N inputs to the soil. Livestock are not grazed on the cover crops. We assume that the cover crops' growth declines under adverse rainfall conditions in the same way as crop yields; thus, in a year with no rainfall, cover crops fail and no N is fixed (Serraj et al., 1999). We assume an annual financial cost equal to the annual cost of insurance. By assuming that the labor required for cover cropping would otherwise be applied to other income-generating activities, this financial proxy for labor is appropriate.

5.2.7 Outcome measures: Poverty reduction and shock absorption

We operationalize climate resilience in two distinct ways. We conceptualize both of these as nested within "development resilience," which describes "the capacity over time...to avoid poverty in the face of various stressors and in the wake of myriad shocks" (Barrett and Constas, 2014). The first measure represents the longer-term capacity of households to avoid poverty (i.e., retain positive livestock holdings) in the presence of climate variability and evolving SOM levels. We refer to this resilience measure as the "poverty-reducing" capacity, R^{pov}:

$$R^{pov} = P(wealth_{t=T_{pov}} > 0) \tag{5.1}$$

where the probability is evaluated over 300 model replications at time T_{pov} (e.g., T_{pov} = 50 years). We conducted a convergence analysis to determine the appropriate number of model replications that ensures our estimates are not strongly influenced by model stochasticity (Appendix D.6). To compare a household's poverty-reducing capacity under cover cropping (CC) and insurance (Ins), we calculate:

$$P(CC \succ Ins)^{pov} = P(R^{pov}_{CC} > R^{pov}_{Ins}) \tag{5.2}$$

where the $\succ$ sign is read as "is preferable to".

The second resilience measure assesses the shorter-term capacity of a household to maintain or increase its income in the wake of a drought. We refer to this as the "shock-absorbing" capacity, R^{shock}. Its measurement requires some explanation. First, we simulate the system under stochastic climatic variability with a single-year "shock" (i.e., drought event) imposed in year T_{shock}. We measure the drought's severity by its percentile in the climate distribution. For example, a 5% drought represents a one in 20-year event. The drought interacts with the model through its effect on food crop and cover crop yields in the same year, as well as any possible insurance payout (Figure 5.1). This can have long-term implications if the household is required to sell livestock, as this both reduces their future buffering capacity and reduces organic N inputs to their field.

To investigate the temporal dynamics of the shock-absorbing capacity, we run experiments that differ across two dimensions. The first dimension represents the point in time at which the shock occurs in the simulation (T_{shock}). Since both strategies (microinsurance and cover cropping) are applied in every year, T_{shock} is equivalent to the amount of time the given strategy has been in use. The second dimension represents the period of time over which the effects of and recovery from the shock are assessed (T_{assess}). Thus, we calculate:[3]

$$R^{shock} = \sum_{t=T_{shock}}^{T_{shock}+T_{assess}} income_t \tag{5.3}$$

To compare the shock-absorbing capacity of a household under the two strategies, we calculate:

$$P(CC \succ Ins)^{shock} = P(R^{shock}_{CC} > R^{shock}_{Ins}) \tag{5.4}$$

To investigate complementarities between the two strategies, we compare the resilience outcomes with both strategies implemented together——i.e., the households engage in both microinsurance and cover cropping and pay the costs for both——against the outcomes of each strategy in isolation. We consider complementarity as a situation in which engaging in both strategies yields additional benefit above that derived from engaging in one strategy alone——either cover cropping or microinsurance——and a tradeoff as a situation in which engaging in both strategies is less beneficial than engaging in a single strategy. Tradeoffs may occur, for example, if the benefits of adding microinsurance to complement cover cropping do not offset the increased cost for the insurance

[3]For simplicity, we do not discount future income to a "net present value". Thus, income in year $T_{shock} + T_{assess}$ is equivalent to income at year T_{shock}. In reality, households may be unwilling (or unable) to forgo short-term losses for long-term benefits. We discuss this in the Discussion (section 5.4).

premiums.

For both measures of resilience, our focus on wealth and income may appear to represent solely economic outcomes and not ecological ones. However, since a household's wealth- and income-generating abilities are mediated over time by SOM, we indirectly incorporate ecological capital into our resilience measures. Additionally, through our dual resilience measurement we combine stability properties with the ability to resist or undergo qualitative changes in structure (Holling, 1973). Thus, a resilient household can both cope with drought-induced disturbance and resist entering a social-ecologically degraded "poor" state. However, as we do not focus on household decision-making or landscape-level processes, we do not consider facets of resilience related to adaptive responses or transformative system-level transitions (Folke, 2016; Walker, 2020).

5.2.8 Simulation experiments

We structure our analysis into four main experiments (Table 5.1). The first and second experiments respectively examine the shock-absorbing capacity (R^{shock}) and the poverty-reducing capacity (R^{pov}) of households under a range of time horizons. In these two experiments, we examine resilience under cover cropping and microinsurance, as well as with both strategies implemented together. In the third experiment, we test how different assumptions about the costs and benefits of the two strategies affect the resilience comparisons (i.e., $P(CC \succ Ins)^{shock}$ and $P(CC \succ Ins)^{pov}$) to identify "robust regions" within the parameter space (Lempert, 2002). Here, we systematically vary the annual costs of both microinsurance and cover cropping, the microinsurance "strike rate" (i.e., percent of years with a payout), and the amount of N fixed by the cover crops. When the microinsurance cost factor is one, the insurance is actuarially fair. A cost factor less than one represents subsidized insurance and a factor greater than one implies net profits to the insurer.

In the final experiment, we explore how the resilience comparisons change under different socio-environmental conditions. To do this, we conduct a sensitivity analysis on the parameters of the model. We employ a meta-modeling approach for global sensitivity analysis (Iooss and Lemaître, 2015) in which we first run our model under a wide range of perturbed parameter configurations and then fit a non-parametric regression model to explain how both resilience assessments change over the perturbed parameter space. From the meta-model we construct a measure of "partial dependence," which describes the relationship between each parameter and the resilience measures as assessed by the meta-model. We describe this methodology in Appendix D.5.

Table 5.1: Simulation parameters under each experiment.

Experiment	T_{pov}	T_{shock}	T_{assess}	Cover crop		Microinsurance		Complem-entarity	Other param-eters
				N_2 fixation (kg N/ha)	Cost factor[2]	% of years with payout[4]	Cost factor		
1: shock absorption	-	1-50	1-15	95[1]	1	10	1	Yes	Baseline
2: poverty reduction	50	-	-	95[1]	1	10	1	Yes	Baseline
3: strategy characteristics	20	10	3	40-200	0.1-4	1-30	0.1-4	No	Baseline
4: socio-environmental characteristics[3]	50	20	5	95[1]	1	10	1	No	Varied

[1] Drawn from empirically measured values in temperate settings (Badgley et al., 2007).

[2] This represents the annual cost of cover cropping relative to the baseline value for microinsurance.

[3] We used different T_{pov} and T_{assess} in this experiment for visual clarity in the plotting. We verified that this does not affect the shape of the relationships.

[4] See Chantarat et al. (2017).

5.3 Results

5.3.1 Model dynamics

Before presenting the results of our main experiments, we first illustrate the representative behavior of the model under three simulations: baseline conditions with regular climate variability (Figure 5.2A), in the wake of a drought (Figure 5.2B), and with the two strategies (Figure 5.2C). To most effectively demonstrate the relevant characteristics of the model, we assess a different time period and different outcome measures in each representative simulation.

First, as specified by the calibration approach, under baseline conditions and regular climate variability, the land-poor households do not earn enough income to satisfy their consumption requirements and so always become poor (i.e., have zero wealth), whereas the middle households sometimes become poor and the land-rich households are never poor (Figure 5.2A). The divergent outcomes for the middle households emphasize the path dependence in the model; all middle households begin the simulation in the same condition, but the randomness in the calculation of crop yields leads to divergent trajectories, particularly when droughts cause some households to either irrevocably lose their wealth reserves or to experience transitory poverty. Households with positive wealth reserves, through external nutrient input from livestock manure, are able to maintain their SOM, but SOM steadily declines for households with no wealth reserves (Figure 5.2A). An imposed drought leads to a decline in wealth that persists for several years (Figure 5.2B). Due to the wealth-SOM feedback in the model, this results in a marginally lower SOM than the

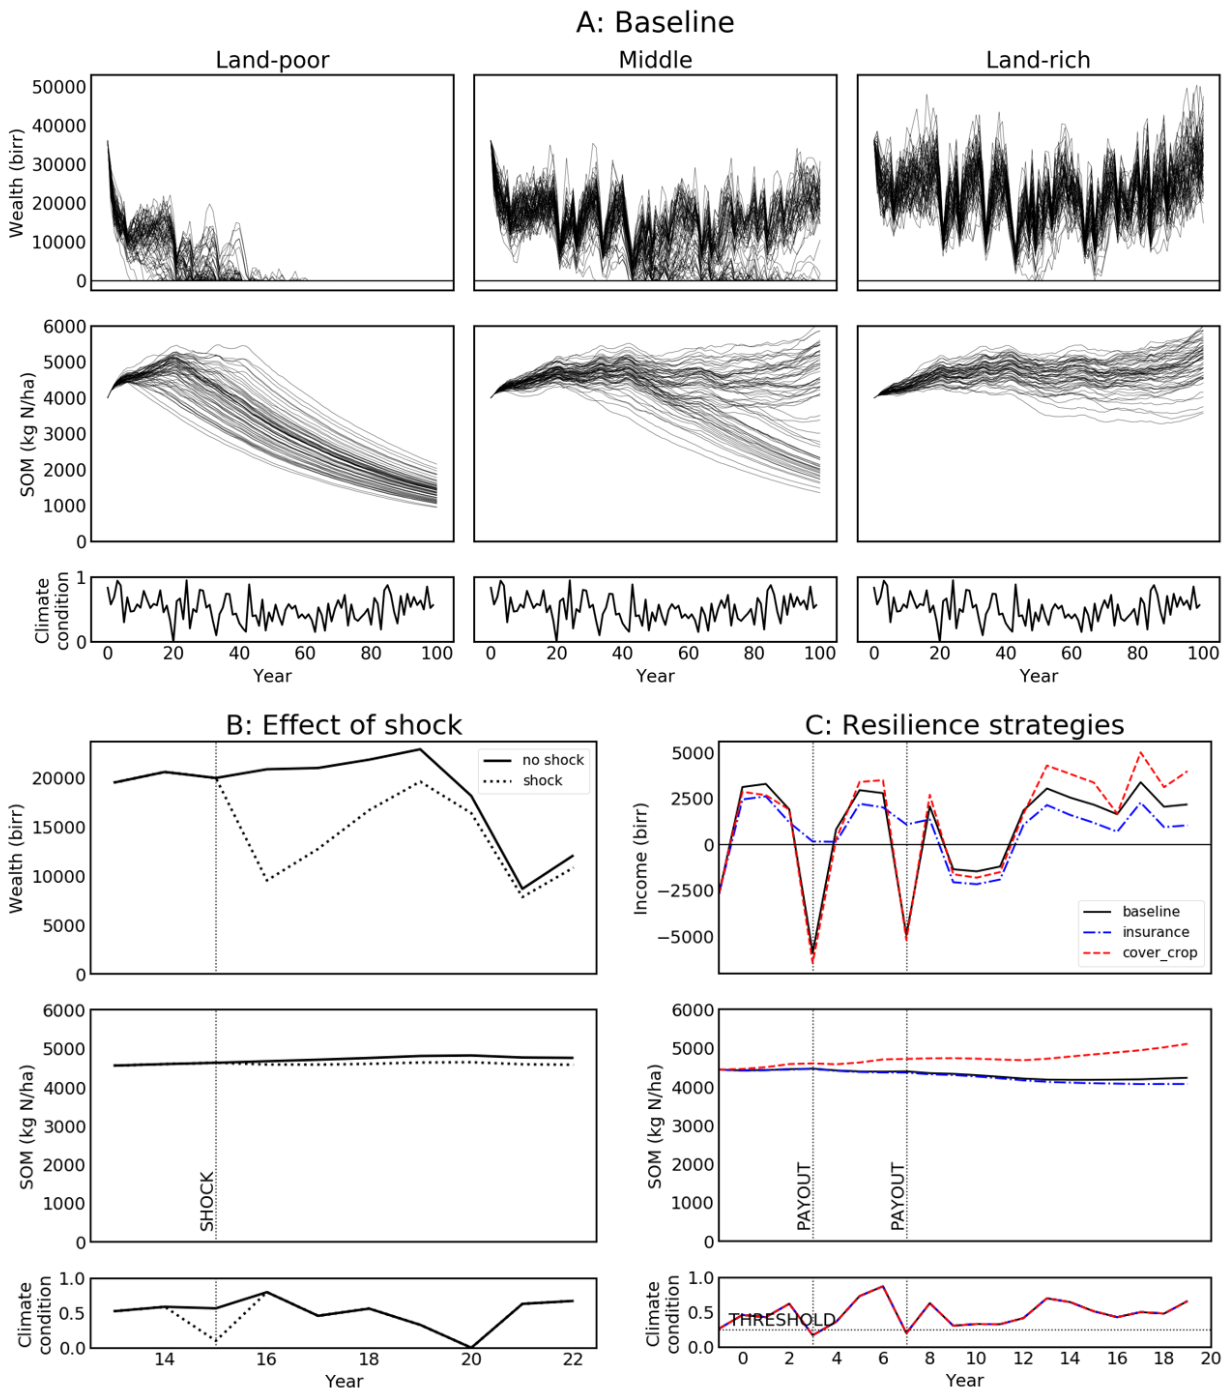

Figure 5.2: Model dynamics under representative simulation runs. "A" shows the evolving wealth and SOM of the three household types under baseline conditions (i.e., without insurance or cover crops) and regular climate variability. Each line represents a single household. Birr is the Ethiopian currency. "B" shows the average effect of an imposed 10% shock in year 15 on the middle household type under baseline conditions. "C" shows the average effect of the two strategies on the middle household type under regular climate variability. The vertical lines in "C" indicate years in which insurance payouts are triggered.

drought-free counterfactual (Figure 5.2B).

Microinsurance and cover cropping affect the model dynamics in several ways. Microinsurance premiums, which cost 10% of average yields, slightly decrease income under regular years, but the insurance payouts effectively buffer the effects of drought when payouts are received (Figure 5.2C). Cover cropping's benefit to income in general increases over time and is strongest in years with higher rainfall (Figure 5.2C). These effects are due to the higher inorganic nutrient availability (from mineralization of cover crop residues) that reduces the extent to which nutrients inhibit crop yields. Because nutrient availability is more critical in high-rainfall years when water is not a constraining factor, the largest benefits are therefore experienced at these times.

5.3.2 Shock absorption

Our results conform with our main hypothesis, showing that insurance as an ex-post coping strategy is preferable in the short-term recovery from a drought, but that there is a time at and beyond which cover cropping provides larger benefits (Figure 5.3). This is not a single point, however, but a line of (T_{shock}, T_{assess}) pairs. When assessing the effects solely in the year of the shock ($T_{assess} = 1$), insurance is the preferable strategy (i.e., $P(CC \succ Ins) < 0.5$) in 100% of the simulations over all time. After 15 years of legume cover cropping, it takes approximately five years following a shock for the cumulative benefits of cover crops to outweigh the benefit of the insurance payout (i.e., transition to red in Figure 5.3). After 25 years of cover cropping, this decreases to three. These effects are qualitatively consistent for each of the three household types (Figure D.3), showing that all types of household strongly benefit from insurance in the wake of a shock. However, when the drought is not severe enough to trigger an insurance payout, cover cropping consistently provides superior shock absorption benefits (Figure D.4).

Due to the strong power of microinsurance in buffering the effects of drought, adding microinsurance to complement cover cropping always increases shock-absorbing capacity (Figure 5.4A). In contrast, adding cover cropping to complement microinsurance leads to tradeoffs in the short-term (black region in Figure 5.4B). This is for two reasons. First, in the year of the drought (i.e., $T_{assess} = 1$), crop yields are constrained by water availability rather than nutrient availability, so cover cropping provides little or no direct benefit to offset its costs. Second, it takes time for cover cropping to build SOM and, consequently, the water retention capacity of the soil. Thus, tradeoffs are stronger when T_{shock} is lower. Nevertheless, as the amount of time for which cover cropping is practiced increases (i.e., as T_{shock} increases), its direct benefits to water retention enabled through higher SOM lead to complementary effects even in the year of the shock (Figure 5.4B). Similarly, as T_{assess} increases, cover cropping provides progressively larger benefits that lead to long-term complementarity. Additional experimentation reveals that the long-term benefits of microinsur-

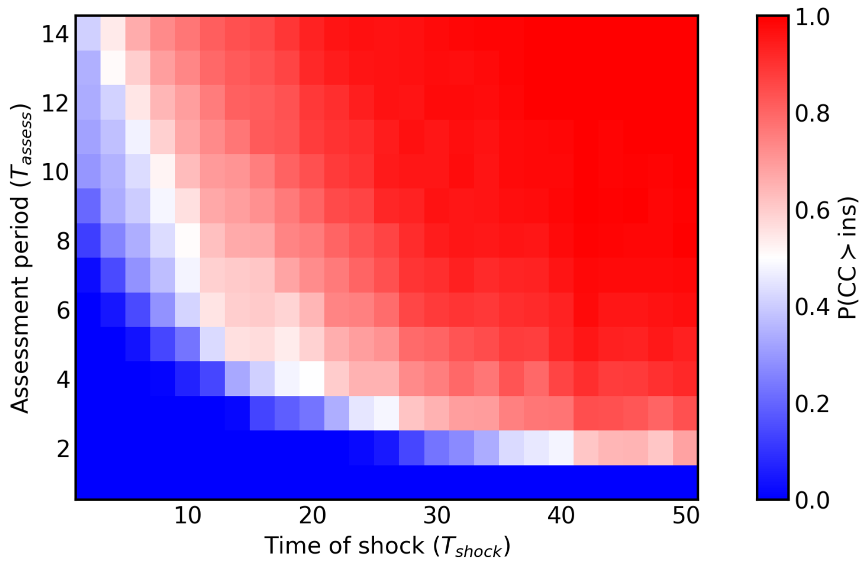

Figure 5.3: Comparison of strategies' shock-absorption benefits. Probability that cover cropping provides larger benefits to shock absorption ($P(CC \succ Ins)^{shock}$) for a middle household type as a function of the year at which the shock occurs (T_{shock}) and the number of years over which the effects are assessed (T_{assess}). Red areas represent situations in which cover cropping provides larger benefits than microinsurance.

ance and legume cover crops are greater than the sum of both strategies in isolation—–i.e., are synergistic (Appendix D.3).

5.3.3 Poverty reduction

Under regular climate variability, legume cover cropping leads to greater poverty reduction than microinsurance (Figure 5.5). The effect is strongest for the land-poor households, who after 50 years of cover cropping are 21% more likely to avoid poverty. For the middle households, cover cropping almost eliminates poverty altogether. These strong effects are explained by the ecological feedback that cover cropping enables; higher SOM increases the productive ability of the households, thus increasing income over time (Figure D.1A). However, there is a 1-2 year period in which the costs of cover cropping outweigh the benefits, resulting in decreased income for all household types (Figure D.1A).

The results show a very different effect for insurance; for both the land-poor and middle households, insurance—–as it is modeled, with ex-post coping benefits only—–is not effective as a poverty alleviation mechanism (Figure 5.5). Despite reducing income variability, the lower mean income in non-drought years due to required insurance premium payments leads to lower mean levels of wealth and SOM (Figure D.1). This demonstrates that although the insurance scheme is actuarially fair, the required premium payments can enable an ecological feedback in the model that results in the payouts in shock years not adequately compensating the income losses in regular

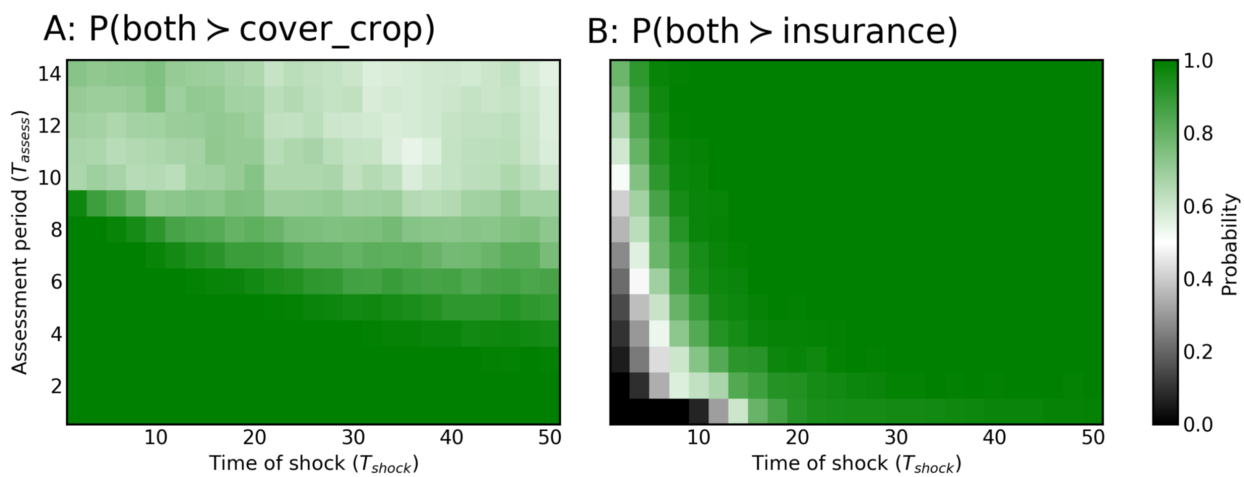

Figure 5.4: Complementarity of strategies for shock absorption. Probability that implementing both strategies together provides greater shock absorption benefit than (A) cover cropping in isolation and (B) insurance in isolation. Green areas indicate complementarity, black indicates tradeoff.

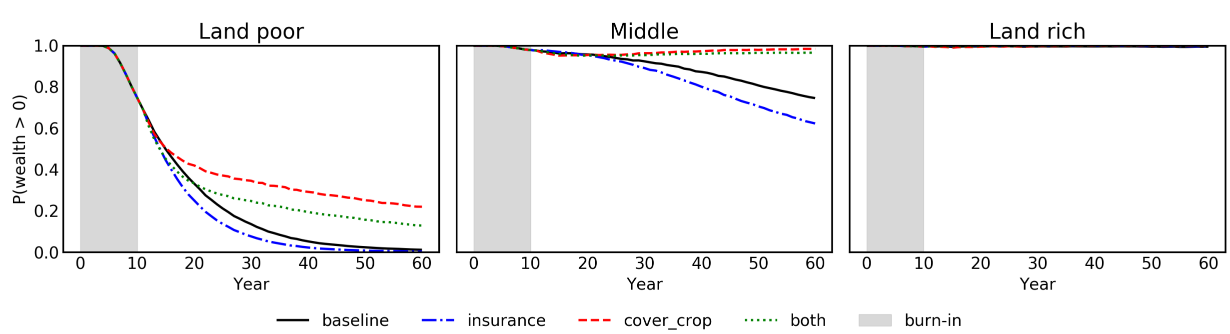

Figure 5.5: The effects of the strategies on poverty reduction. $P(wealth > 0)$ represents the probability that a household has positive wealth reserves over time. We ran a 10-year burn-in period before implementing the strategies to reduce the sensitivity to initial wealth levels (see Figure 5.7C).

years.

With respect to complementarity, for both land-poor and middle households, adding cover cropping to complement microinsurance successfully reduces poverty (Figure 5.5). However, particularly for the land-poor households, the converse is not true; adding microinsurance to complement cover cropping increases poverty above the levels seen with cover cropping by itself. Hence, under the conditions of the model, increasing mean incomes——in this case, through cover cropping——is a more effective strategy for poverty alleviation than reducing income variability.

The measure of poverty reduction assessed in Figure 5.5 is not relevant for the land-rich households, as they are not at risk of poverty under baseline conditions. Supplementary experimentation reveals that, in contrast to land-poor and middle households, microinsurance enables a positive ecological feedback with higher levels of wealth and SOM (Figure D.1). Thus, households not vulnerable to poverty derive some benefit from the reduced income variability provided by microinsurance. To examine this more deeply for a land-rich household, in Appendix D.2 we assess the strategies' effects on a measure of risk-averse utility. Over a range of levels of risk aversion, microinsurance provides welfare benefits to land-rich households. This benefit is initially greater than that of cover cropping, but over time cover cropping's utility benefit surpasses microinsurance's.

5.3.4 Influence of insurance and cover crop characteristics

The superiority of microinsurance for shock absorption is robust to changes in the assumed strategy characteristics (Figure 5.6B and 5.6D). When evaluating shock absorption over a 3-year recovery period, insurance provides on-par or superior benefits to cover cropping up to cost factors of around 2 (i.e., a case in which the annual premium is twice the expected annual payout). Cover crops would need to be both freely available through household production (i.e., cost factor of zero) and fix very high levels of N to provide benefits equivalent to insurance (top-left of Figure 5.6B). When effects are assessed only during the year of the shock (i.e., $T_{assess} = 1$), insurance remains strongly preferable for shock absorption under all conditions in which a payout is received (Figure D.5).

The superiority of cover crops for poverty reduction is also robust (Figure 5.6A and 5.6C). Only at high cover cropping costs and low N_2 fixation rates does insurance become preferable (Figure 5.6A). Similarly, the cost factor for microinsurance generally has to be lower than one for it to reduce poverty more than cover cropping (Figure 5.6C). Interestingly, more frequent microinsurance payouts appear to provide better poverty reduction benefits (top-left of Figure 5.6C). Additional experimentation with the microinsurance payout frequency revealed a tradeoff: providing more regular payouts effectively buffers income losses from moderate shocks, but requires a higher annual premium that leads to increased vulnerability during more extreme shocks even when payouts are received (Appendix D.4).

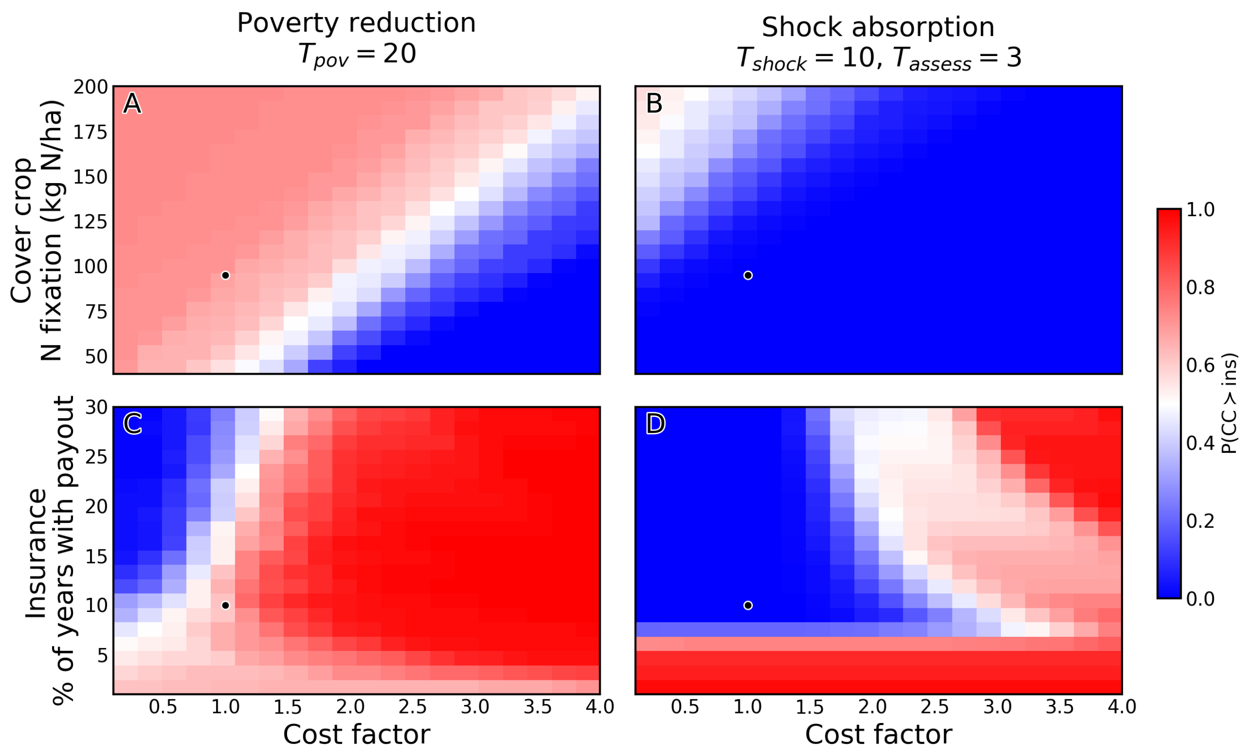

Figure 5.6: Influence of ecological and financial strategy characteristics on the resilience comparisons. The black dots represent the baseline settings used in other experiments. For cover cropping (A and B), the cost factor represents the annual cost of cover cropping relative to the baseline annual cost of insurance. For insurance (C and D), the cost factor represents the ratio of the annual premium to the expected annual payout. When this equals one, the insurance is actuarially fair. The vertical axis for insurance represents the percent of years in which an insurance payout is received. In all cases, we show only the results for the middle household type; additional results are shown in Figure D.5 and Figure D.6.

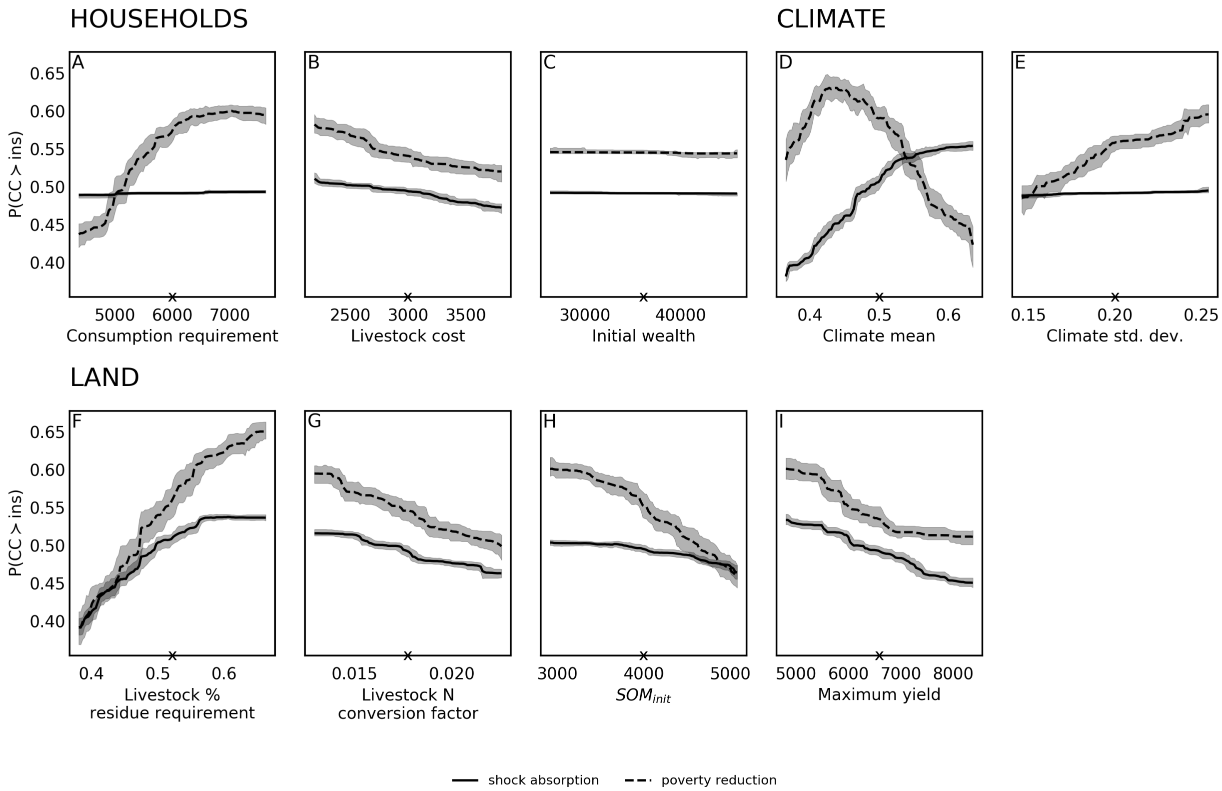

Figure 5.7: Sensitivity of the resilience assessments to changes in model parameters. "x" marks indicate the default parameter values used in the other experiments. Uncertainty bands represent 95% confidence intervals from 100 bootstrapped replications of the model outputs. The method used to generate these plots is described in Appendix D.5. We plot only four "LAND" characteristics, which were selected based on sensitivity and social-ecological relevance.

5.3.5 Sensitivity to socio-environmental characteristics

We use the sensitivity analysis (Figure 5.7) both to assess the sensitivity of the model to its parameters and to draw insights about which resilience-enhancing strategy may be more preferable in different socio-environmental contexts. In Figure 5.7, the slopes of the lines give an indication of the magnitude and direction of the sensitivity of the $P(CC \succ Ins)$ assessments for each parameter. Because this was generated under a single set of settings for T_{assess}, T_{shock}, and T_{pov} (Table 5.1), in this section we are more interested in the slopes of the lines than the absolute $P(CC \succ Ins)$ values.

As consumption requirements (i.e., household living costs) are increased in the model, cover cropping becomes a more preferable strategy with respect to poverty reduction (i.e., the dashed line is upwards sloping in Figure 5.7A). This complements the results of Figure 5.5; higher consumption requirements result in more households becoming poor (Figure D.6A), thus accentuating the poverty-reducing effects of cover cropping and further demonstrating cover cropping's pro-poor benefits. Other household-level parameters do not exert considerable influence on the comparisons

(Figure 5.7B and C), and this low sensitivity provides strength to our results in the above sections.

Changes to the average climate condition have divergent and nonlinear effects on the resilience strategy comparisons (Figure 5.7D). Cover cropping provides the largest relative poverty reduction at moderate climate conditions. This is because under low climate conditions (i.e., low rainfall), cover crops fix less N and so do not provide long-term SOM benefits (Figure D.6B), reducing their relative ability as a poverty reduction strategy. Conversely, with high climate conditions (i.e., more rainfall), more households have livestock and so are able to maintain SOM in their fields without cover crops (Figure D.6B), also reducing cover crops' relative poverty reduction effect. For shock absorption, microinsurance is more beneficial than cover cropping under drier conditions (i.e., lower climate mean). Here, cover cropping more effectively buffers shocks under conditions of higher average rainfall, due to SOM stabilizing yields during the more moderate shocks.

Under higher climate variability, cover cropping provides larger relative benefits to resilience (Figure 5.7E). This is because cover cropping, through building of SOM, moderates the relationship between climate variability and yield variability. Although microinsurance provides payouts when climate conditions fall below the threshold, it does not buffer against climatic variability in non-payout years. Thus, when climate variability is higher, microinsurance has a lower relative benefit on average.

Cover cropping offers larger relative benefits to resilience under more adverse land characteristics—specifically, situations with low external rangeland availability (Figure 5.7F), low soil fertility returns from livestock (Figure 5.7G), low soil fertility (Figure 5.7H), and low yield potential (Figure 5.7I). This result is not surprising, as cover cropping progressively builds the system's natural capital. Relationships are qualitatively consistent between the two resilience measures.

5.4 Discussion

5.4.1 Microinsurance alone may not reduce poverty

Our results suggest that—when used solely as an ex-post risk coping strategy—microinsurance alone may not help households to escape poverty (Figure 5.5; Figure 5.6). The premium payments required for microinsurance pushed poor households into poverty traps, thus increasing poverty relative to baseline conditions. The lack of benefit for poor households highlights potential concerns regarding equity (Fisher et al., 2019) and is in accordance with some empirical research on index-based livestock insurance (Chantarat et al., 2017). In addition, we found that the vulnerable non-poor (i.e., "middle") households also experienced higher poverty levels with the insurance alone. In part, this result is explained by our exclusion of ex-ante effects of insurance that would enable risk-averse households to engage in higher productivity livelihood activities (e.g., fertilizer use, crop choice, and other drought management strategies) (Müller et al., 2011; Mobarak and Rosenzweig, 2013; Karlan et al., 2014; Cole et al., 2017; Kramer et al., 2019). Inclusion of these effects may change the outcomes for the middle households. Nevertheless, the potential for microinsurance to cause vulnerable non-poor households to enter (transitory or chronic) poverty warrants further consideration in models with more complex household behavioral representations (including issues of moral hazard and interaction with other behavioral adaptations (O'Hare et al., 2016)), as well as empirical investigation in different socio-environmental contexts.

5.4.2 Ecologically based farm management enhances resilience over time

The robustness of the relative benefit of legume cover cropping for poverty reduction in our model is largely due to its assumed long-term benefits for agricultural productivity, which enable poor households to "step up" out of poverty (Dorward, 2009). Other production technologies, such as improved crop varieties, cropping system diversification, irrigation, or conservation agriculture practices, may offer similar risk- and productivity-related benefits to cover cropping (Lin, 2011; Hansen et al., 2019). Additionally, other studies have argued for fertilizer subsidies to break soil quality poverty traps (Barrett and Bevis, 2015). Future research could evaluate and compare the resilience effects of such productivity-enhancing technologies and policies.

Importantly, our analysis highlights the value of an integrated social-ecological perspective. Our results show that legume cover cropping—investing directly in soil fertility itself—offers substantial combined potential for long-term environmental improvement and poverty reduction for smallholder farms, which may not exist with non-ecological technologies like inorganic fertilizer. Beyond the modeled effects, ecologically based management strategies offer numerous benefits to field- and landscape-level ecosystem services (Bommarco et al., 2013; Dainese et al., 2019), as well as reduce dependence on external inputs (Shennan, 2008). Reduced externalities and ancillary benefits may be difficult to quantify and slow to build, but ultimately contribute to social-ecological synergies and resilience of a more "general" nature than the "specified" version assessed by our model (Cabell and Oelofse, 2012; Jacobi et al., 2018; Stratton et al., 2020; Weise et al., 2020). Thus, we recommend that future policies, projects, and programs for smallholder poverty reduction empirically examine the benefits of integrated ecological and economic approaches (Müller and Kreuer, 2016; Beck et al., 2019).

Our results revealed a 1-2 year period before cover cropping provided net benefits—i.e., a "transition period" (Martini et al., 2004; Lamine and Bellon, 2009). We did not focus on decision-making or barriers to cover cropping adoption, but these results highlight that liquidity constraints and large time discounting rates could make households unable or unwilling to forgo these short-term losses to engage in cover cropping or similar productivity-enhancing practices (Quaas et al., 2019). Thus, a long-term view may not be pragmatic if focusing exclusively on cover crops. Capacity-building, educational opportunities, and subsidies for cover crop seeds and labor during the transition period may help to overcome this barrier (Baumgärtner and Quaas, 2010; DeLonge et al., 2016; Duff et al., 2017). Integration of dynamic decision-making and interactions with other institutional structures are avenues for future research on ecological resilience-enhancing strategies.

5.4.3 Harnessing ecological and financial complementarities for climate resilience

Our results illustrate the strong complementarity of microinsurance and cover cropping: when implemented together, the strategies can provide greater benefit than either in isolation (Figure 5.4). Climate resilience and poverty reduction programs, development agendas, and empirical studies could further test this complementarity and investigate "bundling" of adaptation strategies (Kramer and Ceballos, 2018; Kramer et al., 2019; Wong et al., 2020). Our study demonstrates the promise of simulation models——whether empirically calibrated to specific locations or stylized as in this study——as tools for ex-ante examination of resilience dynamics and interactions between strategies over long timescales. Particularly in situations where empirical evidence is lacking, simulation modeling can provide important information about time lags, barriers to adoption, and required investments, which can help to inform the design of poverty reduction programs and aid allocation.

Different types of household may require different forms of intervention; our results showed that chronically poor (i.e., land-poor) households benefited greatly from the ecological strategy of cover cropping, which acted as a necessary "cargo net" to mitigate risk and increase asset bases (Barrett, 2005), but that adding microinsurance to complement cover cropping did not provide complementary poverty reduction benefits (Figure 5.5). Thus, risk mitigation strategies such as cover cropping could be emphasized for enabling chronically poor households to step up out of poverty. However, because cover cropping alone did not bring all land-poor households out of poverty (Figure 5.5), bundling with additional interventions, such as social protection measures (Hansen et al., 2019), may be necessary and should be investigated in future research. Bundled cover cropping and microinsurance appears to offer the greatest benefit for the vulnerable non-poor (i.e., middle) and non-poor (i.e., land-rich) households. For the middle households, the bundled strategies reduced poverty by a comparable amount to cover cropping in isolation (Figure 5.5), as well as provided long-term complementarity in the wake of a drought (Figure 5.4). For the land-rich households, particularly those with higher risk aversion, the bundled strategies provided immediate welfare improvements (Appendix D.2).

Environmental context can exert additional influences on the appropriate combination of financial and farm-based strategies. For example, legume cover cropping had a comparative advantage in harsher and more degraded landscapes (Figure 5.7). Yet, annual cover cropping may not be an appropriate agricultural practice in contexts with very low rainfall, as this can limit potential biomass accumulation and N fixation, as well as potentially reduce soil moisture content and subsequent vegetable crop yields (Unger and Vigil, 1998). In these contexts, drought-tolerant cover crops or other sustainable agriculture practices, such as mulching or agroforestry (Shankarnarayan et al., 1987; Ewansiha and Singh, 2006; Bayala et al., 2012), may be more effective—both in isolation and in combination with insurance. Additionally, future case-based studies should target the insurance strike rate to the given social-ecological context (Lybbert and Carter, 2015; Kramer et al., 2019), as context will affect climate-yield relationships, cover cropping performance, and poverty dynamics.

5.4.4 Generalizability of our results

We made several strong assumptions in our model that may influence the generalizability of our results. Most importantly, the wealth-based feedback loop in which wealth (livestock) directly fosters organic nutrient imports and improves crop productivity is critical to our model. In situations where financial resources other than livestock are available (e.g., savings accounts), wealth would not be as strongly linked to field-level nutrient import. Additionally, large areas of grassland may be required to graze livestock to sustain nutrient applications on cropland, which might be infeasible given socio-political constraints on land ownership and access (Dell'Angelo et al., 2017). Furthermore, perfect import of nutrients from rangelands is an optimistic assumption due to competing uses for nutrients (Tittonell and Giller, 2013; Berre et al., 2021). In all cases, the implication is that the wealth-based feedback loop in our model may be exaggerated and thus the strategies' effects on poverty overestimated. However, this exaggeration is the same under each strategy, so by focusing on the relative benefits of the two strategies, we reduce (though do not eliminate) the implications of this bias for our assessment.

Our modeled system most closely approximates an isolated rural community in which non-farm employment opportunities do not exist, use of fertilizer is low, and wealth is constrained by local environmental conditions (i.e., no access to savings accounts or fodder for purchase). Smallholder systems globally are undergoing diverse structural transformations driven by population growth and globalization, leading to increased livelihood diversification both within agriculture and into non-agricultural activities, increased intensification, and commodification and consolidation of land ownership (Barrett et al., 2010; De Schutter, 2011b; Alobo Loison, 2015). Inclusion of such processes would affect our results. For example, including inorganic fertilizer as another mechanism to increase productivity would likely diminish the relative benefits of cover cropping, though fertilizer does not directly build SOM. Moderate fertilizer application and cover cropping could therefore be complementary practices (Giller et al., 1997). Non-farm employment opportunities may help to increase smallholder resilience under baseline conditions by providing a means through which the poor can step out of poverty (Hansen et al., 2019). Additionally, households may be willing to buy fodder to smooth their asset stocks even at the expense of their own consumption (Morduch, 1995), which would reduce the effects of drought on asset stocks seen in our results. Future research could expand the scope of this stylistic model to include additional livelihood activities, behaviors, or exogenous drivers and better match it to specific empirical contexts.

Our study focused on potential benefits if support systems existed such that smallholders were able to adopt legume cover cropping and microinsurance. We did not incorporate household decision-making with respect to uptake of the strategies or their spillover effects on other management practices. In reality, there exist financial, social, and informational barriers to the adoption of both ecological and financial strategies that have led to limited uptake in smallholder systems to date. Integrating decision-making and approaches from ecological economics with the resilience perspective in this chapter is a promising avenue for future research.

5.5 Conclusions and recommendations

We assessed the effects of microinsurance and legume cover cropping on climate resilience in a stylized mixed crop-livestock smallholder agricultural system. Our study offers a fresh, reconciliatory approach to the current debate on pathways toward climate risk management and poverty reduction (Hansen et al., 2019). Distinct agricultural development communities and organizations advocate for microinsurance and ecologically based management, sometimes with strong ideological disagreements. By providing a rigorous comparative assessment of these strategies, we hope to bring these paradigms together, illuminate their complementarity, and seed future collaborative empirical assessments and integrated applications to programs and policies for sustainable development.

Our model results can essentially be boiled down to this: insurance provides an important buffering effect to climate shocks, while legume cover cropping progressively decreases poverty and the impacts of shocks over time. Together, these benefits underscore the potential complementarity of economic and ecological adaptation strategies for smallholder resilience. Future development programs and empirical research could test this complementarity in different socio-environmental contexts, including how it develops over time and throughout a heterogeneous population of households. Finally, development resilience provides a useful conceptual framework for quantitative resilience analyses that jointly considers the capacities for poverty reduction and shock absorption (Dou et al., 2020). An integrated approach to resilience assessment shows promise to mitigate tradeoffs and harness complementarities so as to improve smallholder livelihoods and social-ecological functioning.

Chapter 6

Smallholder Agency in Large-Scale Land Acquisitions

Large-scale land acquisitions (LSLAs) are a means for agricultural intensification in low- and middle-income countries, but are frequently implemented by dispossessing smallholders of their land, thereby generating tradeoffs between market-oriented production and smallholder livelihoods. Given that LSLAs remain a prevalent global phenomenon, the question remains of how to reconcile these conflicting goals. For this chapter, we examined the potential effects of contract farming (CF), an arrangement compatible with LSLAs that preserves some smallholder land rights, on smallholder food security and regional productivity. To do so, we developed an agent-based model of mixed crop-livestock smallholder livelihoods and calibrated it using household survey data collected in four LSLA-affected areas of Ethiopia. Using the model, we examined smallholder livelihood adaptations under different LSLA/CF implementation conditions and their effects on household food security and regional productivity. Our results point to the importance of supporting smallholder land rights and hence autonomy over CF participation to avoid the tradeoffs generated by LSLA-induced displacement. CF schemes that retained smallholder land ownership within the LSLA led to positive food security outcomes and comparable regional productivity increases. Further, allowing smallholders to choose to participate in the CF scheme (i.e., retaining smallholder autonomy) resulted in stronger accordance between the two outcomes and larger food security benefits, particularly for poorer households. To realize these benefits, it is imperative to ensure contract compliance to maintain smallholder trust in the contracting firm. Going forward, our results highlight that LSLA governance, implementation, and monitoring could aim to foster conditions that empower smallholder land rights and autonomy and thereby contribute to sustainable development.

6.1 Introduction

Recent years have seen rising levels of commodification in smallholder agricultural systems. One important contributing mechanism is "large-scale land acquisitions" (LSLAs), where domestic or international actors (states and/or private investors) purchase or lease large tracts of land for agricultural development, primarily in the Global South (Deininger and Byerlee, 2011). The pace and scale of contemporary LSLAs are inarguably massive, and are estimated to impact up to 90 million hectares globally (Müller et al., 2021). Although proponents argue that LSLAs can benefit rural livelihoods by closing yield gaps and prompting infrastructure investments and technology spillovers (von Braun and Meinzen-Dick, 2009; World Bank et al., 2010; Rulli and D'Odorico, 2014), ample empirical evidence reveals negative impacts on adjacent populations, frequently due to dispossession of smallholder land and livelihoods (a.k.a. "land grabbing") (Oberlack et al., 2016; Dell'Angelo et al., 2017; Schoneveld, 2017; Nanhthavong et al., 2021). To reduce the future risks of LSLAs, we need a more nuanced understanding of mechanisms for addressing potential tensions between smallholder wellbeing and market-oriented agricultural development. In this chapter, we focus in particular on contract farming (a smallholder-centric approach to market-oriented agricultural production) to assess how it may affect smallholder wellbeing and productivity in the context of LSLAs.

Contract farming (CF) describes an arrangement in which a buyer of agricultural produce (henceforth, "firm") forms contractual relationships with individuals or collectives of small-scale growers ahead of production for a specific crop (Little and Watts, 1994). The precise contractual arrangements are incredibly diverse (Bellemare and Lim, 2018), but often guarantee a fixed buying price and require the small-scale producers to sell specified levels of production from the contracted land to the firm. By providing smallholders with access to commodity markets as well as agricultural inputs, credit, and technology, CF addresses many challenges facing smallholders in the Global South (Wiggins et al., 2010; Barrett et al., 2012). A growing base of empirical evidence shows that, although context matters, CF can effectively increase smallholder income and food security relative to traditional subsistence arrangements (Prowse, 2012; Bellemare and Bloem, 2018; Ton et al., 2018; Meemken and Bellemare, 2020). In the context of LSLAs, CF is most relevant in places with considerable existing smallholder agriculture. It allows for increases in market-oriented production while retaining smallholder land ownership and management and therefore has been praised as a desirable middle ground between traditional smallholder agriculture and large-scale agribusiness (De Schutter, 2011a; Chamberlain and Anseeuw, 2019).

Contract farming also poses risks to smallholder livelihoods. Several mechanisms are relevant. First, and particularly pertinent in the context of land acquisition, it is possible for smallholders to retain land ownership but be forced to participate in a CF scheme, thus losing autonomy in decision-making. This has been observed in Ethiopia, where households forced to participate in a sugarcane outgrower program had significantly lower income and asset stocks than non-participants (Wendimu et al., 2016). The term "contract farming" therefore can encompass a variety of levels of smallholder autonomy. Second, even when smallholder autonomy is retained, power imbalances can limit accountability if the firm breaches the contract (e.g., does not accept the production at the agreed price), particularly when smallholder side-selling of crops is difficult due to limited alternative output markets (Barrett et al., 2012). In these contexts, smallholders' trust in the contracting firm can be an important determinant of their decision to participate (Nguyen et al., 2019). Third, higher labor requirements and production costs for contracted crops (e.g., sugarcane) may trade off against other livelihood activities that traditionally support food production and income diversification (Bellemare, 2018; Ragasa et al., 2018; Bottazzi et al., 2018). This can exclude resource-constrained farmers from the benefits of participation, thereby generating inequitable effects on livelihoods. Fourth, as CF schemes often involve production of non-food crops for export markets (Oya, 2012), productivity benefits may not improve local food security.

Given these risks, there is a need to better understand the conditions under which CF may be a desirable LSLA governance strategy by promoting dual benefits to smallholder wellbeing and market-oriented production. This is a necessary step toward better LSLA governance and an important area for research (Agrawal et al., 2019). For example, identifying smallholder-inclusive tenure arrangements and LSLA structures can jointly help to inform decision-making of value chain actors (e.g., contract conditions set by contracting firms) and state actors regulating land governance and investment (Debonne et al., 2021). However, identifying these conditions ex-ante is difficult due to the many mechanisms at play, coevolving outcomes, and heterogeneity of smallholder populations.

Process-based simulation models are a promising tool to extend beyond the range of empirical data and mechanistically compare alternative governance arrangements. They have been extensively applied for agricultural policy assessment (Kremmydas et al., 2018), yet are relatively underutilized in the context of LSLAs. Previous household-level models of LSLAs have primarily focused on "representative" households (Baumgartner et al., 2015; Kleemann and Thiele, 2015; Schuenemann et al., 2017) and thus do not incorporate interaction or other behavioral dynamics. Such factors are particularly important for CF, where heterogeneity, trust, and learning are critical determinants of smallholder decisions and outcomes (Barrett et al., 2012; Nguyen et al., 2019).

Ethiopia, the empirical context for our study, has experienced one of the largest rates of LSLA in Africa (Schoneveld, 2011). Similar to other countries, empirical studies in Ethiopia have found negative effects of LSLAs on smallholder livelihoods and food security (Shete and Rutten, 2015; Wendimu et al., 2016), little evidence of technology spillovers from large farms (Ali et al., 2019), as well as both gendered (Hajjar et al., 2019) and ethnic (Moreda, 2015) differentiation in LSLA impacts. In contrast to more highly forested regions targeted by LSLAs in places like South America and Southeast Asia, Ethiopia's dominant agro-pastoral lands hold potential for LSLAs to be implemented using CF schemes. Despite this, little CF has been reported in Ethiopian LSLAs to date. There is therefore an underexplored potential for CF to alleviate some of the concerns of LSLAs in Ethiopia, as well as similar contexts targeted by LSLAs globally.

We leveraged household survey data collected from four LSLA-affected regions of Ethiopia to develop and calibrate an agent-based model (ABM) of smallholder livelihoods. We applied the ABM to examine how regional productivity and household food security (measured as access to a single, representative staple crop) may change under different LSLA and CF arrangements. Motivated by the tradeoffs that CF may pose, we focus on the following questions:

1. How do LSLA/CF arrangements that retain different levels of smallholder autonomy over land management differentially affect regional productivity and household food security?

2. How are the effects on food security distributed between richer and poorer households?

3. How does contract breaching by the firm affect households' trust and thereby mediate the regional effects?

Given the global scope and predominance of negative implications of LSLAs thus far, our analysis provides an important step toward reconciling the goals of market-oriented "development" and smallholder livelihoods in the Global South. Our household-level focus and assessment of multiple outcomes enables us to identify synergies and pro-poor governance arrangements. We use the model——empirically grounded in the Ethiopian context——as a virtual laboratory to examine how effects might be distributed under different LSLA/CF arrangements, as well as to generate hypotheses that could be tested in future empirical studies.

6.2 Data and methods

6.2.1 Site selection and empirical data

We focused our assessment in Oromiya (OR), Ethiopia. Within this context, we selected four LSLA-affected "sites," which we label as OR1-OR4 (Figure 6.1). These sites were selected to represent the diversity of land uses targeted by LSLAs in Oromiya: OR1 solely on smallholder agriculture, OR2 solely on shrubland, and OR3 and OR4 on a mix of agriculture, shrubland, and forest. Tenure changes included both transitions from previous state farms to foreign ownership (OR2 and OR3) and LSLAs enacted by the Ethiopian government and managed by domestic actors (OR1 and OR4) (Table 6.1). In none of these cases was a CF scheme implemented. Given the

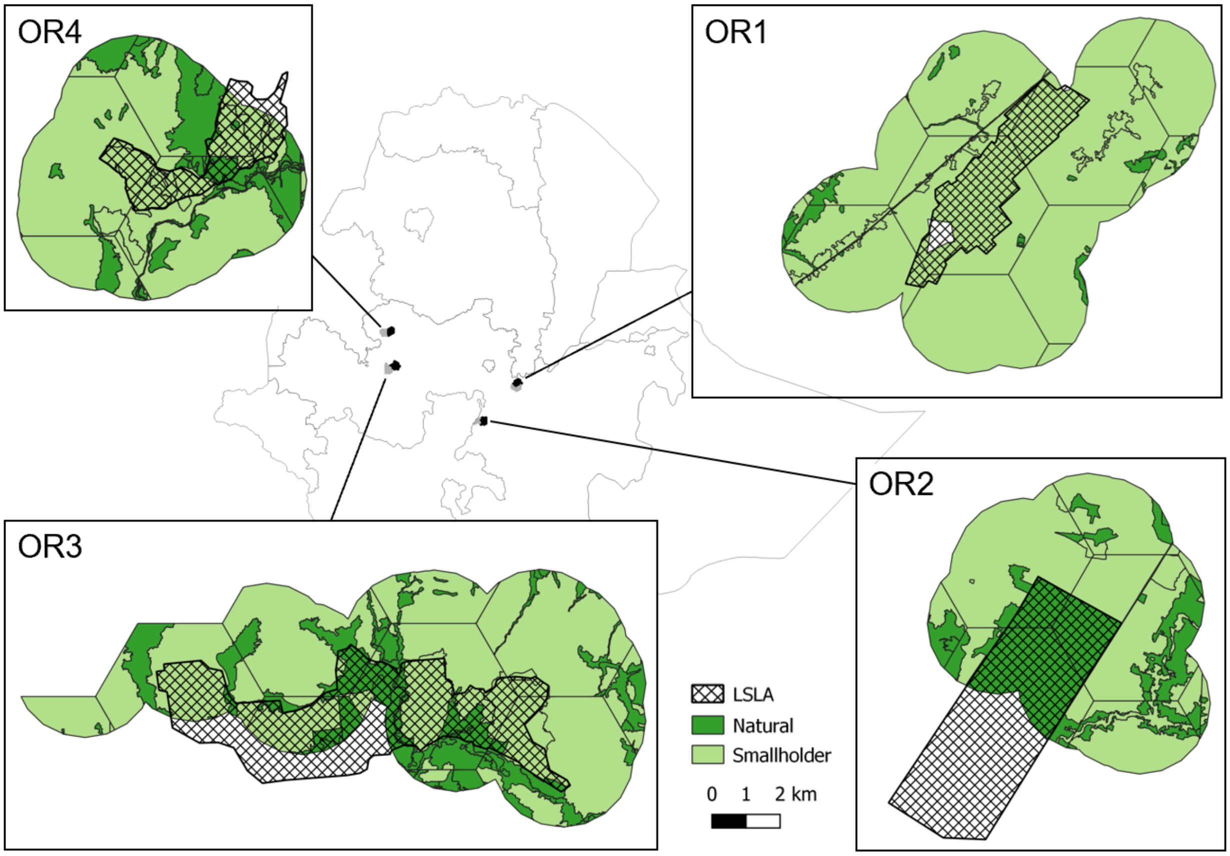

Figure 6.1: Location and pre-LSLA landcover of the four sites within Oromiya (OR), Ethiopia. "Natural" land represents woodland/shrubland, grassland, and forest, which we assume to be communally managed. The displayed data represent 2km buffers around the surveyed households, which are not shown to protect privacy.

different land use histories, we expected the levels of smallholder displacement and livelihood effects to vary between sites.

We conducted approximately 100 household surveys at each site, asking questions about agricultural management and yields, livestock holdings, income, labor allocation, and food security. Across all sites, most households pursued mixed crop-livestock livelihoods. Agriculture was rainfed and the primary crops were maize, teff, and beans, which were grown for both subsistence and sale to market. Some households engaged in non-farm employment, but this was not a dominant livelihood activity, where crop sales provided the majority of most households' income.

6.2.2 Agent-based model

A complete model description following the ODD+D (Overview, Design concepts, Details, and Decisions) protocol (Grimm et al., 2006, 2020; Müller et al., 2013) is provided in Appendix E.6. The purpose of the model is to understand and support generalized conclusions about how alternative LSLA and CF configurations may affect household food security and regional productivity

Table 6.1: LSLA information and effects on land cover within a 2km buffer of surveyed households.

Site	LSLA information				Land cover					
					LSLA		Smallholder agriculture		Common land	
	Ownership	Year	Previous use	Crop	% of site	% in smallholder ag.	% of site	% lost to LSLA	% of site	% lost to LSLA
OR1	Domestic	2012	Smallholder	Sugarcane	12	100	97	13	3	0
OR2	USA	2012	State farm	Maize	20	0	65	0	35	56
OR3	India	2003	State farm	Sugarcane	17	19	82	4	18	76
OR4	Domestic	2008	Forests	Maize	15	30	77	6	23	45
ORX*	-	2010	-	-	15	Variable	80	Variable	20	Variable

* This is a synthetic representative site created for the subsequent simulation experiments.

within mixed crop-livestock smallholder systems.

The agents in the model represent a population of smallholder households (Figure 6.2; (Frelat et al., 2016)). Each agent manages a fixed amount of agricultural land and an evolving herd of livestock, which they graze on a communal rangeland and on their own and others' crop residues. The agents interact directly with their neighbors in forming their beliefs about climate, prices, and the availability of off-farm employment, as well as indirectly through competition for fodder in the rangeland and jobs in the off-farm employment markets. All agents have equal access to the rangeland and employment markets (i.e., they are not spatially explicit), though each agent has a distinct set of neighbors for sharing beliefs and crop residues.

The simulation proceeds at an annual time step. The main model processes are shown in Table 6.2. Principal model outputs include region-level agricultural production and average household food security status. Regional production is calculated as the sum of all agents' crop production and production within the LSLA, which is influenced by the exogenous climate as well as agents' fertilizer decisions. Household food security emerges from a combination of model inputs (e.g., household-level landholdings) and dynamic processes (e.g., receiving non-farm employment). Region-level food security is the probability of food security across all agents. We note that the food security measure represents staple food availability and does not imply that food secure households have access to adequate nutrients to meet their dietary needs.

We used the empirical data to initialize, calibrate, and validate the ABM (see Appendix E.4). These procedures verified the ABM's ability to generate sets of empirical patterns from each of the four sites, as well as approximate the effects of the LSLAs on food security. However, given the similarities between the sites and their small number, we utilized a single Oromiya-level ABM structure and parameterization to represent all four sites for our main experiments. To account for the uncertainty introduced by the calibration procedure, we identified six plausible model calibrations and estimated scenario outcomes with each plausible model (Williams et al., 2020b).

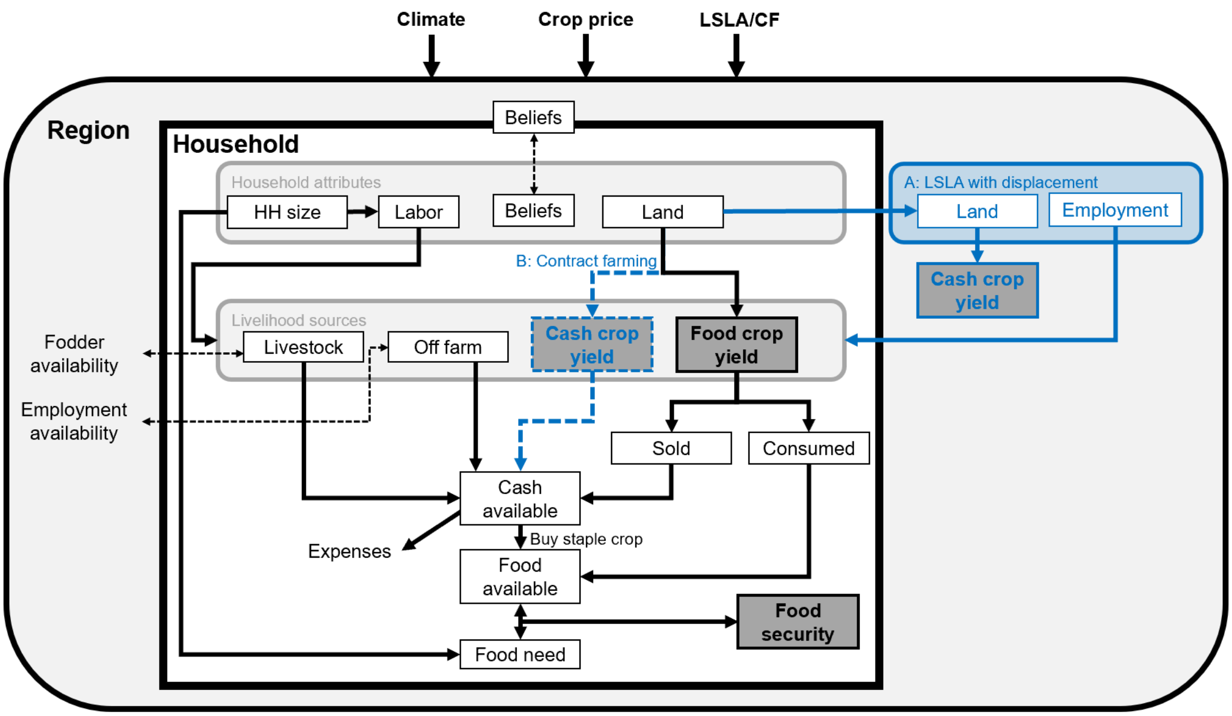

Figure 6.2: Overview of the exogenous drivers, region-level interactions, and household livelihood activities within the ABM. Black boxes and lines represent the baseline subsistence structure. LSLA with displacement (A; blue solid lines) takes smallholder land and employs smallholders on the large-scale farm. Under contract farming (B; blue dashed lines), smallholders produce cash crops on their own land. Grey shaded squares represent the primary model outputs: food security and agricultural productivity.

Table 6.2: ABM process overview and scheduling. Processes are listed in the order in which they are executed at each annual time step. Principal model outputs are in bold. More detailed descriptions are included in Appendix E.6.

Process	Description
Update model environment	• Simulate regional climate condition. • Simulate regional crop prices (food crop and cash crop).
LSLA/CF implementation	• Only run this process in the LSLA/CF implementation year. • See details in section 6.2.3.
Agent decision-making	• Under baseline conditions, decision options include all feasible combinations of fertilizer purchase (Y/N), invest savings in stocking livestock herd (Y/N), and off-farm salary employment (no change, decrease, increase). • When agents have the choice to participate in the contract farm, their decision set is expanded to incorporate the CF participation decision (no change, increase by one field, or decrease by one field). • The spatial resolution of the agents' agricultural land is 0.25 hectares, conceptually representing a field. • The decision options are evaluated over each agent's uncertain beliefs about in climate, crop prices, and the probability of finding off-farm employment. • Agents pursue an explicit objective, which is a risk-averse utility function. • Utility is calculated on anticipated wealth levels (cash and livestock) after attempting to satisfy food requirements.
Allocate regional salaried employment	• There is a limited regional availability of off-farm employment. • Agents can retain salaried jobs over consecutive years. • Jobs available in the market are randomly allocated between agents that seek them.
Crop yields	• Households grow a single staple crop. • Crop yields are calculated using the "yield gap" concept (van Ittersum et al. 2013), in which yields can be limited by water and/or nutrient availability. • Under baseline conditions, crop water is provided exclusively by rainfall (i.e., no irrigation) and is homogeneous each year across all agents. • Households have heterogeneous levels of soil productivity that are held static throughout the simulation. • Nutrients available for crop growth depend on an agent's soil productivity and their inorganic fertilizer application. • We proxy an additional effect of labor on crop yields, in which households with low labor availability relative to their landholdings experience yield reductions. • Finally, the calculated yield is perturbed by a random error (household and time dependent). • Crop yields within firm-operated land (i.e., within the LSLA) are calculated using the same procedure, assuming a fixed fertilizer application rate and no labor constraints. • **Total regional crop production** is calculated as the sum of all agents' crop production and (if relevant) the production within the LSLA.
Income, food consumption, and food security	• Households have an annual non-food expenditure requirement and an annual food consumption requirement for a single food product, which they can satisfy through their own crop production as well as purchase from the market. • The buying price is higher than the selling price, representing transaction costs. • Agents that are unable to meet their consumption requirements through their livelihood sources are classified as "food insecure". • **Regional food security** is calculated as the probability of food security across all agents. • Note: because we model the production and consumption of a single, staple food crop, our measure of household food security represents staple food availability and does not imply that food secure households have access to adequate nutrients to meet their dietary needs.

Process	Description
Livestock grazing, reproduction, and stocking	• Each head of livestock represents a large ruminant and has a fixed annual food consumption requirement. • There is no purchase of fodder. • Livestock consumption is preferentially met through on-farm crop residues. • If a livestock consumption deficit remains, livestock are then grazed on others' leftover crop residues and then on communal grassland, which produces a fixed amount of biomass at the regional level. • Livestock that cannot be supported through these mechanisms are destocked without compensation. • Each year, each animal has an exogenous probability of reproducing. Animal age and sex are not modeled.
Agent coping measures	• Food insecure agents can reduce their food consumption to a limited extent. • If the severity of food shortage is larger than this, agents resort to two coping measures to (attempt to) make up the deficit. • First, agents seek wage-based off-farm work. There is a limited availability of wage employment, which is allocated randomly between agents on a (pseudo-)daily basis. • Second, if food insecurity remains, agents destock from their livestock herds.
Update agent beliefs	• Beliefs are represented probabilistically and are updated using Bayesian conjugate priors with agents' and their neighbors' experiences.

6.2.3 LSLA/CF scenarios and simulation experiments

We designed our simulation experiments to address the three questions outlined in the Introduction. For all experiments, we used the six region-level ABM calibrations with household inputs pooled from the four sites and landcover inputs from the synthetic representative ORX site (Table 6.1). We ran each simulation for 30 years and 30 replications (Appendix E.3).

6.2.3.1 LSLA/CF scenarios and smallholder autonomy

We constructed three LSLA/CF scenarios that represent increasing levels of smallholder land management autonomy (Figure 6.3):

A $Displacement$. This represents the observed situation (i.e., no CF), where the LSLA displaces existing land uses and employs people on the firm-operated farm.

B CF_{forced}. Households within the transacted area are forced to participate in the CF scheme. Affected households must cultivate a non-food crop on all land within the LSLA, using inputs provided by the firm, and sell all production to the firm (Figure 6.2).

C CF_{choice}. No land acquisition occurs. Smallholders can choose their level of participation in the CF scheme each year.

Within these three LSLA/CF scenarios, we examined a variety of implementation and contract configurations (Table 6.3). Because our empirical data did not contain CF, these dimensions were selected to represent variation documented in empirical literature (Oya, 2012; Bellemare and Lim,

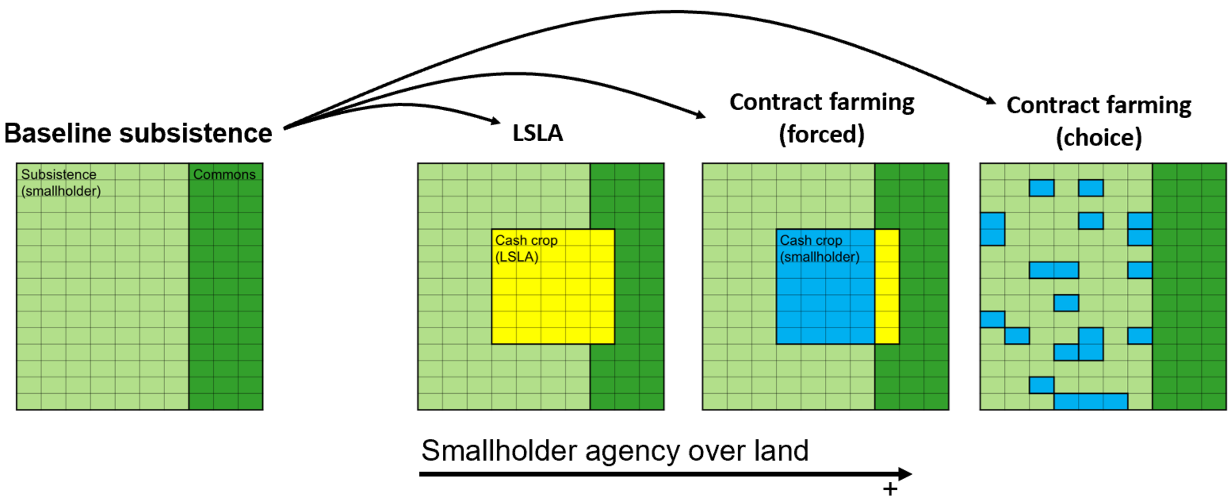

Figure 6.3: Conceptual illustration of land use under baseline conditions and the three LSLA/CF scenarios of increasing smallholder autonomy. Yellow represents firm-managed cash crop production. Blue represents smallholder-managed cash crop production.

2018) as well as factors relevant to agricultural productivity and food security. We ran the ABM for each valid combination of these factors, which yielded 4,644 unique simulation conditions. We measured the effects on regional production and average agent food security relative to baseline subsistence conditions. We also examined which implementation characteristics most strongly influenced the outcomes in each scenario.

6.2.3.2 Distributional impacts

To examine distributional effects on household food security, we stratified the agents into quintiles using their cash income under baseline subsistence conditions.

6.2.3.3 Contract breaching and trust

Within the CF_{choice} scenario, we allowed the firm to breach the contract.[1] We assume that the firm has an exogenous "trustworthiness," i.e., probability of honoring each contract. If the firm breaches an agent's contract, the agent must sell their crop at the subsistence market price (Nguyen et al., 2019) and incurs some lost production (representing, for instance, losses in transportation or marketing). Each agent has a trust in the firm (i.e., trustworthiness belief), which they update at the end of each year based on both their own and their neighbors' experiences with the firm (Michelson, 2017). Trust factors into the agents' decision-making by reducing the expected returns

[1] In reality, smallholders sometimes can breach their contract by "side selling" their produce to alternative markets (Barrett et al., 2012). For simplicity, we do not focus on this behavior in our analysis.

to CF. For this experiment, we systematically varied the firm's trustworthiness and smallholders' initial trust in the firm.

Table 6.3: Set of LSLA implementation conditions. All conditions are exogenous to the ABM. Bold values denote the default settings used for model calibration and in subsequent experiments. 'x' symbols represent the relevant scenario(s) for each implementation condition.

Condition	Settings	Scenario			Description
		$Displa-cement$	CF_{forced}	CF_{choice}	
Implementation	0%, 50%, **100%**	x	x		When the LSLA is not fully implemented, displacement occurs but the benefits (through employment or increased productivity) are not experienced.
Fraction smallholder †	0, 0.5, **1**	x	x		The fraction of smallholder agricultural land in the LSLA (0=no agricultural displacement, 1=full displacement). Equivalent to "% in smallholder ag." in Table 6.1.
Employment	0, **0.15**, 0.3 jobs/ha	x			Employment is available to all households in the region, including those that are not displaced.
Land taking type	Random field, **random percent**	x	x		This describes the allocation of the site-level land-use change between the agents. Random field: select fields at random. Random percent: randomize the agents, then take a random percent of each agent's land until the overall land-use change is met.Note: random percent is the default option used in the model calibration as it better approximates the available empirical data.
Irrigation	**False**, True	x	x		Irrigation removes the water limitations on crop yields.
Intensification	1, **1.5**, 2	x	x	x	Fertilizer application for cash crop relative to baseline subsistence conditions.
Price premium	1, **1.5**		x	x	Average selling price for cash crop relative to food crop.
Harvest period	**1**, 2 years		x	x	With a two-year harvest period, agents receive the cumulative crop yield at the end of the second year.
Labor requirement	1, **1.5**		x	x	Agricultural labor requirement for cash crop relative to food crop.
Production costs	1, **1.5**		x	x	Fixed costs for production of cash crop relative to food crop.
Land requirement	**0.25**, 0.5 ha			x	Minimum amount of land in a contract. Each field is 0.25 ha.
Trustworthiness	0.5, 0.75, **1**			x	Probability that the firm honors each agent's contract in each year.
Production losses	0, **0.25**, 0.5			x	Fraction of production lost when the firm does not honor the contract.
Initial trust in firm	0, 0.5, **1**			x	Agents' initial expected trustworthiness belief.

†The site-level values from Table 6.1 were used for model calibration.

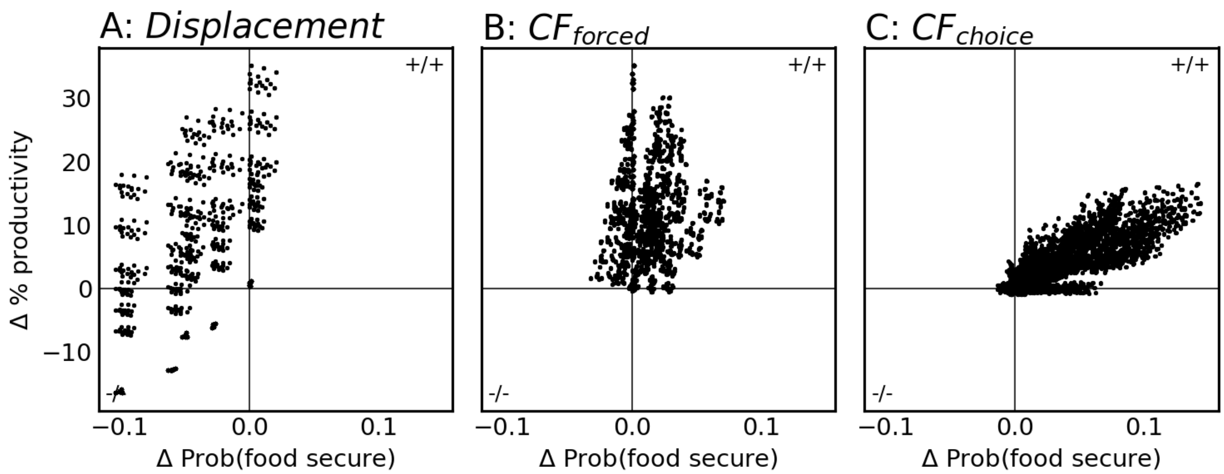

Figure 6.4: Spread of region-level food security and productivity outcomes under the three LSLA/contract farming (CF) scenarios, relative to baseline subsistence conditions. Each point represents the mean region-level output under a single model calibration and set of LSLA/CF implementation conditions (from Table 6.3). All six model calibrations are included in this plot and did not generate discernably different patterns of effects.

6.3 Results

6.3.1 Smallholder autonomy and LSLA/CF characteristics

The three LSLA/CF scenarios generate distinct patterns of effects on regional food security and productivity (Figure 6.4). LSLAs involving displacement ($Displacement$) lead to the strongest tradeoffs—–the "paradox" of LSLAs (Müller et al., 2021)—–with increased regional productivity at the expense of smallholder food security. Involving smallholders within the acquired land as outgrowers (CF_{forced}) provides comparable productivity increases while better supporting smallholder food security. CF_{forced} also reduces the risk of decreasing regional productivity, as observed the *Displacement* scenario. Finally, allowing smallholders to decide to join the CF scheme (CF_{choice}) generates the largest potential food security benefits (up to $\sim$14% increase), but lower maximum productivity increases (up to $\sim$17% increase) than the other scenarios. Under this CF_{choice} arrangement, the food security effects are more strongly correlated with regional productivity gains than in the other arrangements (Figure 6.4), suggesting a closer alignment of firm and smallholder preferences.

Within each LSLA/CF scenario, the effects are mediated by a number of implementation characteristics (Figure 6.5). Under the arrangements involving LSLA ($Displacement$ and CF_{forced}), high levels of implementation in non-smallholder land generate the largest productivity increases. This is because these conditions bring non-agricultural land into production (i.e., entail agricultural expansion; Figure 6.3). In contrast, the CF_{choice} scenario does not allow for agricultural expansion and so does not provide as large benefits to regional productivity (Figure 6.4). Under CF_{choice}, the

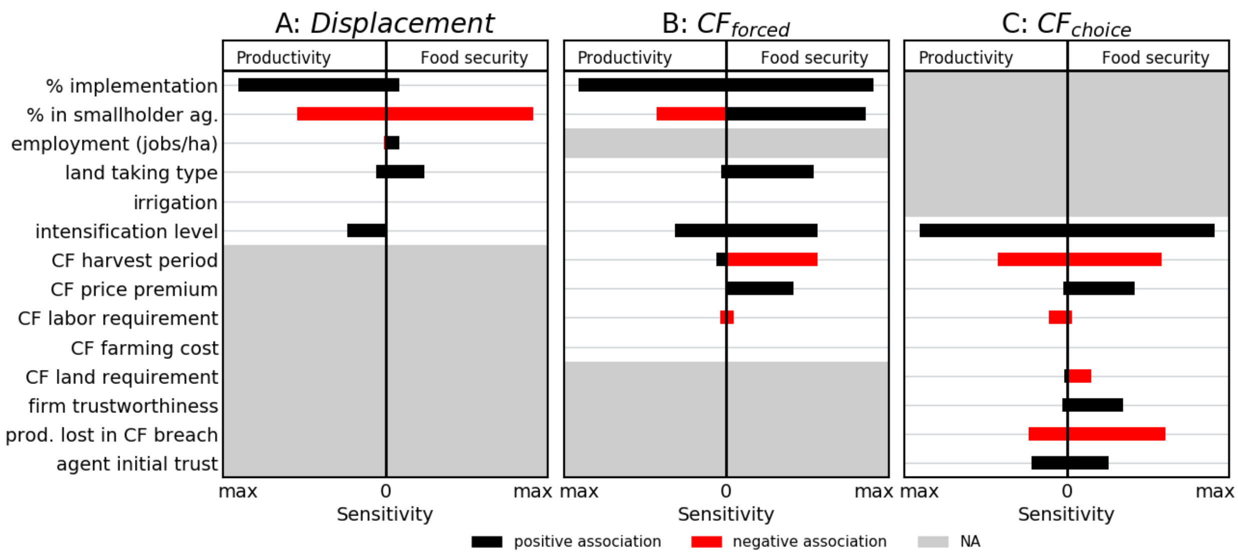

Figure 6.5: Sensitivity of the food security and productivity outcomes to the LSLA/contract farming (CF) implementation characteristics. The sensitivity for each characteristic x is calculated as $E[Y(x = H) - Y(x = L)]$, where E is the expectation operator over all model replications and calibrations, Y is the outcome, and L and H are respectively the characteristic's lowest and highest settings (Table 6.3). Sensitivity values are scaled between zero and their maximum value for each outcome. Black and red bars represent positive and negative sensitivity, respectively.

level of intensification exerts the largest effect on regional productivity (Figure 6.5C).

For smallholder food security, the amount of displacement in the $Displacement$ scenario (through "% in smallholder ag.") is much more important than the amount of employment offered by the LSLA (Figure 6.5A). This implies that LSLA-based employment does little to offset the negative effects of displacement. Within the CF_{forced} scenario, the largest food security benefits are experienced when the LSLA is fully implemented in smallholder land with high levels of intensification, large price premiums, and using a crop with a single-year harvest period (Figure 6.5B). The same general patterns are seen for food security under the CF_{choice} scenario, where both firm trustworthiness and agents' initial trust also support better outcomes.

For some implementation characteristics, the direction of sensitivity changes between LSLA/CF scenarios. This is most clearly illustrated by the cash crop harvest frequency ("CF harvest period" in Figure 6.5). Within the ABM, cash crops with a two-year harvest period (e.g., sugarcane) increase income variability and thereby increase food insecurity in non-harvest years. This effect is particularly strong for households that are forced to contribute a large fraction of their land (Figure E.3). Under CF_{forced}, the harvest frequency has little influence on the regional productivity (Figure 6.5B). However, when smallholders can choose to participate (CF_{choice}), a two-year harvest period is associated with smaller regional productivity increases (Figure 6.5C), as fewer smallholders decide to join (Figure E.2). In this case, allowing the smallholders to choose to join internalizes the negative effects of the two-year harvest period to the firm.

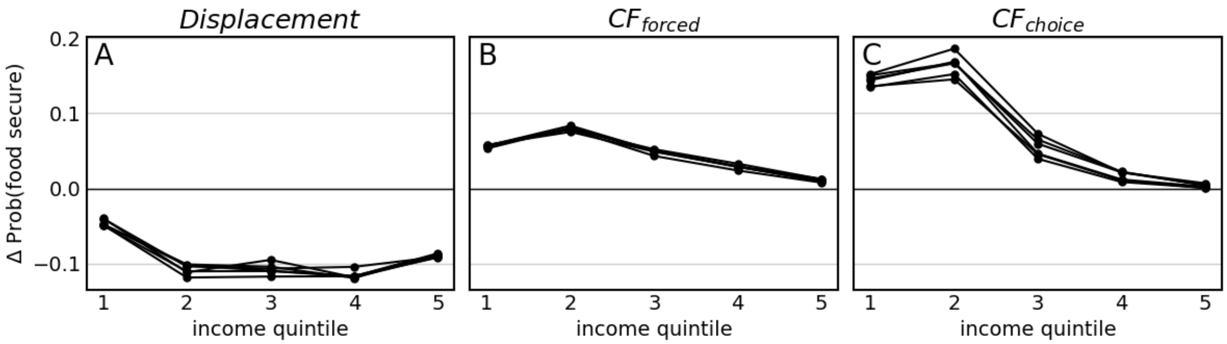

Figure 6.6: Distribution of food security effects under each LSLA/contract farming (CF) scenario, relative to the baseline subsistence conditions. Households are grouped into income quintiles under the baseline subsistence conditions, where quintile 1 contains the poorest households. Each line plots the mean response under a different model calibration.

6.3.2 Distribution of effects

Under the default settings (Table 6.3), the $Displacement$ scenario negatively affects food security across the entire population (Figure 6.6A). Yet, the poorest households are least strongly affected. This is for two reasons. First, due to the empirical data, approximately 45% of these agents do not own any land and so are not directly affected by the LSLA-induced displacement. Second, the poorer agents, with less land on average and hence a larger dependence on off-farm income, benefit from the employment generated by the LSLA. Further inspection reveals that levels of employment above 1.0-1.5 jobs/ha begin to offer net-positive effects on average (Figure 6.7A), which are experienced primarily by the poorest group of households (Figure 6.7B).

Under the scenarios involving CF, the food security effects are non-negative across all quintiles and the CF_{choice} scenario yields strictly better effects than CF_{forced} (Figure 3.10B and C). In both scenarios, poorer households generally experience larger benefits. This is for two reasons. First, poorer households have higher food insecurity under baseline conditions and so there is more room for improvement. Second, the richest households are less likely to join the CF scheme (Figure E.2) because they already experience high utility under baseline conditions and so the relative utility of joining is not as high. Yet, the lowest income quintile does not benefit from CF as much, which is due primarily to the land constraints described above (i.e., 45% of this group are landless, so cannot participate in CF).

6.3.3 Trust and contract breaching

Varying the firm's trustworthiness and smallholders' initial trust reveals distinct patterns of sensitivity; moderate declines in smallholders' initial trust do not substantially affect the food security outcomes, whereas small declines from perfect firm trustworthiness have comparatively large ef-

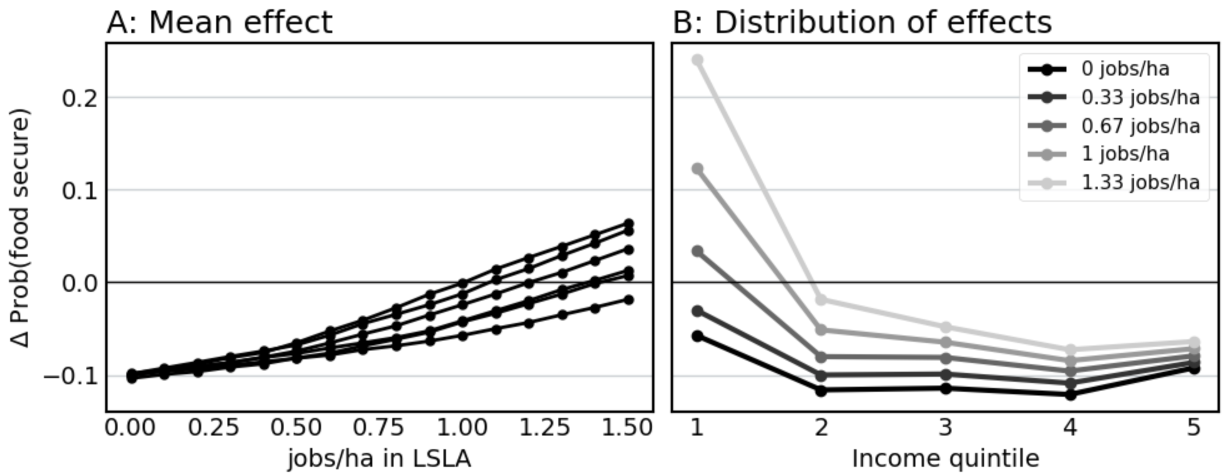

Figure 6.7: How much employment is required to offset the negative effects of displacement? Effect of the LSLA on food security in the $Displacement$ scenario, under different levels of LSLA employment. (A) shows the effect over all agents, where each line represents the mean response under a different model calibration. (B) shows the effect for each income quintile, where each line represents the mean over all model calibrations.

fects (Figure 6.8). The lower sensitivity to declines in agents' trust is because trust is repaired over time, assuming the firm consistently honors the contract. In contrast, moderate reductions in the firm's trustworthiness generate a feedback between food security and trust: by the model's design, contract breaches degrade both food security and trust, consequently reducing smallholders' subsequent likelihood of participation in CF. At very low levels of trustworthiness, food security outcomes are worse than under baseline conditions (Figure 6.8). This is due to the mismatch between the firm's trustworthiness and smallholders' initial trust, which causes agents to still enter the CF scheme despite it being welfare-reducing.

6.4 Discussion

6.4.1 Supporting smallholder autonomy for pro-poor development

The main message emerging from our analysis is the potential benefits of supporting smallholder land rights and agency in LSLA policy and design. By shifting the locus of decision-making toward the smallholder, either by implementing the LSLA as a CF scheme (CF_{forced} scenario) or through CF with no land acquisition (CF_{choice}), we observed larger food security benefits as well as greater alignment of food security and agricultural production outcomes. Similar results have been observed in other modeling studies (Arndt et al., 2010; Baumgartner et al., 2015; Schuenemann et al., 2017) and underscore a need to reflect on the balance of agency smallholders have over agricultural land management (Debonne et al., 2021; Preston et al., 2015). In line with our

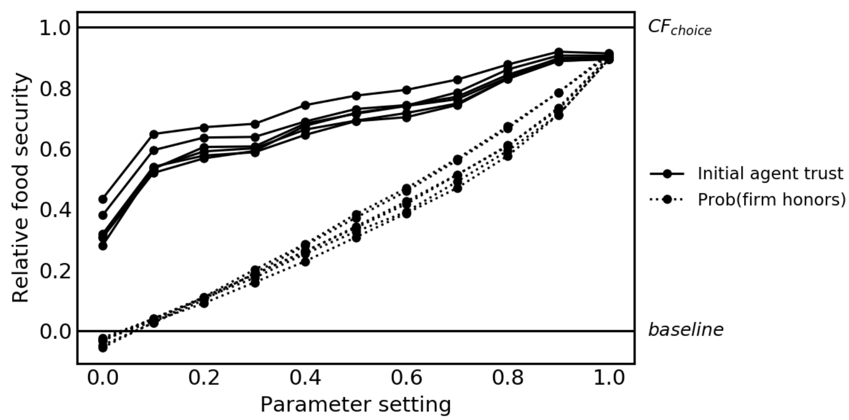

Figure 6.8: Impact of varying smallholder trust in the firm (solid lines) and the firm's trustworthiness (dashed lines) on regional food security. Outcomes are scaled between the baseline subsistence arrangement (the lower bound, conceptually representing ubiquitous contract breaches) and the CF_{choice} arrangement with perfect trust (the upper bound). The smallholder trust experiments (solid black) assume that the firm always honors the contract. The firm trustworthiness experiments (dashed black) assume an initial smallholder trust level of 1. Each line plots the mean response under a different model calibration.

results, empirical evidence has separately shed light on the tradeoffs generated by LSLAs with displacement (Müller et al., 2021), the risks of forcing smallholders to participate in CF (Wendimu et al., 2016), and the potential welfare benefits of opt-in CF schemes (Meemken and Bellemare, 2020). Our study is the first to investigate this continuum within a consistent framework.

The hypothesis that higher smallholder agency leads to better alignment of smallholder and firm outcomes merits further scrutiny in both modeling and empirical research. Future simulation-based studies could expand the scope of smallholder decisions—–such as input intensity (Bellemare and Lim, 2018), side-selling of production (Barrett et al., 2012; Nguyen et al., 2019), or other forms of resistance to LSLAs (Hall et al., 2015)—–to understand the importance of agency *over what* in affecting CF outcomes, as well as how social-political context conditions smallholders' reactions to LSLA/CF institutions (Tamura, 2021). Beyond this, deviations in the firm's behavior may undermine the empirical plausibility of the win-win outcomes we observed. Because smallholders will not accept a contract that they expect to decrease their welfare, allowing smallholders the choice to join theoretically incentivizes large-scale actors to support smallholder welfare. Yet, contract breaching is a pertinent concern in CF schemes worldwide (Barrett et al., 2012), and our results demonstrate the potential feedbacks between contract breaching, smallholder trust, and smallholder food security. Accountability mechanisms therefore need to be designed and enforced, in conjunction with arrangements, such as cooperatives, for distributing power more equitably between parties (De Schutter, 2011a).

Our results also show that opt-in CF schemes (CF_{choice}) could increase regional productivity, though the maximum increases were smaller than under conditions of land acquisition (Figure

6.4). This is because productivity increases in CF_{choice} occurred solely through agricultural intensification and not expansion, as in $Displacement$ and CF_{forced} (Figure 6.3). Beyond these results, empirical research on outgrower schemes in Ethiopia has found higher crop yields on smallholder-managed fields than factory-operated fields (Wendimu et al., 2017), and more generally there is convincing evidence of higher productivity of small farms (Ricciardi et al., 2021). Therefore, there could be large productivity benefits to opt-in CF and enhanced smallholder agency without the need for agricultural expansion. Yet, in any case, there are financial costs and risks involved in contracting with a large number of independent smallholders (Barrett et al., 2012) and production volumes must be sufficient to ensure firm profitability. Future work could expand the modeling scope to include the firm's profits (Nguyen et al., 2019) to understand how profitability constraints mediate the results observed in this study.

6.4.2 Alternative approaches to transformation of smallholder systems

Both LSLAs and CF have the potential to play important roles in the broader processes of rural development and structural transformation. Although our model is not a predictive tool, the levels of employment required to offset the negative effects of displacement (Figure 6.7A) are considerably larger than previous empirical estimates, which extend to at most 0.7 jobs per hectare (Arndt et al., 2010; Deininger and Byerlee, 2011; Baumgartner et al., 2015). A similar result was observed in another simulation-based study of LSLAs (Hailegiorgis and Cioffi-Revilla, 2018). Thus, LSLAs involving displacement likely need to be coupled with other forms of rural development or value addition to succeed as a pro-poor development strategy (Nanhthavong et al., 2021).

Contract farming, in contrast, offers a more smallholder-centric approach to structural transformation (Wiggins et al., 2010; De Schutter, 2011a). Yet, this pattern introduces some unique considerations. First, our results show that households contributing only a fraction of their land benefit the most from CF (Figure E.3). CF schemes should therefore allow market-oriented agriculture to coexist with subsistence production and traditional livelihoods (Bellemare, 2018; Bottazzi et al., 2018). Second, although CF improved the average household's food security in our model, land-poor households could not join or benefit. CF does not solve problems of limited land access, which even in the absence of LSLAs are increasingly prevalent in the Global South (Anseeuw and Baldinelli, 2020). Third, CF does not preclude other forms of dispossession and accumulation from occurring, for example through elite capture (Oberlack et al., 2016). As we did not include such processes within our model, our results potentially overestimate CF's benefits to the poor.

All of our scenarios represent transitions toward agricultural intensification and integration with commodity markets. However, such approaches to development have been criticized for furthering integrating into exploitive market-based economies (Borras and Franco, 2010; Martiniello,

2020) and legitimizing agribusiness under the guise of smallholder inclusivity (Oya, 2012; Oliveira et al., 2020). CF can result in other kinds of losses to smallholder agency, leading to relations of dependence (Oliveira et al., 2020) and compromising food sovereignty (Moreda, 2018). Moreover, agricultural intensification can undermine underlying ecosystems and contribute to environmental degradation (Singh, 2002; Debonne et al., 2021). Future empirical and modeling work could expand the scope of processes and outcomes to contrast rural development through intensification with alternative paradigms, such as sustainable intensification or agroecology (Matson et al., 1997; Anderson et al., 2019), and contribute to broader debates about land sparing versus land sharing (Vongvisouk et al., 2016).

6.4.3 Generalizability of findings

Our model is a simplification of LSLA and CF processes, and does not include all the social-ecological complexity that qualitative and ethnographic approaches tend to articulate (Oya, 2012; Tamura, 2021; Oliveira et al., 2020). A "win-win" outcome as observed in our model therefore does not imply that effects are strictly positive, especially for dimensions not included in our simulation. With this in mind, we have used our model results to identify hypotheses and patterns that warrant further empirical and model-based investigation (Schlüter et al., 2019a) and for assessing generalizability across contexts.

We modeled a mixed crop-livestock smallholder system. This type of system represents at most two thirds of the areas targeted by LSLAs globally (i.e., densely populated croplands and moderately populated shrub- and grasslands (Messerli et al., 2014)). The question remains as to how our results generalize to LSLAs targeted at forested landscapes. In such contexts, processes other than displacement of agricultural land (e.g., deforestation and agricultural expansion) dominate (Debonne et al., 2018) and the LSLA/CF scenarios we tested are not as directly relevant.

Beyond this, the ABM necessarily excludes processes that may be important in some real-world agricultural systems. For example, due to lack of data on decision-making processes, we assumed that households are risk averse and seek to increase their income, subject to meeting food requirements. Other factors may dominate decision-making in some places, such as social norms and pressures. Additionally, our model did not include land degradation, which is a major challenge facing smallholder agriculture globally (Sanchez et al., 1997). CF could either exacerbate soil degradation issues (Debonne et al., 2021) or help to prevent natural resource-based poverty traps (Barrett and Bevis, 2015; Haider et al., 2018), thereby affecting the relative benefits of CF over time. Smallholder agricultural systems face many other challenges that often intersect with the drivers and effects of LSLAs (e.g., climate change (Franco and Borras, 2019)). We did not attempt to include such processes, but instead flag these as additional opportunities for model extension and future empirical research.

6.5 Conclusions

There is a need to better align the development preferences of large-scale actors with the wellbeing of smallholder populations. LSLAs represent a globally pertinent conflict between these objectives (Müller et al., 2021). In this study, we sought to examine the potential for contract farming (CF) as an alternative implementation of LSLA to generate outcomes beneficial to both regional productivity and smallholder food security. We developed an agent-based model of smallholder livelihoods and calibrated it using household survey data from Oromiya, Ethiopia, a region representative of many areas targeted by LSLAs globally. Our simulation-based approach enabled us to cover a large set of experimental conditions and to examine the distributional effects of LSLAs and CF across the smallholder population.

Our results show that CF offers potential to simultaneously increase agricultural productivity and support livelihoods in mixed crop-livestock smallholder systems. In particular, arrangements that gave smallholders greater agency over their land led to better food security outcomes that aligned more closely with productivity increases. CF therefore should be seriously considered as an alternative to forms of intensification by dispossession enacted by LSLAs. Nevertheless, we neither claim that CF is "the solution" nor seek to promote the proliferation of "sustainable" LSLAs through CF. Rather, we suggest that moving the needle toward enhanced smallholder agency is a step toward greater smallholder benefit, within the paradigm of commodity agricultural production.

Chapter 7

Conclusions

7.1 Dissertation contributions

7.1.1 Significance

This dissertation attends to the overarching question of how to improve resilience in smallholder agricultural systems. It contributes to this at two levels. First, substantively, I quantitatively compare development approaches in smallholder agricultural systems, finding that there exists a degree of complementarity between different paradigms. Second, methodologically, I develop and demonstrate approaches that enable more robust and equitable policy analysis for socio-environmental systems. As a whole, the research is underpinned by a resilience perspective and integrates approaches from risk analysis, complex systems, and operations research. Given the broad use of both resilience thinking and policy analysis in social sciences, natural sciences, and engineering, the advances within this thesis are relevant to a broad academic community.

7.1.2 Reconciling development perspectives

In smallholder agricultural systems, there exists a wide range of perspectives on how to manage the intertwined domains of livelihood, environment, and economy (Pretty et al., 2018; Bommarco et al., 2013; Gaffney et al., 2019). Disputes surrounding these perspectives have become highly political. In this dissertation, I used agent-based modeling to quantitatively evaluate, compare, and integrate disparate development approaches. An emerging conclusion is that there is a degree of complementarity between alternate development paradigms.

In chapter 5, I compared the effects of cover cropping and microinsurance. Cover cropping is an ecologically-based farm management practice traditionally discussed within agroecology and conservation agriculture, whereas microinsurance is a financial support for sustainable development. I found that, due to the different mechanisms through which these two approaches operate, they offer strongly complementary benefits to smallholder drought resilience. Thus, this research advocates for a greater integration of ecological and financial development approaches. This could have significant impacts by informing future research and programs aimed at agricultural development.

In chapter 6, I examined the potential for contract farming to simultaneously support smallholder livelihoods and increase market-oriented agricultural production. This research was situated in the context of large-scale land acquisitions (LSLAs), which are a strategy for increasing commodity agricultural production but pose a large risk to smallholders through dispossession of land rights. I found that contract farming offers potential as a more smallholder-centric development approach, relative to LSLAs. Thus, this research offers a reconciliatory perspective on market-oriented agricultural development, particularly underscoring the importance of preserving smallholder land rights and autonomy.

A broader consideration within this theme is how to measure resilience and development. A diverse range of objectives can be prioritized within smallholder systems, and these objectives may favor or discount certain groups or development strategies. For example, in Chapter 3 I found that assessments based on "poverty reduction," measured as the reduction of food insecurity over time, suggested that climate forecasts and increased non-farm employment availability could lead to equitable effects, i.e., most strongly benefit the most food insecure households. However, similar assessments based on "shock absorption," measured as the reduction of food insecurity in the wake of a drought, led to inequitable effects, i.e., benefits to the less vulnerable at the expense of the food insecure. Thus, future research and programs for smallholder resilience need to consider a suite of indicators that measure distinct objectives. "Resilience thinking" is a promising holistic lens through which to understand development as a set of distinct yet intertwined objectives that encompass the capacity to persist, adapt, and transform.

7.1.3 Robust and equity-oriented policy analysis

This thesis presents methodological advances to quantitative resilience assessment and model-driven policy analysis. Despite the very real potential for policy recommendations to be maladaptive (to lead to unintended increases in vulnerability), concepts such as equity, temporal tradeoffs, and model uncertainty are rarely considered in model-based analyses. These contributions, described below, facilitate more robust policy assessments and therefore reduce the potential for maladaptation. They are relevant beyond the context of smallholder agriculture and could be applied to models in a diverse range of socio-environmental systems.

First, in Chapter 2, I developed an encompassing framework for the interface of equity and agent-based modeling. This framework can be used to guide future agent-based modeling research, both in agriculture and elsewhere, to integrate equity more thoroughly into its design and application. This has two principal benefits. First, it facilitates the use of ABMs to understand and identify strategies for ameliorating societal inequities. Second, by critically reflecting on potential inequities within the modeling process more generally, it reduces the risk of inadvertently perpetuating societal inequities through agent-based modeling, which has led to highly publicized criticism in other modeling fields (i.e., machine learning).

In Chapter 3, I applied an ABM to examine equity in the effects of development interventions within a smallholder agricultural system. To do so, I stratified the smallholder households by demographic characteristics (e.g., land area, household size) and baseline vulnerability (using a measure of food security). These simple analytical extensions allowed me to assess equity in the distribution of the interventions' effects, thereby providing a more robust assessment. Other ABMs could easily use a similar approach.

In Chapter 4, I presented an approach for acknowledging equifinality in model-based policy analysis. My approach identifies a set of diverse model calibrations, which can be independently applied to a policy analysis, resulting in a more robust policy assessment that incorporates the effects of model equifinality. Equifinality is a common concern in complex systems, yet is not frequently accounted for in model calibration. This approach is therefore applicable to a wide range of application areas.

A final analytical consideration that proved to be important in comparative policy assessment was the time horizon. In Chapter 5, I compared the relative effects of cover cropping and microinsurance over two time dimensions: the time at which a drought occurs and the time over which the drought's effects are assessed. I found that these features exerted a strong influence on the comparison, leading to drastically different recommendations at different points in time. In other application areas, interventions likely work over different time scales and so comparative assessments should carefully consider the effect of time.

7.2 Future research directions

Building from the findings within this dissertation on complementarity, more work is needed to encourage productive conversations between communities with paradigmatically different development perspectives. Although this is not solely a research frontier, agent-based modeling shows promise to act as a boundary object in such efforts (Reilly et al., 2018). By formalizing and contrasting alternative system representations and interventions, agent-based models can be used to facilitate discussions between stakeholders and work toward common understanding. There already exists a large amount of work on participatory modeling in socio-environmental research and beyond (Steger et al., 2021a; Voinov et al., 2016; Loureiro et al., 2020), so leveraging participatory processes with diverse actors to contrast agricultural system futures is a fruitful area for future research.

With respect to robustness and maladaptation, there exist vast literatures on robust decision-making, equifinality, and uncertainty quantification. In a socio-environmental context, my impression is that these ideas and methods have been developed and applied mainly within hydrology, water, and energy systems (Kasprzyk et al., 2013; Reed et al., 2013; Ekblad and Herman, 2021; Vrugt and Beven, 2018), but not sufficiently integrated into agriculture and food systems. There is therefore a great deal to be learnt at these intersections, which can help to develop robust models and identify robust policies for resilient food systems.

Beyond the content of the model analysis itself, there exist many ways in which agent-based modeling can more thoroughly engage with and seek to mitigate societal inequities. This includes embedding equity considerations within standardized agent-based modeling practices and presentation, such as the ODD and TRACE protocols (Augusiak et al., 2014; Grimm et al., 2014, 2020). More consistent emphasis should be placed on integrating stakeholder values and participation within model-based analysis, including using models as boundary objects to stimulate debates and work toward consensus within diverse stakeholder groups (Steger et al., 2021b).

Finally, in keeping with the resilience perspective of this dissertation, more work is needed in the realm of modeling food system "transformation". Much research has focused on system persistence and response to shocks, as well as adaptive actions within current system configurations (Egli et al., 2018). Yet, the future will bring unprecedented threats that will require (or, alternatively, cause) social-ecological systems to fundamentally reconfigure. There exists a growing body of social science perspectives on system transformation (e.g., (Geels, 2002; Moore et al., 2014; Scoones et al., 2020)) that can be learned from and operationalized in current and future modeling frameworks. Some transformative modeling work is beginning to emerge (Zagaria et al., 2021), but there is a need for more empirically informed assessments of multiscale agricultural transformation drivers, inhibitors, and pathways, including how these are conditioned by context.

Appendix A

Supplement to Equity in Agent-based Modeling

A.1 Example bias and positionality statement

This statement was written by the principal researcher and reflects on bias and positionality with respect to the ABM research in Williams et al. (2020a), which examines strategies for increasing the resilience of smallholder farmer livelihoods.

A.1.1 Positionality

What are the racial and cultural backgrounds and identities of the modeler(s)?

The principal researcher identifies as a White male of European descent. He is a Ph.D. student within an engineering department at a research institute in the United States. The other members of the research team identify with a range of identities. All identify as male and work at predominantly White academic institutions in the United States. Two of the research team identify as European American and one identifies as a Marwari Bihari out of place. Although the entire team played a role in shaping the direction of the research, the following reflection is from the perspective of the principal researcher.

How might these identities have influenced how the modeler(s) experience the world and approach research?

The principal researcher acknowledges the privilege he has experienced throughout his life and seeks to use this position of relative power to work to deconstruct global and societal inequities. Conditioned by his academic training, the principal researcher assumes a primarily positivist paradigm but seeks to integrate critical theory to acknowledge the subjectivity in how people perceive reality.

How do these identities, worldviews, and objectives relate to the participants and/or context of the research? (e.g., in what ways are the modelers insiders or outsiders?)

Given his advantaged position, the principal researcher has never personally experienced food insecurity or lived in poverty and so lacks understanding of the lived experiences of Ethiopian smallholder farmers. Thus, he approaches the research context from an inherently outsider's perspective.

What is the social, institutional, and historical nature of inequity in the context of the research?

Imperialism has historically imposed Western values on non-Western countries, such as Ethiopia, frequently through forms of domination and subjugation. More recently, globalization has radically transformed smallholder farming systems. Thus, there is a pervasive antecedent landscape of inequity within the research context. The principal researcher acknowledges this and how the modeling exercise may risk repeating historical colonialist practices, such as information extraction and cultural appropriation.

A.1.2 Framing

*What **narratives** underlie the formulation of the challenge, problem, or research questions?*

The framing of the research contends that the livelihoods of smallholder farmers in Ethiopia (and elsewhere) are vulnerable to the effects of climate. Further, it contends that this vulnerability is problematic——i.e., that it should be reduced.

*What kinds of **solutions** does the problem framing open itself to?*

Through its focus on region-level 'resilience-enhancing strategies,' the research contends that this vulnerability could—and thereby should—be reduced through top-down external intervention. The strategies that are tested (seasonal climate forecasts and increased non-farm employment) both originate from the "West" and were not co-developed with local farmers or decision-makers to inform this modeling work. For these reasons, these solutions may perpetuate recognitional and/or procedural inequities. For example, seasonal climate forecasts are technology-centric and may not agree with local belief systems. Non-farm employment displaces labor from agriculture and catalyzes industrialization, which do not necessarily align with community priorities.

*What **entities** are included/excluded? Who are the actors involved within the framing of the problem and solutions?*

The principal actor within the model is the smallholder farmer. Farmers are modeled as heterogeneous with respect to their land endowment and family size, but we do not represent heterogeneous ethnocultural groups, gender, or other dimensions of identity.

Due to the region-level focus of the resilience-enhancing strategies, the problem's framing places the onus for adaptation on the government or higher institutions. We do not explicitly model these actors, the political processes for implementing the strategies, or the politics of unequal access to the strategies.

*What **outputs** are included and prioritized?*

The principal model output is a measure of food insecurity. Because the most vulnerable households are food insecure, our model outputs are most sensitive to the experiences of these households and we therefore assume a needs- or vulnerability-based perspective on equity.

*What is the **scale** of focus within the problem framing?*

The scale of intervention is the region- or community-level, whereas the scale of modeling is the household-level. Yet, the scale of the narrative (i.e., that smallholder farmers are vulnerable to the effects of climate) frames the problem as one of global concern. This framing does not mention context-specificity in drivers of vulnerability or the appropriateness of the interventions within specific smallholder systems, which are incredibly diverse.

*What **theories** and/or relationships is the conceptual model predicated on? If relevant, are there alternative explanations?*

The model is predicated on the notion that climatic disturbances (i.e., drought) affect smallholder food security, but that these effects are mediated by both household attributes and chosen livelihood strategies. Livelihoods comprise farming, livestock rearing, and non-farm employment. We assume that smallholder farmers have some agency to affect their own outcomes through their livelihood decision-making, but are to some extent constrained by the structure of the system.

A.1.3 Data

How could historical patterns of inequity exist within the data?

The data used for model calibration were drawn from the World Bank's Living Standards Measurement Study (LSMS). In particular, we focused on households' reported experience of food shortages to form our food insecurity measure. This variable specifically aims to measure relative household vulnerability and so, by design, represents historical distributional inequities.

How could marginalized people or groups be misrepresented in or excluded from the data?

The LSMS surveyed one person in each household (the household head). Yet, there are frequently intra-household gender differentials in power and food access, so the survey responses may not include these perspectives. Our research team did not conduct these surveys and so we are not knowledgeable about whether survey enumerators could have biased participants' answers or selected households in a biased way.

How could the process of data collection have perpetuated inequity?

As we did not conduct these surveys, we are not well positioned to answer this question. However, because the LSMS surveys are global in scope, it is unlikely that the World Bank partnered with local communities to collect the data and so critics may view the collection process as extractive.

A.1.4 Process quantification

What subjectivity is involved in defining model variables and/or translating information from data sources into the model format? (e.g., are model variables latent constructs?)

First, food insecurity is a subjective measure and households likely have different perceptions around what it means to "experience a food shortage" (in the language of the LSMS survey). For example, households may compare their experience to others in their networks and therefore the empirical survey measure is likely geographically biased. In particular, it may underrepresent food insecurity in vulnerable regions.

Second, our model-based representation of food consumption is also imperfect; we model production and consumption of a single cereal crop and develop a threshold of food consumption below which a household is considered food insecure. Thus, there may be a mismatch between the theoretical understanding of food insecurity and our operationalization of it.

Could the inclusion or exclusion of model processes misrepresent or lead to bias against certain groups?

There are several potentially important processes not included within the model. First, the model does not include forest-based livelihoods, which are particularly important for land-poor and resource-constrained households. Second, we do not model social networks or sharing of resources (e.g., food) between households. Third, we do not model other social support systems, such as the productive safety net program in Ethiopia. Fourth, we do not model land degradation, which may affect crop yields and thereby the relative prevalence of food insecurity over time. These mechanisms have potentially divergent implications for equity and so it is difficult to speculate about their net effect on food insecurity.

More generally, we did not involve Ethiopian stakeholders during model development. This likely limits the model's acceptability in the modeled context and does not ameliorate power imbalances between the researchers (i.e., scientists in the Global North) and the research subjects (smallholder farmers in Ethiopia).

A.1.5 Model interpretation

How could calibration and validation procedures prioritize models that (dis)advantage certain modeled subgroups?

We used a genetic algorithm to calibrate the model to data from the LSMS. The model's fit was relatively accurate over most patterns, but the model overestimated the proportion of households with no livestock. Given its accurate fit to the food security outcome and the less accurate fit to livestock, the model potentially contains a biased representation of the relationship between these two variables. It may therefore attribute food insecure households with more or less adaptive capacity than they possess in reality.

More generally, the optimization procedure within the genetic algorithm sought to minimize the total discrepancy between the empirical distributions and comparable model-derived distributions. This therefore treated all fitting patterns and all households equally and was not designed to ensure accurate representation of the most vulnerable households.

How could pre-conceived understandings or objectives affect which model structures and outputs are considered acceptable and subsequently communicated?

The model calibration process was highly iterative and required many rounds of model development and experiments until we arrived at an acceptable level of fit. The 'acceptable' level of fit was not defined ahead of the modeling exercise and was ultimately determined by pragmatic constraints (researcher time availability).

With respect to the resilience-enhancing strategies, the research team did not enter the project with pre-conceived opinions around which of the two strategies is preferable, so we believe our assessment was relatively unbiased in this respect. However, we viewed the focus on equity and distributional effects to be a novel scientific contribution and so were motivated to communicate and frame our results around this story. Thus, we focused on the model outputs that yielded the most interesting story about equity and, through this process, discounted some outputs. We do not believe that this process discriminated against particular socio-cultural groups and therefore is unlikely to be inequitable in this way, but it is an issue in science more generally.

Appendix B

Supplement to Resilience and Equity

B.1 Data sources

Table B.1: Data used to initialize the model. "rast" means that the data are represented as a raster. "dist" means that a distribution of values is used.

Object	Parameter	Unit	Value	Description/source
Market	Wood sell	$birr/m^3$	515	CSA Monthly Retail Prices (Amhara, 2015-16 average) [1]
	Livestock sell	birr/head	3536	""
	Milk sell	birr/100L	1554	""
	Fertilizer cost	birr/ha	1000	LSMS 2015 median urea expenditure
	Maize prices	birr/kg	-	FAO GIEWS portal [2]
Environment	Landcover	-	rast	Cropland / non-cropland classification, 30m resolution (Xiong et al., 2017)
Agents	Household size	Adult eq.	dist	LSMS 2015/16 survey
	Landholdings	#plots	dist	""
	Livestock	head	dist	""
	Water collection	hrs/pp/day	dist	""
	Fuelwood collection	hrs/pp/day	dist	""
	Distance to market	km	49.8	""
	Sustenance requirement	quint/pp/mo	0.18	(CAADP, 2013; Worku et al., 2017)
	Initial food stores	months	12	-
	Decision horizon	years	5	-
	Discount rate	%	70.6	(Holden et al., 1998)
	Labor days per month	days	30	-
Climate	Rainfall[3]	mm (/day)	rast	CHIRPS (Funk et al., 2014)
	Temperature[3]	^oC (daily high/low)	rast	GDAS [4]
Yield	Maize yields	kg/ha	dist	Annual agricultural sample survey, Ethiopian Central Statistical Agency

[1] http://www.csa.gov.et/monthly-retail-price

[2] http://www.fao.org/giews/food-prices/tool/public/#/dataset/domestic

[3] All data were bias-corrected and downscaled to a 25km resolution using MicroMet formulations (Liston and Elder, 2006) in the NASA Land Information System (Kumar et al., 2006) and then interpolated over the model grid to give a spatially explicit representation of climate.

[4] http://www.emc.ncep.noaa.gov/gmb/gdas/

B.2 Decision algorithms

```
farming = {no, yes}
livestock = {sell, nothing, buy}
fertilizer = {no, yes}
off-farm % = {0,25,50,75,100}

livelihood options = [farming] X [livestock] X [fertilizer] X [off-farm %]

For each livelihood option:
    Allocate domestic, then farm, then livestock, then off-farm labor

    For each yield forecasting model:
        Estimate monthly income
        Estimate monthly food security
        Calculate utility (using wealth or leisure)

    Calculate expected risk averse utility

Retain option(s) with the lowest perceived food insecurity
If more than one option remains:
    Select livelihood option with maximum utility
```

Figure B.1: Start-of-year decision-making process employed by the agents. Agents iterate through *livelihood* options, which are landuse-livestock combinations. The planting date decision is described in section B.7.5 in SM B.7. The yield forecasting models refer to a discrete number of heuristics agents may use to predict yields (see section B.7.5, SM B.7).

```
for m in 1:12:
    add monthly wage labor income
    add monthly livestock `food' (milk)
    if m == harvest month:
        add crop yields

    subtract monthly expenditures
    subtract monthly food consumption

    if food stores > 12x monthly consumption:
        sell excess food

    # COPING MEASURES
    if cash < 0:
        sell food stores to neutralize cash
    if cash < 0:
        sell livestock to neutralize cash
    if food stores < 0:
        buy food from market
    if (food stores < threshold) AND (previously food insecure):
        sell livestock to neutralize food stores
```

Figure B.2: Monthly food consumption and coping measures process.

B.3 Additional figures

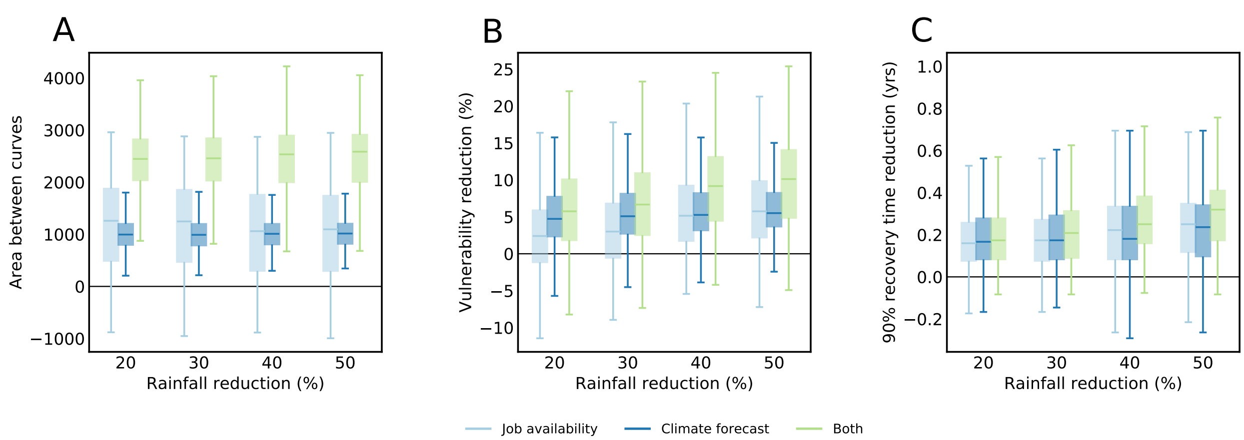

Figure B.3: The benefits of each intervention on different dimensions of resilience: (A) overall resilience (area between food security curves), (B) vulnerability (the maximum damage), and (C) 90% recovery time. The boxplots indicate the distribution of benefits over all simulations.

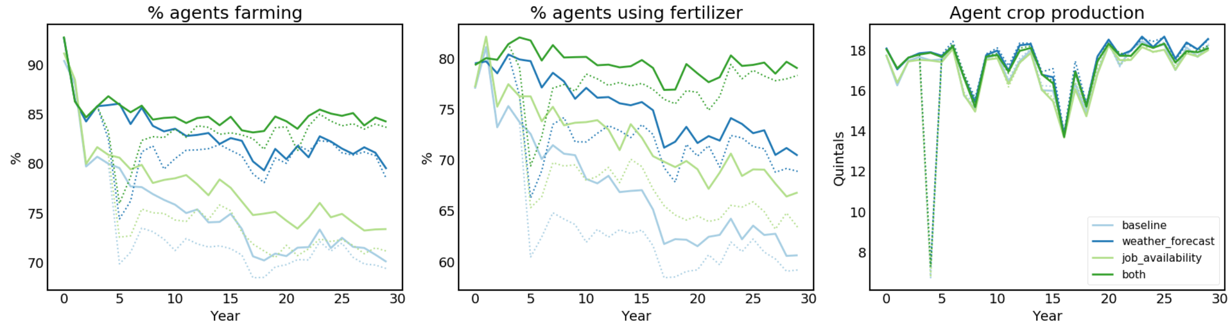

Figure B.4: Impact of interventions on farming practices. Dashed lines represent the shock simulations and solid lines the no-shock simulations.

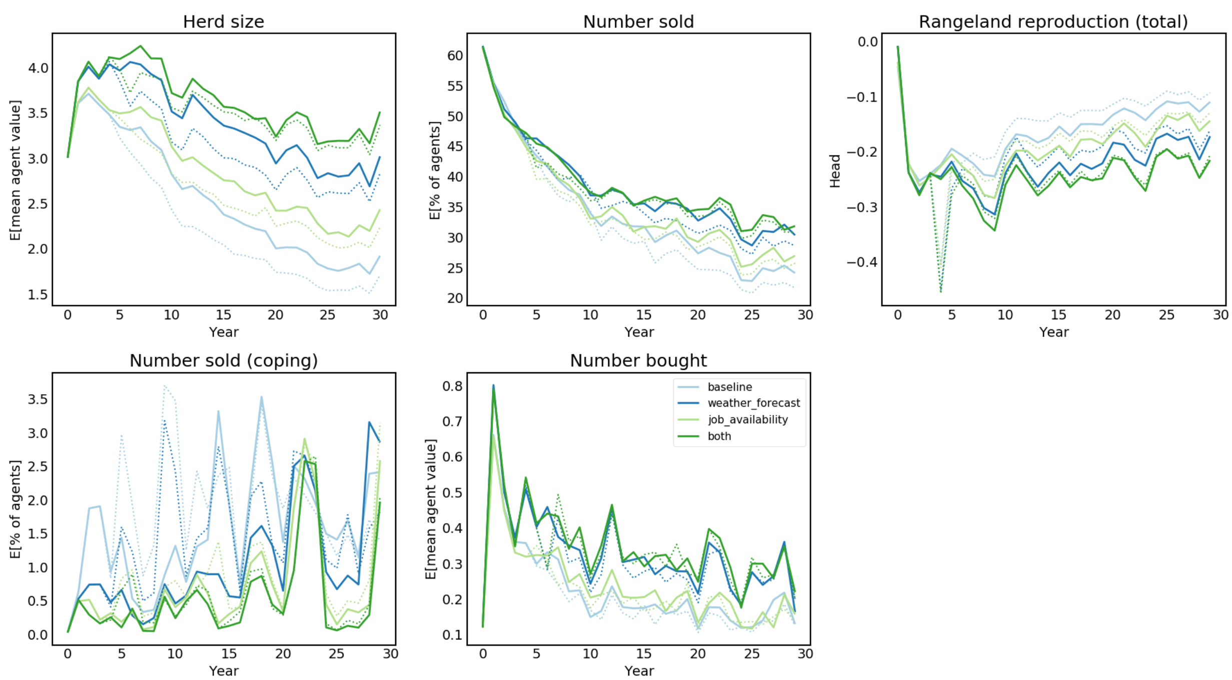

Figure B.5: Impact of interventions on livestock. Dashed lines represent the shock simulations and solid lines the no-shock simulations.

B.4 Cross validation of calibration distributions

Overall, the ABM is qualitatively in accordance with the empirical distributions (Figure 3.5). However, it does not accurately recreate the shape of the herdsize distribution: the ABM overestimates the number of households with no livestock. Given the structure of our model, we found that to generate levels of food insecurity comparable to the empirical values (i.e., approximately 20% of households food insecure) required a substantial amount of agents to have no livestock. In the empirical data, there are more food insecure households than households with no livestock, yet our model does not conform with this. By making households hesitant to sell their livestock (i.e., only selling livestock after experiencing repeated months of food insecurity) we hoped to improve this fit, yet were not able to do so. Further refinement of the coping mechanism heuristic, as well as potentially the rangeland model, may be required to improve this fit.

One of the recommendations for pattern-oriented modeling is that the patterns against which model outputs are compared are independent (Grimm et al., 2005; Latombe et al., 2011); that is, each fitting distribution should offer some additional information. However, given the flexibility of the ABM, there is a risk that we have overfit the model to these distributions. There is therefore a trade-off between selecting patterns that are sufficiently independent and designing a model structure that can robustly reproduce patterns on which it has not been calibrated. To test this, we conducted a cross validation by running the genetic algorithm parameterization procedure five times, each time excluding one of the fitting distributions from the objective function. Our results (Figure B.6) show that there is a reduction in fit when each distribution is excluded. The largest effect - almost a 4000% relative decrease in fit - is observed in the subsistence fraction measure, suggesting that the fitting procedure relies very heavily on this distribution. The smallest absolute effects are seen for the two labor distributions (on- and off-farm labor), suggesting some redundancy in the information contained in these data and/or robustness in the structure of the labor components of the ABM. For food insecurity - the main outcome of interest for the resilience assessment - the loss metric increases by 70% when this distribution is excluded.

Together, these results suggest that in the spectrum of this trade-off between independent, useful patterns and model structure robustness, our calibrated model more closely aligns with the former; that is, each distribution offers considerable value to the fitting process. We therefore rely on our justifications for the structural assumptions within the ABM in our resilience analysis.

B.5 Assessing synergistic effects

To investigate the potential for synergistic effects, we calculated the additional benefit realized by having both policies in conjunction compared to the sum of the two in isolation (i.e. both -

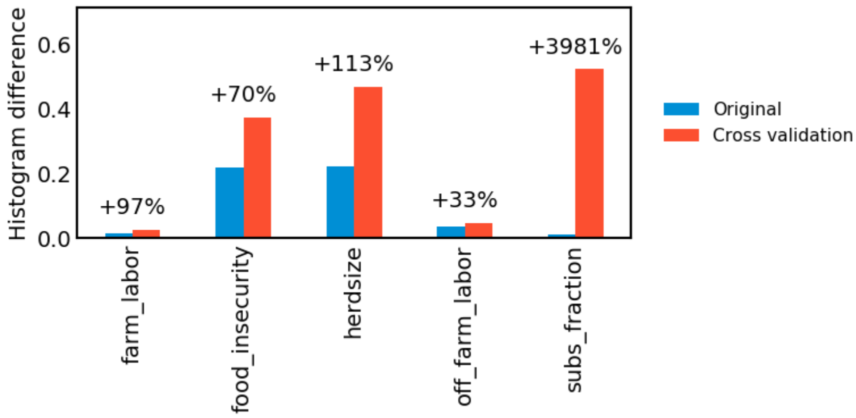

Figure B.6: Cross validation of fitting metrics. Comparison of the ABM-generated and empirical distributions for each fitting characteristic when that characteristic was *not* used in the parameterization process. The "histogram difference" represents the loss metric used in the genetic algorithm.

(job_availability + climate_forecast)). The majority of the simulations exhibited a synergy with respect to our measure of overall resilience (Figure B.7A), with, for example, 65% of simulations under a 50% magnitude shock showing synergistic effects. This is because the two interventions work through different pathways, with the climate forecasts helping to inform agricultural management decisions and wage labor availability working as an additional coping strategy throughout the year. Though not shown here, the combined interventions only resulted in a tradeoff (i.e., net benefits lower than an individual intervention) in a maximum of 4% of the simulations (for the 50% magnitude shock).

The combined interventions are less synergistic with respect to vulnerability reductions; for the 50% shock, only 30% of the simulations were synergistic with respect to this measure (Figure B.7B). Given that the measure of vulnerability is equivalent to our overall resilience measure assessed over a single year time horizon (i.e., maximum additional population food insecure in a single year versus total additional population food insecure over all time), this demonstrates that the synergistic effects of the interventions take time to accrue. In general, the potential for synergies is an important result and agrees with other studies recommending that agricultural interventions can be most beneficial in portfolios (Berger et al., 2017; Wossen et al., 2017). We note that we do not display synergies in the time to 90% recovery measure because this calculation relies on the vulnerability value, which is different under each intervention.

B.6 Sensitivity analysis

To explore the robustness of our results to the calibrated values of the uncertain parameters we conducted a one-way sensitivity analysis in which we set six of the uncertain parameters to 60%,

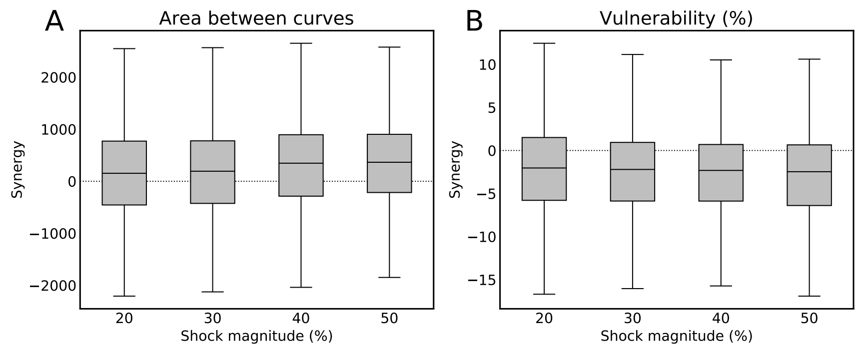

Figure B.7: Distribution of the synergistic benefits realized through implementing both policies in conjunction for: (A) overall resilience, measured as the area between the food security-time curves (Figure 3.4); and (B) vulnerability, measured as the maximum impact of the shock. In this graph, resilience Equation 3.2 is used.

80%, 120% and 140% of their fitted values and reran our experiments. Here, due to computational time requirements, we selected only these six parameters and ran only five replications of the model.

We explored the impact of these changes on the relative effect of the two policies on overall resilience (Figure B.8) and on the distribution of food insecurity throughout the population (Figure B.9). We note that the cases in which it appears that no data are plotted in Figure B.9 represent situations where the given parameter change had no effect on the distribution of food insecurity, thus providing evidence for the robustness of the distributional results.

A: Planting fraction: This represents the fraction of each agent's land that can be farmed. The reason that this is an uncertain parameter is due to a mismatch between the LSMS data used to parameterize agent-level land holdings and the Ethiopian Central Statistical Agency (CSA) field-level data used to parameterize the crop yield model; assuming that all land could be planted and that crop yields were in line with those reported by the CSA resulted in production that was much too great to achieve levels of subsistence in accordance with that calculated from the LSMS data. As a result, we assumed that agents are only able to plant some fraction of their land. This could also represent yield losses or fallowing practices.

The results show that when little land is available for agriculture (or yield losses are very high), the agents with large land holdings become more food secure (Figure B.9). This is because, under the calibration presented in the main body of the paper, agents with large land endowments are in some cases not able to farm their land due to large labor requirements. A reduced planting fraction therefore lowers this labor requirement and allows these agents to engage in farming, thus improving their food security. At the same time, agents with lower land holdings become more food insecure due to lower planting fractions limiting their production capacity. Together, this contributes to an increased relative benefit in the climate forecast intervention (Figure B.8A); agricultural productivity is more limited, so increased information about climate conditions is more important in the wake of a drought.

B: Grass regeneration rate: This parameter represents the rate at which the rangeland is able to regenerate. The results are insensitive to changes in this parameter.

C: Farm labor factor: As the labor requirements for farming are increased, the relative benefit of the climate forecast policy increases (Figure B.8C). An explanation of this effect is that as more labor is required for farming, agents have less remaining time to engage in other activities and therefore less of an ability to diversify their incomes. As such, they rely heavily on agriculture and thus are strongly impacted by a drought. The availability of climate forecasts, even though they do not provide perfect information, enables agents to divert effort away from farming in the year of the drought, which has larger implications when farming labor requirements are high.

Changes to this parameter have substantial distributional effects (Figure B.9), with reductions in the labor requirements for farming greatly benefiting those with large amounts of land. The reason for this is the same as the effect described in the planting fraction above; lower labor requirements enables more of these households to engage in farming.

D: Risk aversion: The model is insensitive to changes in the risk attitude of the agents. This is interesting, and suggests that risk aversion does not have strong implications for decision-making in the model.

E: Job availability: We have already discussed the impacts of an increased availability through-

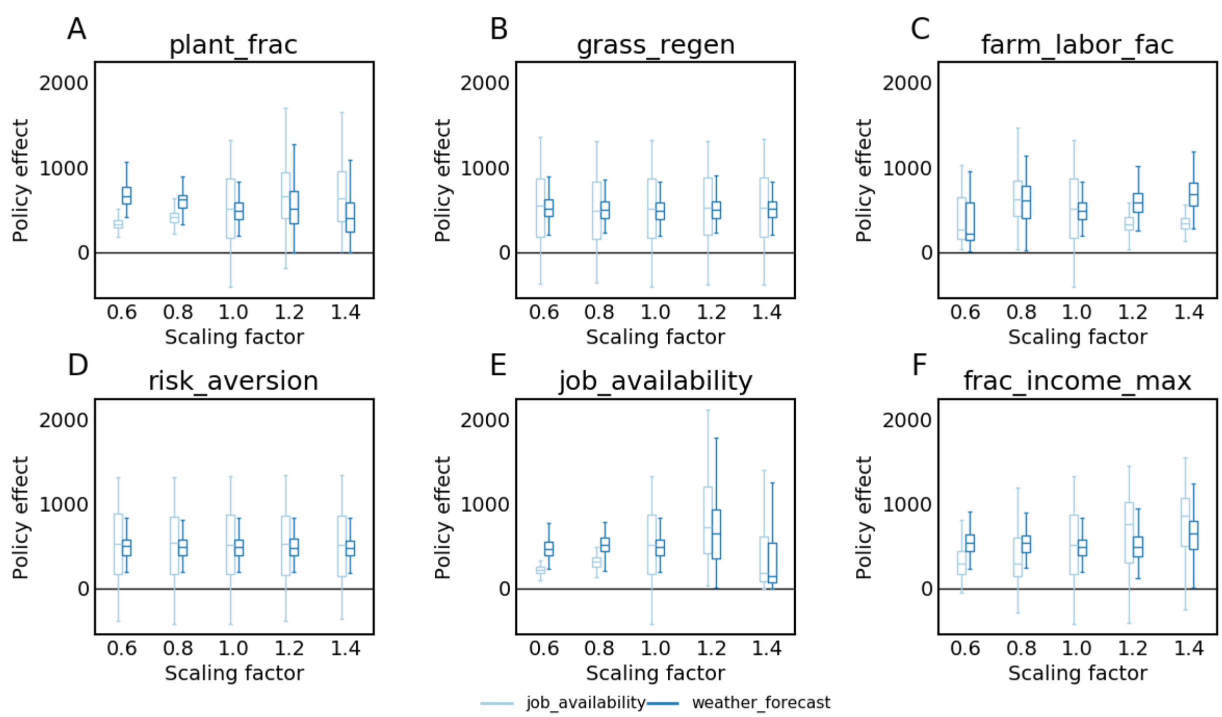

Figure B.8: Sensitivity analysis: Effect of varying selected parameters on overall resilience under each intervention.

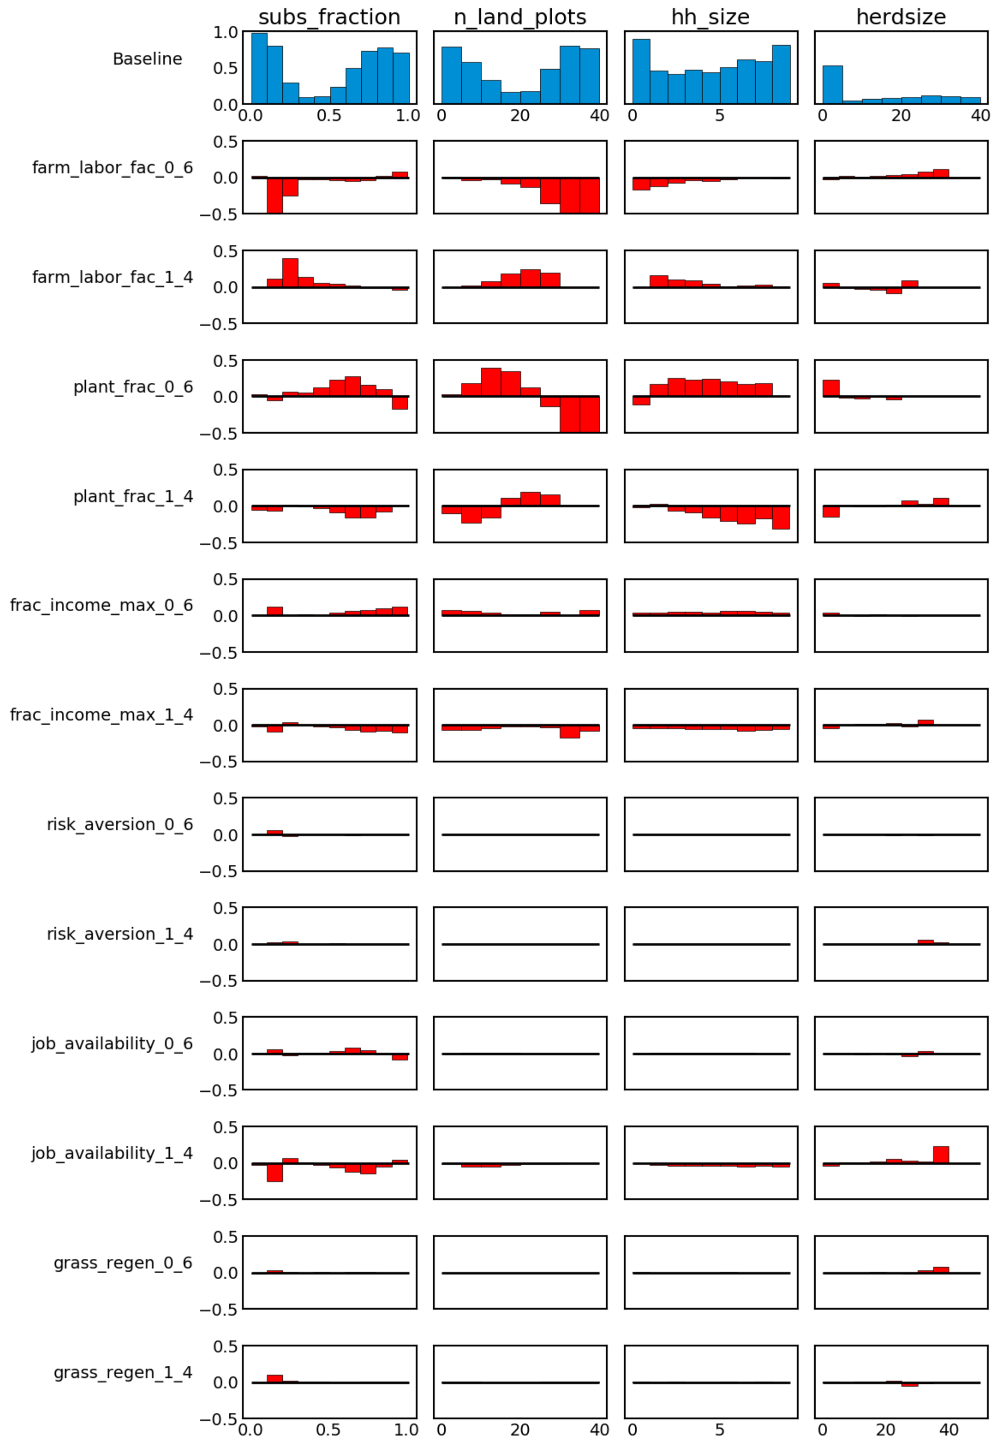

Figure B.9: Impact of varying selected parameters on the distribution of food insecurity. The units of the vertical axis represent the fraction of each population group experiencing food insecurity at some point in the simulation (median over all replications) (Baseline, blue) and the effect of parameter changes on this (all others, red). Variable names correspond to those in Table B.4 and the 0_6 and 1_4 respectively represent the variables being set to 60% and 140% of their original values.

out the rest of the paper. Here we test how these conclusions may change as the assumed baseline job availability is varied. Interestingly, the job availability intervention has a lower benefit when the baseline job availability is decreased (Figure B.9E). This is likely simply a result of the way in which the job availability intervention was implemented (i.e., as a percentage increase from the baseline value); when baseline job availability is low, a 20% increase amounts to a lower absolute increase in jobs, and hence the absolute benefits are lower.

F: Fraction income maximizers: This parameter represents the fraction of the agent population that have a preference for income maximization. Here, the job availability intervention has increasingly strong benefits as the proportion of income-maximizing population increases (Figure B.9F). Given that the income-maximizing agents are assumed to be independent of the limited availability job market, increasing the proportion of these agents therefore increases the number of jobs that are available to the "leisure-maximizing" agents. When this change in the makeup of the agent population is combined with the effects of an externally-imposed increase in job availability, the combined effect is that the leisure-maximizing agents are able to more effectively leverage increased job availability in the wake of a shock.

B.7 Additional ABM details

B.7.1 Crop yields

B.7.1.1 Climate yield factors

The climate yield factor (CYF) is defined using the ratio of the simulated actual evapotranspiration (ET) (ET_A) to the potential ET (ET_C):

$$CYF = 1 - K_y\Big(1 - \frac{ET_A}{ET_C}\Big) \tag{B.1}$$

where K_y is a yield response factor, which is crop-specific and varies over the growth stages of the crop (vegetative, flowering, yield formation, and ripening) (FAO, 1984, 1998). The meteorological inputs required to calculate the CYF include precipitation, temperature, solar radiation, cloud cover, humidity, and wind speeds. The calculation of actual ET is made on a daily basis throughout the growing season, accounting for crop water demand, soil moisture, and root growth. Critical values of the CYF are taken for each crop growth stage to give an overall seasonal CYF. We note that similar representations of the effect of water on yields (i.e., some ratio of potential and actual evapotranspiration) are used in other process-based crop yield models such as FAO's AquaCrop (http://www.fao.org/aquacrop/), CENTURY (USDA, 1993), APSIM (https://www.apsim.info/), and EPIC (https://epicapex.tamu.edu/epic/).

B.7.1.2 Calculating agent-level yields

The following procedure was implemented to calculate agent-level yields:

1. Derive a region-wide log-normal distribution from which to sample Y^p values.
2. Randomly assign each agent an inherent "position" in this distribution (i.e., a quantile).
3. At each time step (t) in the simulation, add some noise to this quantile (i.e., $q_{t=t,a=a} = q_a + \epsilon_{t,a}$, where a represents an agent and $\epsilon_{t,a} \in (-\epsilon, +\epsilon)$ is some noise) .
4. Calculate the agent's plot-level yield as $Y^{act.}_{t,a} = CYF_{t,a} * Y^p_{t,a}$ where $Y^p_{t,a}$ represents the perturbed quantile $q_{t,a}$ from step 3.

First, to calibrate the magnitude of the potential yield (Y^p), we used observational yield data from Ethiopia's Central Statistical Agency's (CSA's) Annual Agricultural Sample Survey (AgSS). We created Y^p distributions under fertilized and non-fertilized conditions. In doing so, we proxy the effects of nutrient limitations on crop yields, constraining these effects such that they recreate historically-observed discrepancies between fertilized and non-fertilized crop yields. For both

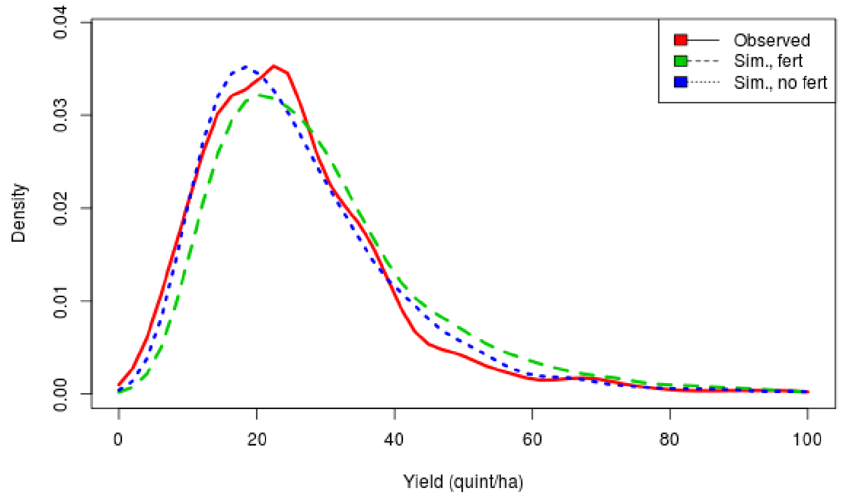

Figure B.10: Comparison of empirically-observed maize yield distribution (overall) and simulated yields with and without fertilizer.

fertilization options we fitted a log-normal distribution using plot-level maize yield data from the CSA's AgSS from 2005-2010 and 2012. The data were cropped to a 50km buffer around the model region. First, empirical Y^p values were estimated as $Y^p_{emp} = Y^{obs.}/CYF^{obs.}$. The mean of this distribution was then taken for fertilized and unfertilized fields and used as the mean parameter in the log-normal distribution. The standard deviation was chosen to fit the spread of data from the simulated log-normal distribution (adjusted by the observed CYF values) to the spread of the empirically-observed yields. Figure B.10 compares the empirically-observed distribution of maize yields in the model region with simulated yields using the corresponding historic CYF values. Figure B.11 shows the simulated yield distributions for fertilized fields over a range of potential CYF values.

Our approach for calculating crop yields preserves the population-level distribution of yield. However, we do not evaluate it based on its ability to predict field-level yields. We admit that this is an imperfect method for calculating crop yields, and other process-based methods that are validated using observational data from the modeled region may give more defensible yield estimates or be used to incorporate the effects of nutrient limitations in a process-based sense. However, given our focus on drought and lack of site-specific observational yield data, we judge our method to be sufficient for the purposes of this paper. Additionally, it is more computationally feasible within an ABM.

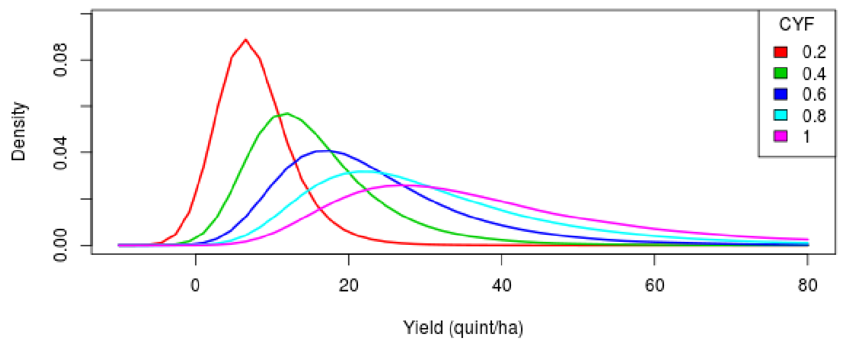

Figure B.11: Simulated crop yield distributions for a range of CYF values. Note: density smoothing was applied to the simulated distributions - there is no probability of negative yields.

B.7.2 Crop prices

Crop prices can vary considerably over time. In countries like Ethiopia, cereal prices are driven by many factors, ranging from local supply-demand dynamics to global market forces (Baffes et al., 2017; Brown et al., 2017). To represent temporal price dynamics in the ABM and proxy the effects of region-level market dynamics, we created a regression model that relates a variety of climate covariates (Table B.2) to monthly crop prices for the region of Amhara[1]. The dataset spanned from 2005 to 2015. Although the LSMS contains data about crop prices, we did not use this for two main reasons: (1) given there are only three LSMS waves, it would be difficult to develop temporal relationships for a single region; and (2) reported prices in the LSMS are not consistent between households, thus adding an additional layer of heterogeneity that may obscure the climate-induced drought effects.

Figure B.12 shows the application of a multiple linear regression model to predict these historic prices using the variables in Table B.2. These figures show the results for *in-sample* predictions for maize in Amhara, so are over-estimating the effectiveness of the model at predicting values that it has not seen, but the results show that the model does well at predicting maize prices using solely climatic information.

Figure B.13 shows the *out-of-sample* predictive accuracy of the linear model, determined using a 50-fold random holdout analysis. The model struggles to predict the high/low prices, but captures the general trends and deviations (though notably underpredicting the price spike in 2008, due to the global food crisis). Inclusion of a greater number of covariates generally improves the predictive accuracy (Figure B.13b), so a model using all 10 covariates was used in the ABM. Finally, Figure B.13c shows the predictive accuracy of a model that contains solely a *lag* term,

[1]Sourced from the FAO GIEWS portal: http://www.fao.org/giews/food-prices/tool/public/#/dataset/domestic

Table B.2: Covariates used to predict crop prices. All covariates were developed at monthly intervals for the region of Amhara.

Column	Description
$Price$	Price (birr / 100kg)
$Month$	Month
$EDDs$	Regionalized average extreme degree days (EDDs) for this month
$GDDs$	Regionalized average growing degree days (GDDs) for this month
avg_sum_pre	Regionalized average total precipitation for this month
$rain\_lag$	Regionalized average total precipitation from the previous year
EDD_lag	Regionalized average total EDDs from the previous year
GDD_lag	Regionalized average total GDDs from the previous year
$rain\_YTD$	Total year-to-date sum of rainfall
EDD_YTD	Total year-to-date sum of EDDs
GDD_YTD	Total year-to-date sum of GDDs

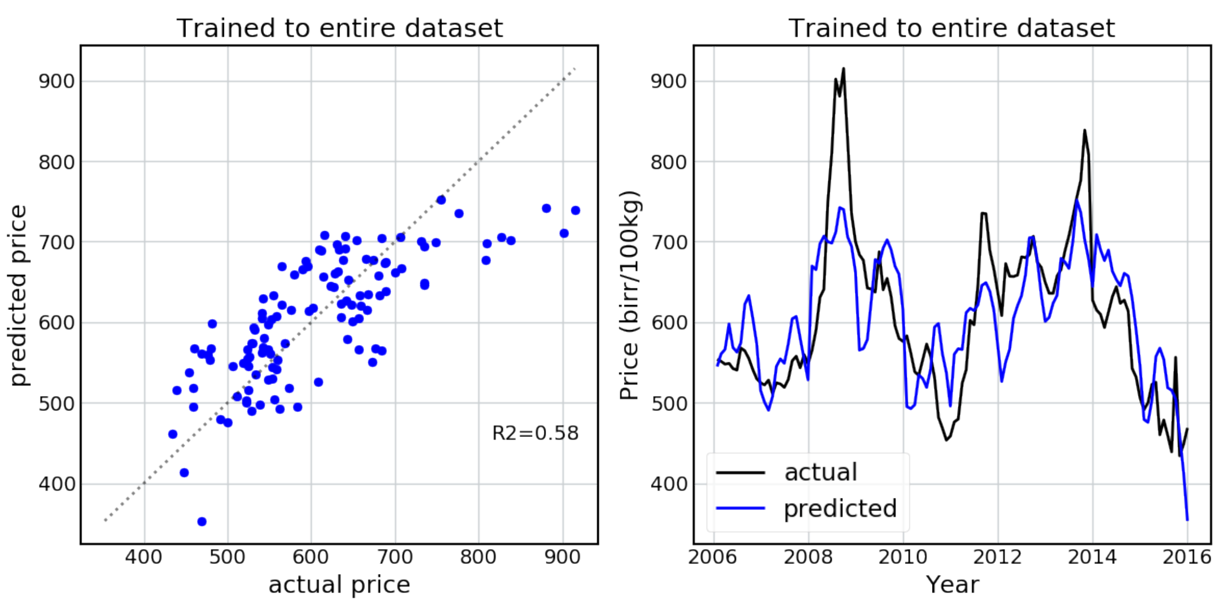

Figure B.12: In-sample predictive accuracy of the price regression model.

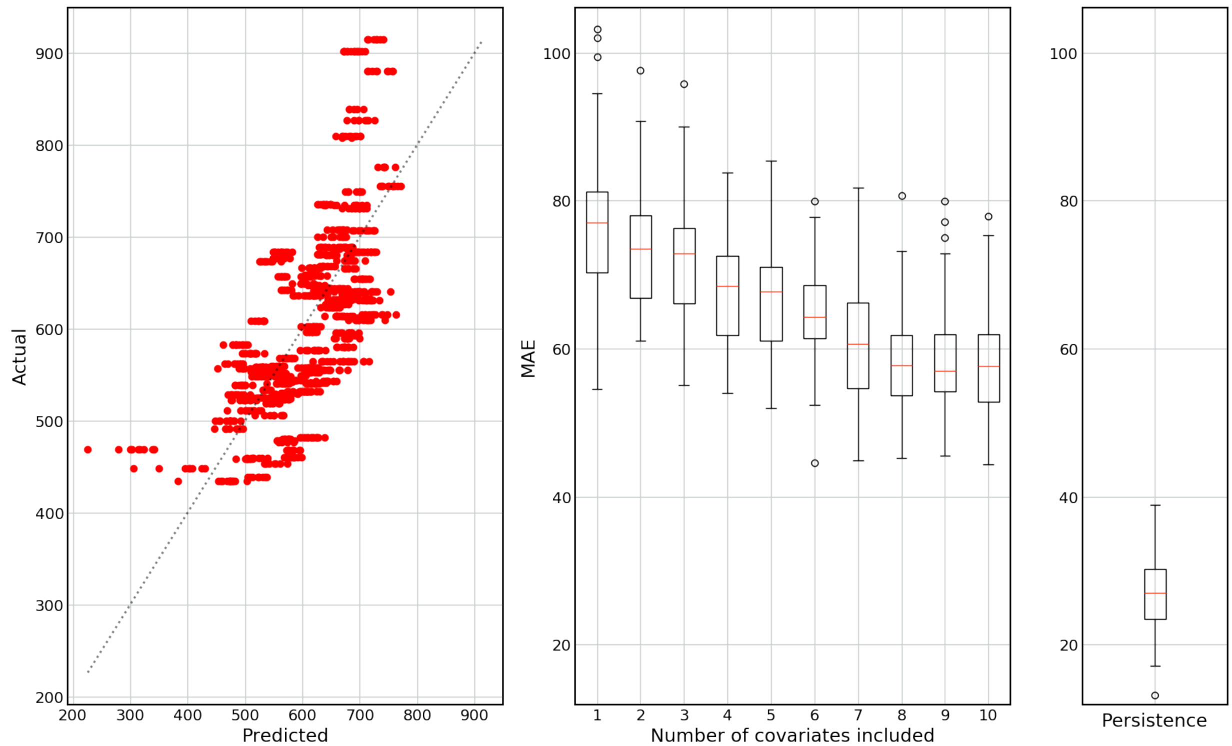

Figure B.13: Out-of-sample predictive accuracy of the price regression model, assessed using a 50-fold repeated random holdout. (a) predicted vs. actual values. (b) mean absolute predictive error (MAE) as a function of the number of covariates included. (c) the out-of-sample MAE of a persistence model, which uses the previous price to predict the subsequent price.

i.e., a model that simply predicts that next month's price will be the same as this month's. This can be used as a benchmark against which to compare the results of the climate-based regression models. We can never achieve as accurate predictions as a lag-only model, and our median MAE is approximately twice as high as the persistence model. However, for the purposes of the ABM, we will not know what last month's price was, and therefore desire a method that can predict price with reasonable accuracy independently of antecedent price levels. Hence, we conclude that the model developed is sufficient for its purpose.

Finally, we note that due to our representation of droughts solely through reductions in rainfall, the effects of temperature on crop prices during the drought are not incorporated.

B.7.3 Livestock dynamics

The state of the aggregated livestock herd density (head/ha) is updated at the end of each year, representing the net effect of rainfall, grazing pressure, and buying/selling on regional livestock reproduction/mortality. Our rangeland model, including the values for some parameters (Table B.3), was inspired by Janssen et al. (2000, 2004). However, we use the genetic algorithm calibration

Table B.3: Coefficients used in the livestock-rangeland dynamics model. '**' values were determined by the genetic algorithm (section B.7.7).

Symbol	Description	Value	Source
α_1	Maximum livestock reproduction rate	**	-
α_2	Maximum livestock mortality rate	0.1	(Janssen et al., 2000)
g_0	Initial grass value	**	-
β	Grass regeneration rate	**	-
g^{max}	Maximum grass biomass	**	-
a	Rainfall intercept value	**	-
b	Rainfall slope value	**	-
cf	Livestock consumption rate	400	(Janssen et al., 2000)

process to set values for several of these due to lack of site-specific environmental data.

Not all livestock are grazed on the rangeland; we assume that livestock are preferentially grazed on crop residues, and that the portion of livestock fed by residues does not face any climate- or sustenance-related mortality and thus reproduce at the maximum rate. The amount of crop residue available for livestock consumption at the agent level is two times the harvestable crop yield in the same year, with a 10% loss factor applied (Bogale et al., 2008; Assefa et al., 2013). The portion of livestock that are not grazed on residues are grazed on the communal rangeland. The rangeland model operates at the regional scale and is characterized by a livestock density (x_t) and a grass biomass (g_t).

In updating the livestock density, we assume that reproduction is constrained by the scarcity of grass and that mortality is higher under low rainfall conditions (i.e., due to lower water availability). We included this rainfall-dependent mortality effect so that drought affects livestock in the year of the drought and not only indirectly through grass biomass (Angassa and Oba, 2007). The updating of the livestock density (x_t) is completed as follows:

$$x_{t+1} - x_t = \alpha_1 * x_t\Big(1 - \frac{cf * x_t}{g_t}\Big) - \alpha_2 * x_t * (2 - rf_t) \tag{B.2}$$

where α_1 is a constant representing the unconstrained reproduction rate, α_2 is a constant representing the unconstrained death rate, cf is a constant consumption factor, and rf_t is the rainfall modifier effect. In using this representation, we assume that grazing pressure and rainfall affect reproduction and mortality respectively, and that these effects are independent.

In line with (Janssen et al., 2000), climate is assumed to affect livestock herds solely through annual precipitation amounts. The rainfall modifier effect is calculated as:

$$r_t = a + b * R_t \tag{B.3}$$

where a and b are constants and R_t is the rainfall in year t. Note that it is possible for r_t to be either negative or positive.

The growth or decline of grass biomass is influenced by both rainfall and livestock grazing. To update the grass biomass, we use:

$$g_{t+1} - g_t = \beta * r_t * g_t * \left(1 - \frac{g_t}{g^{max}}\right) - cf * x_t \tag{B.4}$$

where β is the unconstrained grass regeneration rate and g^{max} is the maximum possible grass biomass.

B.7.4 Beliefs and learning

Agents have beliefs about the probability they will receive a job, crop yields, crop price, livestock growth, the effectiveness of fertilizer, and the timing of rainfall onset. In most cases, beliefs are modeled using probability distributions, which are updated using Bayesian updating - a commonly-used approach for representing beliefs for farmer decision-making (Ng et al., 2011).

Priors: The extent to which new information influences a belief is given by the strength of the agent's prior. Priors indirectly represent the "learning rate" of the agents; agents with stronger priors will have slower learning rates, as new information has a smaller effect on their overall beliefs. Agents are initialized with heterogeneous prior strengths, which are generated from a uniform distribution ($\sim U(1, 10)$). The bounds on this uniform distribution were chosen so that there is considerable heterogeneity over the agent population, and so that agents' beliefs significantly develop over the course of the simulation period. It is assumed that each agent's initial prior is the same for all beliefs, but is scaled by a factor in some cases. A 13-year burn-in period (2003-15) is run in each simulation, which reduces the sensitivity of the model to the belief initialization.

Job probability: This represents the agent's belief about the probability that they will receive a non-farm job on a given day. Since this is a belief about a probability, it is contained on the interval $[0, 1]$, and is modeled using a beta distribution:

$$\pi(\theta) \sim Beta(\alpha, \beta) \tag{B.5}$$

where $\pi(\theta)$ is the density of the agent's believed job probability and α and β are hyperparameters. These are updated using the beta-binomial conjugate prior, where, given n successes out of m trials, the parameters are updated as:

$$\alpha' = \alpha + n \tag{B.6}$$

$$\beta' = \beta + m - n \tag{B.7}$$

The agents' initial mean beliefs (i.e., expected job probability, $E[\pi(\theta)]_{t=0}$) are generated from $U(0.1, 0.9)$, and the initial values for α and β are chosen such that this expectation is honored and such that $\alpha + \beta$ (a representation of prior strength in a beta distribution) is equal to their prior strength.

Fertilizer effectiveness: From the agents' perspective, fertilizer is assumed to have some scaling effect on baseline yields (i.e., $Y_{fert} = f * Y_{no\_fert}$, where f is the scaling factor). This is a simplistic representation, but creating a belief that is conditional/dependent on the yield belief or perceived soil conditions would require additional parameters and complexity. However, we note that a constant scaling effect will result in larger *absolute* yield differences in high-yield situations; for example, if the baseline value is 1 t/ha, a factor of 1.3 amounts to a 0.3 t/ha difference, whereas with a baseline value of 2 t/ha the same factor amounts to a 0.6 t/ha difference. Given that fertilizer generally provides larger benefits in non-drought years, our model qualitatively captures this property.

The scaling factor (f) is represented using a normal model with known variance (assumed to be 0.1). After each growing season, agents observe their and their neighbors' yields with and without fertilizer, and calculate an observed fertilizer effectiveness (x) representing the observed ratio of fertilizer:no-fertilizer yields. Their prior belief is given by:

$$f \sim N(\mu_0, \sigma_0^2) \tag{B.8}$$

and is updated as:

$$f|x \sim N\left(\frac{\sigma_0^2}{\sigma^2 + \sigma_0^2}x + \frac{\sigma^2}{\sigma^2 + \sigma_0^2}\mu_0, \left(\frac{1}{\sigma_0^2} + \frac{1}{\sigma^2}\right)^{-1}\right) \tag{B.9}$$

where σ^2 is the assumed known variance (0.1).

Agents are initialized with a $\mu_0 \sim U(1, 2)$, and $\sigma_0^2 = 0.1 * prior\_strength_{t=0}$.

Crop yields: Yields take on continuous, positive values. Agents evaluate the anticipated yield of a potential farming option as a perceived *baseline* yield, multiplied by a perceived *fertilizer effectiveness* value (if applicable). What follows is a description of agents' *baseline* yield beliefs.

Drawing from (Magliocca et al., 2013), agents employ a number of backward-looking expectation models. Each of these models makes a prediction about an upcoming observation (of yields), given a history of observations. All agents employ every one of these models, but have differing levels of "trust" in each of them. Trust is influenced by the historical accuracy of the models at predicting their yields, and determines the weight that the agent places on each model's prediction when making their decisions.

The motivation for representing these weights is so that agents consider multiple potential realizations of future yields when making their decisions. Given that the purpose of this model is to assess climate resilience, explicit representation of agent-level uncertainty about future climate is important. The range of models implemented here represent potential heuristics that farmers could use when predicting what their (climate-influenced) yields will be in the upcoming year. Future extensions of the ABM could include further or refined forecasting models.

The models employed are as shown below. Y refers to yield (quintals per hectare).

1. *Last period model:* $Y(t+1) = Y(t)_{obs}$
2. *Mean model:* predict that $Y(t+1)$ will be the same as the mean of the last n periods. For simplicity, let $n = 2$.
3. *Trend model:* predict that $Y(t+1)$ will be an extrapolation of the previous two years' values, i.e., $Y(t+1)_{pred} = Y(t)_{obs} + (Y(t)_{obs} - Y(t-1)_{obs})$.
4. *Neighbor model:* predict that $Y(t+1)$ will be the same as the average of the agent's immediate neighbors' previous beliefs.

Agents' trust in each of these models is represented using weights (w), which are categorical probabilities. These weights are interpreted as the probability that a given forecasting model is the most accurate. Weights are represented using a Dirichlet-categorical distribution:

$$w_1, ..., w_n \sim Dir(\alpha_1, ..., \alpha_n) \tag{B.10}$$

$$y \sim Cat(w_1, ..., w_n) \tag{B.11}$$

where the subscripts $1, ..., n$ refer to each of the forecasting models, w_i are weights/probabilities, α_i are hyperparameters, and y is an observation. These hyperparameters are updated as:

$$\alpha_i' = \alpha_i + \mathbb{1}(y_t = i), \forall i \tag{B.12}$$

where $\mathbb{1}$ is an indicator function, and y_t is the model that *most accurately* predicted the observed value at time t. When evaluating a believed probability, the expected value is used, i.e.

$$P_i^{belief} = \frac{\alpha_i}{\sum_j \alpha_j} \tag{B.13}$$

All agents begin the simulation with equal trust in each of the forecasting methods (i.e., $\alpha_{i,t=0}$ are equal $\forall i$ for each agent). An agent's values for $\alpha_{i,t=0}$ are equal to their initial prior strength (i.e., $\in (1, 10)$). Additionally, when initializing the model, it is not possible to calculate previous yields for the forecasting models, so an initial value is drawn for each agent $\sim U(15, 25)$ and assigned to all prediction models. The 13-year burn-in period reduces the sensitivity of the resilience analysis to this initialization.

When including the climate forecast in the simulation, this acts as a fifth forecasting model and behaves exactly the same as the other models. The predicted yield is evaluated using the actual CYF for the year. However, their information is not perfect, as they do not know the random perturbation (ϵ_t from step 3 in section B.7.1).

Prices: Prices also take on continuous, positive values and are modeled in the same way as yields, using an independent set of weights. However, explicit incorporation of uncertainty in price beliefs would increase the computational burden of the decision module without justifiably enriching the model, so agents only consider one value when making their decisions. This is the weighted average prediction from each model, i.e.

$$Price(t)_{pred} = \sum_{i=i}^{n} P_i^{belief} * Price(t)_{pred,i} \tag{B.14}$$

where i indexes the n forecasting models.

Agents are initialized with prior beliefs (for all models) equal to a simulated price using climate data from the year preceding the simulation period. Initial model weights are set equal to agents' initial prior strength.

Although prices are calculated on a monthly basis (section B.7.2), agents' beliefs are equivalent to the *average* annual price.

Livestock growth: Agents have a belief about the growth of their livestock herds. This growth represents the net effect of reproduction and mortality, and is expressed as a fraction, where, for example, -1 represents 100% of the herd dying, 0 represents no net change, and 0.5 represents a 50% increase in herd size. These beliefs are also modeled using the forecasting models described above, and agents have an independent set of weights for their livestock growth beliefs. Again, to limit the computational burden, the expected value over all forecasting models is used in the decision-making process.

Agents are initialized with prior beliefs (for all models) $\sim U(-0.05, 0.05)$ and model weights equal to their initial prior strength.

Rainfall onset: The onset of rainfall is modeled discretely as either early, on-time, or late. Agents have a belief about the probability of each of these three events. These are modeled using a Dirichlet-categorical distribution as described above. The expected value is used when evaluating the probabilities (Equation B.13).

B.7.5 Decision-making

We employ a version of bounded rationality to represent the agents' decision making.

Agent objectives: A common measure of prosperity in Ethiopia is the accumulation of livestock. Livestock can act as a "walking bank account" (Bellemare and Barrett, 2006), providing income in times of need and growing through reproduction. As such, the accumulation of livestock is analogous to the accumulation of capital, or wealth, and can be said to be a primary goal of many Ethiopian smallholders. However, a variety of factors often may prohibit the ability of households to achieve this goal, and trade-offs must be made between wealth accumulation and shorter-term priorities such as stabilization of welfare by reducing risk (Dercon, 2004), for example ensuring the provision of food to sustain the family. This tendency has been empirically observed, in which wealthy households are primarily motivated by profit maximization and poor households by risk minimization and income stabilization (Demissie and Legesse, 2013).

In tandem to this, Chayanovian models represent peasant households making tradeoffs between consumption and labor (Chayanov, 1986), generally seeking to minimize the "drudgery" of work. Here, income maximization is not a priority; rather, households seek to satisfy some level of consumption with the minimal required labor input.

Given this, we assume that all agents' primary concern is the satisfaction of their own food consumption requirements (i.e., fundamentally, agents wish to be food secure). Above this, we assume that each agent has a preference for either wealth accumulation or leisure (i.e., given an agent is not at risk of being food insecure, they seek to maximize either their wealth or leisure). Food insecurity is represented as a perceived number of months with food shortages. Wealth is composed of cash, food stores, and livestock holdings converted to a monetary value.

To operationalize this into an objective function for the agents, we draw from Kaufman's (1990) theory of satisficing, and our agents have two levels of objective:

1. First, agents restrict their considered options to those that they believe will lead to the lowest level of food insecurity.
2. If more than one option remains, agents then choose the option with the highest expected utility.

Since decisions are made under uncertainty, the wealth and food security components of a decision option both represent subjectively-perceived values. In addition, it is well-documented that smallholder farmers are highly risk averse (Devereux and Sussex, 2000; Wossen et al., 2015; Jumare et al., 2018) and that this can affect their decision-making (e.g., leading to so-called "poverty traps" (Dercon and Christiaensen, 2011; Barrett and Carter, 2013)). Our model assumes that there is no uncertainty in the allocation of labor. Hence, agents with leisure preferences simply maximize free time in step 2. For agents with wealth preferences, the returns to each decision option are uncertain. These agents calculate the expected utility of the wealth component of a decision option d using the following exponential utility function:

$$EU = \sum_{i=1}^{n} w_i * (1 - exp(\frac{-Wealth_i}{R})) \tag{B.15}$$

where w_i is the weight (trust) that the agent places on yield forecasting model i (see section B.7.4), and R is the agent's risk tolerance parameter, the overall distribution of which is determined by the genetic algorithm.

When evaluating the value of a livestock herd, a net present value (NPV) is calculated. Empirical evidence suggests that rural households in Ethiopia have very high discounting rates, and drawing from Holden et al. (1998), a discount rate (over a single year) of 71% is used [2].

Decision options: At the beginning of each year, agents make decisions about land, labor, and livestock. The livelihood options considered are all feasible combinations of:

- Land-use - $\{farm, not\_farm\}$. This assumes that agents either plant crops on all or none of their plots.
- Fertilizer - $\{yes, no\}$. Agents can choose to fertilize their fields. This is a binary decision.
- Livestock - $\{sell, do\_nothing, buy\}$. Here, *sell* represents an agent selling half of their existing herd, and *buy* represents buying the maximum number of livestock possible, given their cash and labor availability.
- Non-farm - $\{0, 25, 50, 75, 100\}$. This represents the percentage of remaining labor that is allocated to non-farm wage labor.

Additionally, agents choose the time at which they plant their crop (either early, normal, or late). This decision is given by their perceived probabilities of rainfall onset time. For example, if $P(Late) = 0.2$, there is a 20% chance that they will choose to plant late. Shifting of planting date is a common ex-ante adaptation mechanism employed in an attempt to reduce the impacts of climatic shocks (Bryan et al., 2009; Amare and Simane, 2017). Policies such as the provision of seasonal climate forecasts can help to give farmers more information about the expected timing of rainfall onset (Ziervogel et al., 2005; Luseno et al., 2003).

The livelihood options considered by the agents represent the major livelihood components in the Ethiopian highlands and are similar to previous agent-based modeling efforts (Berger et al., 2017). With respect to agriculture, agents have the option to apply fertilizer to their fields. Fertilizer

[2] In the survey carried out by Holden et al. (1998), respondents were asked how much money they would have to be given today to subsitute for 100 birr in a year's time. In Ethiopia, the average present equivalent was 58.6 birr, which is equivalent to a discount rate of 70.6%

use is relatively high in Ethiopia (for example, 75% of maize fields in Ethiopia were fertilized in 2015 (CSA, 2017)). We therefore do not model fertilizer choices using an explicit technology diffusion model similar to other studies (Kiesling et al., 2012; Swinerd and McNaught, 2014). Households also have the option to completely disengage from agriculture and leave their fields unfarmed. Given the growing population pressure in Ethiopia, this is relatively uncommon in practice (and indeed uncommon in the model), and households who do not themselves farm are more likely to lease or rent their land out to others or engage in some form of sharecropping (CSA, 2017). However, we do not believe that these practices would enhance the achievement of the objectives of this paper and so have excluded them from our model. Additionally, as we are not modeling land degradation, we do not include notions of leaving plots fallow or engaging in land conservation measures in our model. This does not mean that we consider these decisions irrelevant to long-term resilience of agricultural communities.

We have chosen to model a single crop. In reality, farmers often grow a diverse portfolio of crops in order to minimize risk. However, in the modeled region, soil type and quality pose a major constraint to crop production (Simane et al., 2013), and it is rare for farmers to switch crop types. Hence, for the purposes of this paper we have made this simplification. If anything, this leads to an underestimation of resilience.

Decision process: The decision-making process proceeds as shown in Algorithm B.1 (Appendix B.2). For each livelihood option, agents begin by sequentially allocating their labor between different activities. First, labor is allocated to domestic activities (water and firewood collection). The time required for each of these activities is constant throughout the year and over the simulation, but is heterogeneous over the agent population. Time spent on these activities was drawn directly from the LSMS survey data when making the artificial populations of agents. Second, agents allocate time to agriculture (if relevant). Agricultural labor requirements vary throughout the year, and are divided into land preparation, planting, weeding, and harvesting. Labor requirements scale linearly with land area, but we assume that plowing requirements are 50% lower for agents that have at least two livestock (loosely drawn from (Lawrence et al., 1997)). This indirectly builds in a preference for the leisure maximizing agents to maintain livestock herds. Third, labor is allocated for livestock grazing. The LSMS contains no specific questions about livestock-related labor allocations, so we assumed that labor requirements followed a simple linear model with a herdsize-dependent value (slope). The coefficient for this model was determined by the genetic algorithm. Finally, any remaining labor is allocated to off-farm wage-based labor based on the percentage of non-farm labor allocation imposed by the decision option in question (ranging from 0% to 100%). The labor allocation for the above activities is not a decision per se, as the requirements for domestic and farm labor are fixed for each agent and the requirement for livestock is fixed, given the herd size.

After allocating labor for the decision option being assessed, agents then estimate their monthly wealth and food security status under each potential yield realization (given by the yield belief forecasting models (section B.7.4)). They do this by assessing whether they will be able to satisfy their consumption requirements through their own food stores (from previous crop harvest), expected returns from livestock, and food purchase from the market. Monthly fixed expenditures are imposed, which are higher for larger and wealthier households [3]. Expected wealth is then calculated at the end of the year as the discounted value of all livestock assets, stored crop, and cash. Using this expected wealth, wealth-maximizing agents then calculate an overall expected risk-averse utility (Equation B.15) for the option. Finally, agents select the option that they believe will lead to the lowest level of food insecurity. If there is a tie, the remaining option with the highest utility (based on either wealth or leisure) is selected.

Uncertainty: Because the concept of resilience is inherently related to uncertainty, it is necessary that agent-level uncertainty is specifically represented in the decision-making model. This is often ignored in agent-based models of agriculture. Agent-level uncertainty is represented through beliefs about crop yields. When evaluating the overall utility of a decision-option, each wealth-maximizing agent calculates a utility for each of the different belief forecasting models. These separate utilities are then weighted by the agent's *trust* in each forecasting model to give an overall expected utility. Refer to section B.7.4 for an overview of the representation of agent beliefs.

Although agents also have different levels of trust in the forecasting models for their beliefs on prices and livestock growth rates, only the expectation of these beliefs is evaluated in their decision-making. This simplification has been made for computational reasons, as evaluating the probability of each distinct forecasting model combination (e.g., $(price, yield, livestock)$ = $(linear\_trend, neighbor, last\_period)$) would involve m^3 distinct computations (where m is the number of forecasting models).

Discussion: Clearly, human decision-making is incredibly complex, heterogeneous, and context-specific. We (as a society) are far from a universal, causal mathematical decision model, and the decision-making model in the ABM is far from perfect. Rationality and utility maximization have long been criticized as inaccurate and founded on erroneous assumptions (Simon, 1955, 1979; Kahneman and Tversky, 1979). However, utility maximization remains the primary decision methodology implemented in previous ABMs of agriculture (Kremmydas et al., 2018). This has also been subject to criticism (Schulze et al., 2017; Malawska and Topping, 2016; Levine et al., 2015). Ideally we would conduct specialized farmer interviews to shed light on the decision-making processes relevant for representing livelihoods and household-level responses to shocks.

[3] We fit the following linear model using the 2015 LSMS data: $[cost] = 891.3 + 18.7[nonag\_hrs\_pw] + 305.1[adulteq] + 83.6[livestock]$, where $cost$ is the annual fixed cost for the household, $nonag\_hrs\_pw$ is the amount of hours the household spends per week on non-agricultural (wage-based) activities, $adulteq$ is the household size in adult equivalents, and $livestock$ is the number of livestock owned.

This would give us a better understanding of (1) which alternatives are considered and by whom, and (2) how options are evaluated and assessed against other alternatives. However, absent of such information, and to not become overly distracted from the purpose of this model (to demonstrate the assessment of agricultural resilience), we have opted for the above framework that incorporates many of the major livelihood processes and models decision-making using a bounded rationality, satisficing framework. Future work could refine and test the sensitivity of the model to the decision-making process to provide more realistic policy-relevant conclusions for Ethiopia.

B.7.6 Food consumption and food security

In the ABM: Food requirements are expressed with respect to calorie intake. The baseline per-capita calorie requirements are set at 2,200kCal/day. This is in line with (CAADP, 2013), as well as values used in previous Ethiopia-specific studies (Clay et al., 1999; Bogale et al., 2005). Since the model is focused on cereal crop production, the assumption is made that agents receive all of their calories from cereal crops. This is a simplification, and ideally dietary diversity would be explicitly modeled to give a more comprehensive picture of food security, but this would complicate the model and the simplification was deemed to be appropriate for the purposes of this study. The per-capita calorie requirement was converted into a weight-based crop requirement using cereal crop calorie densities were taken from the USDA Food Composition Database [4]. This resulted in a monthly per-capita requirement of 0.18 quintals (18 kg) of cereal crop.

Generalizing: The FAO's (2008) widely-employed definition defines food security as consisting of four pillars: (1) The *availability* of an adequate quantity of food (i.e., sufficient supply/production); (2) adequate *access* to this food, both in a physical and monetary sense; (3) sufficient biological *utilization* of any consumed food (which could be compromised by, for example, gastrointestinal viruses); and (4) *stability* of the above three factors over time. The ABM incorporates elements 1, 2, and 4. Food *availability* is influenced by farming decisions and climate, which interact to influence household-level yields. *Access* to food is explicitly represented, as agents must have the (financial) means to procure food for their consumption. Additionally, transport and transaction costs are included in the model, reflecting the fact that buying and selling prices are not necessarily equal. The ABM allows for measurement of household-level and regional food insecurity on a monthly basis, so also incorporates the *stability* pillar. *Utilization* of food is beyond the scope of the model and is an interesting avenue for future work (Nicholson et al., 2019).

Empirically, the measurement of food insecurity is an elusive concept (Barrett, 2010), and there exists a wide variety of food insecurity indicators. One of the oldest (and simplest) indicators of food security is calorie intake and deprivation. However, the relevance of this to food security

[4] https://ndb.nal.usda.gov/ndb/

outcomes (e.g., stunting and wasting) has been subject to debate (Headey and Ecker, 2012). Other methods, such as dietary diversity, are argued to be more robust food security indicators (Headey and Ecker, 2012; Leroy et al., 2015; Pangaribowo et al., 2013), however its quantification is more complex and involved than other measures. Given that the ABM only models a single crop, we represent food security with respect to calorie intake. Cereals make up the majority of consumption in rural Ethiopia (Worku et al., 2017), and the majority of agriculture is cereal-based, so this is an appropriate place to start. Future efforts could increase the number of crops represented in the model, and work towards quantifying food security with respect to dietary diversity or another more complete measure.

B.7.7 Empirical parameterization - genetic algorithm

Overview: An *indirect calibration* approach (Windrum et al., 2007) was taken to parameterize the model empirically. Where possible, parameters (e.g., household size, number of plots, labor availability) were taken directly from empirical data. However, a number of uncertain parameters remained. We employed a parameterization approach using a genetic algorithm (GA).

Uncertain parameters: Table B.4 identifies the uncertain parameters. For each of these, a wide range of plausible values was specified. The purpose of the GA was to refine the range of these uncertain parameters by identifying regions of the parameter space that lead to model outputs in accordance with empirically-observed data.

Empirical data: The Living Standards Measurement Study (LSMS) has been conducted in three waves in Ethiopia (2011/2013/2015) and provides household-level demographic, agricultural, and consumption data that is representative at the regional level (e.g., Amhara). We generated five distributions from the LSMS data to be used for the parameterization procedure (refer to Figure 3.5). These distributions were generated using the 2015 LSMS data, subsetted to Amhara. The distributions selected for fitting were chosen because of (1) their direct mapping to emergent outcomes of the ABM, and (2) their relevance for the assessment of resilience. For example, labor allocated to non-farm activities is not something imposed by the structure of the ABM, but emerges as a result of the interaction between the different processes. Similarly, the distribution of food security throughout the population is not imposed, but emerges, and forms the basis of the assessment of resilience. Ideally, for the most robust validation of the model, we would compare model outputs to empirically-observed effects of shocks on food security. However, as far as we are aware, these data are not available. Hence, the best we can do is to compare food insecurity measures from the LSMS with outputs from the ABM run under historical climatic conditions.

Objective function: The goal of the GA is to optimize an objective function by finding the combinations of parameter values that lead to the best objective values. Here, our goal was to both

	Name	Unit	Range	Description	Fitted value
Agent properties					
1	min_sust	%	$0.2, 0.7$	Sustenance deficit threshold for "severe" food insecurity	0.65
2	$frac\_income\_max$	fraction	$0, 1$	Fraction of agents with income maximization preference	0.21
3	$risk\_aversion\_mult$	-	$1, 1000$	Multiplier on risk aversion coefficient	834
Market					
4	$job\_availability$	hrs/mo/ag.	$0, 100$	Job availability	2.65
5	$crop\_capital$	birr/ha	$0, 500$	capital cost required to engage in agriculture	282
6	$labor\_wage$	birr/day	$8, 100$	Labor wage per day	67.4
Labor					
7	per_cap_labor	hrs/pp/day	$5, 13$	Labor availability	7.21
8	$farm\_labor\_mag$	-	$100, 1600$	Multiplier on farm labor requirements	247
9	$ls\_per\_head\_herding$	hrs/head/mo	$0, 50$	Livestock labor requirements	26.3
Environment					
10	$planting\_fraction$	fraction	$0, 1$	Fraction of land that is planted	0.90
Rangeland					
11	$grass\_regen$	-	$1, 5$	Grass regeneration rate	2.32
12	$rf\_intercept$	-	$-1, 1$	Rainfall effect with no rainfall	-0.91
13	rf_slope	fraction	$0.001, 0.004$	Change in rainfall effect per mm of rain	0.0013
14	ls_max_growth	fraction	$0.2, 0.5$	Unconstrained livestock reproduction rate	0.45
15	g_max	kg/ha	$400, 10000$	Maximum grass biomass	1077
16	$g0$	kg/ha	$400, 4000$	Initial grass biomass	1001

Table B.4: Uncertain parameters refined by the genetic algorithm. All parameters were generated uniformly over the specified range.

minimize the difference between the empirically-observed (E) and ABM-generated (A) distributions and have the model recreate several desired qualitative patterns (Table B.5). We refer to this overall difference as the "loss". For each distribution, we calculate the loss using a measure of comparability between the E and A histograms. To describe this, we refer to x as the histogram value (e.g., 5 livestock) and y as the height of the histogram bin height (e.g., 10% of households). The histogram loss is composed of two components. The first component (L_1) is equal to the sum of the squared differences between each histogram bin height in E and A:

$$L_1 = \sum_{x=1}^{X} (y_x^A - y_x^E)^2 \tag{B.16}$$

where X denotes the maximum x value in E. The second component of the loss is an additional penalty added when A contains data that is *outside* the bounds of E (e.g., if the ABM produced livestock herd sizes greater than any empirically-observed values). This is calculated as:

$$L_2 = \frac{1}{N_{agents}} \sum_{agents: x > X} \left(\frac{x - X}{X} \right)^2 \tag{B.17}$$

These two loss components are added to give an overall loss for each distribution. The losses

	Description
1	The grass biomass does not decrease to zero at any time
2	At least 70% of agents engage in farming on average
3	There is uncertainty in the job market: the probability of getting a non-farm job is always less than 1
4	No agent ends the simulation with more than 80 livestock

Table B.5: Qualitative patterns used for the genetic algorithm

Name	Value
Number of generations	200
Population size	38
Chance of mutation	2%
Type of crossover	Uniform
Type of selection	Binary tournament

Table B.6: Genetic algorithm parameters

are then summed over all distributions.

For each qualitative pattern, a large value (10) is added to the objective function if the model does not generate the pattern. This is a large penalty in comparison to the histogram losses; thus, the GA prioritizes models that generate the qualitative patterns, then seeks to maximize the histogram fits.

For the fitting, the ABM was run from 2003-2015, and the histogram loss functions were computed by taking the average of the final three years' of data in the ABM.

Genetic algorithm: The genetic algorithm settings are displayed in Table B.6. The GA was initialized with a population of "solutions", each characterized by a randomly-generated set of parameters (generated uniformly over the ranges defined in Table B.4). For each of these solutions, the ABM was run and the loss computed. We only ran a single replication of the ABM for each potential parameter set as we found that changes in model parameters contribute much more variability than model stochasticity itself.

After calculating the loss, *parents* are identified through a binary tournament selection (repeatedly randomly identifying two individuals and selecting the fitter of the two to be a parent). A population of parents equivalent to the original population size is created, these parents are paired, and each set of parents generates two children, combining their parameters (genes) using uniform crossover. There is a chance that each gene will undergo mutation, whereby a random value is generated from the defined uniform distribution.

These children then comprise the subsequent "generation" in the algorithm. The ABM is run for the new population, loss values are calculated, and further selection takes place. The GA is run for a pre-defined number of generations.

Results: Table B.4 shows the resulting restricted parameter values and Figure 3.5 in the main body of the paper compares the empirical distributions with the ABM-generated distributions under the final parameterization.

B.7.8 Model convergence

Since the ABM is stochastic, another factor to consider to give confidence in the validity of any estimates generated by the model is convergence. A sufficient number of model replications must be conducted so that the estimates are not excessively influenced by randomness. The required number of simulations, n^*, was determined as suggested in Law (2008), by conducting an initial number of simulations (n) and then determining n^* as:

$$n^*(\gamma) = min\Big\{i \geq n : \frac{t_{i-1,1-\alpha/2}\sqrt{S^2(n)/i}}{|\bar{X}(n)|} \leq \gamma'\Big\} \tag{B.18}$$

where $S^2(n)$ is the variance of the mean in the initial n replications, $\bar{X}(n)$ is the mean estimate in the initial n replications, and $\gamma' = \gamma/(1+\gamma)$ is the adjusted relative error. The outcome of interest for the model convergence testing was the overall estimated resilience of a simulation (i.e., the area above the resilience curve).

Appendix C

Supplement to Assessing Model Equifinality

C.1 EAGA hyperparameter experiments

The specification of hyperparameters for the EAGA presents tradeoffs between model fit, model diversity, and computational burden. Practical considerations will likely influence decisions in most applications, and different ABMs will respond differently to hyperparameter changes. Thus, we do not present concrete guidelines for hyperparameter selection here, but qualitatively explore the effects of hyperparameter changes in our case study application. The experiments we conducted are summarized in Table C.1. Due to stochasticity in the initialization of the EAGA solutions, we repeated each experiment three times (i.e., three different random number seeds). Computational requirements prohibited us from running more replications, but we believe our results reveal the general trends.

C.1.1 Number of subpopulations (N_{SP})

The primary effect that we expect as N_{SP} is increased is that the solutions will reduce in diversity (i.e., the configuration space becomes more crowded). We observe this effect, with a dramatic decrease in diversity above three SPs, and then a progressive decrease beyond this (Figure C.1A). Increasing beyond 20 SPs has little effect on the solution diversity, suggesting that the additional solutions in the 25SP and 30SP experiments populate different regions of the configuration space.

Varying N_{SP} had little effect on the fitness to the empirical data in the selected solutions (Figure C.1B) or the spacing within each SP (Figure C.1C). This suggests that, for all N_{SP} considered, the EAGA is able to find a set of reasonable solutions and that 250 generations are sufficient to converge in the objective function value for these selected solutions.

Finally, at all N_{SP} above three, there were some SPs that contained no feasible solutions (Figure C.1D), thus reducing the number of model configurations that can be selected. This could suggest

Table C.1: Number of model replications required for hyperparameter experiments.

Experiment	**# gens**	N_{SP} **†**	pop_size	**# seeds**	**# model runs**
N_{SP}	250	3	20	3	45,000
		5			75,000
		10			150,000
		15			225,000
		20			300,000
		25			375,000
		30			450,000
					1,575,000
pop_size	250	5	10	3	37,500
			20		75,000
			30		112,500
			40		150,000
					375,000

† Note that this value includes the master SP

issues with the genetic algorithm getting stuck in local minima that do not fit the data well enough to be considered feasible. A larger number of GA replications may help these SPs to escape these regions. Because the fraction is not strictly increasing in N_{SP}, this demonstrates that the configuration space is not yet fully "saturated" even with 30 SPs; there likely exist more model configurations that could satisfy the feasibility requirements.

Choosing the "appropriate" number of SPs, given this information, is clearly a subjective decision. As stated in the main body of the paper, practical considerations also have to be considered (e.g., computational requirements or ease of communication). For ease of display and to enable investigation of each individual SP, we chose four SPs (five, including the master) in the main body of the paper. However, policy-relevant applications investigating the robustness of policy interventions may choose to use a larger number of SPs to more fully cover the model configuration space.

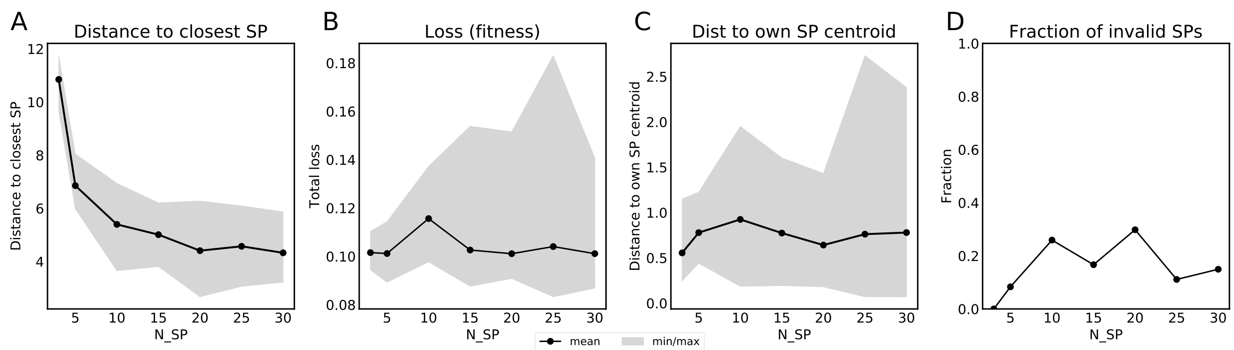

Figure C.1: Influence of the number of subpopulations (N_{SP}) on EAGA convergence. The plotted values in A, B, and C represent the minimum / maximum / mean over the selected final solutions (i.e., the most diverse feasible solution in each SP). D shows the fraction of SPs that contained no feasible solutions.

C.1.2 Population size

The results reveal a slight tradeoff between solution diversity and objective function value as population size is varied; small population sizes are unable to achieve as good fits to the data within the modeled number of generations (Figure C.2B and D), yet foster slightly more diverse solutions (Figure C.2A). The progressive reduction in diversity is because larger population sizes result in a greater degree of "smoothing" of the centroid locations as the parameter values of a larger, more diverse (Figure C.2C) population of solutions are averaged.

Again, choosing an appropriate population size is a subjective decision. We chose to use a population size of 30 in our application because it achieves a good balance between diversity and fitness, yet does not have issues with invalid SPs (Figure C.2D).

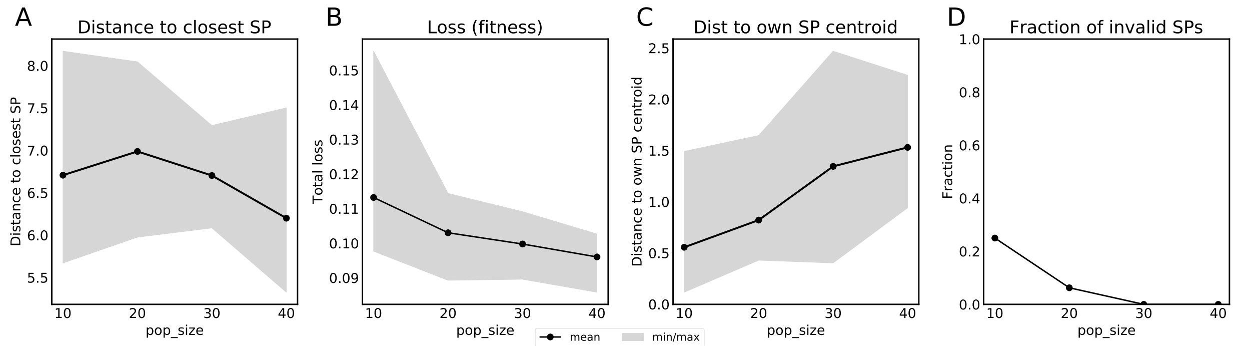

Figure C.2: Influence of the population size within each SP on EAGA convergence. The plotted values represent the minimum / maximum / mean over the selected final solutions (i.e., the closest feasible solution to each SP centroid).

C.2 Comparison of ABM-generated and empirical patterns

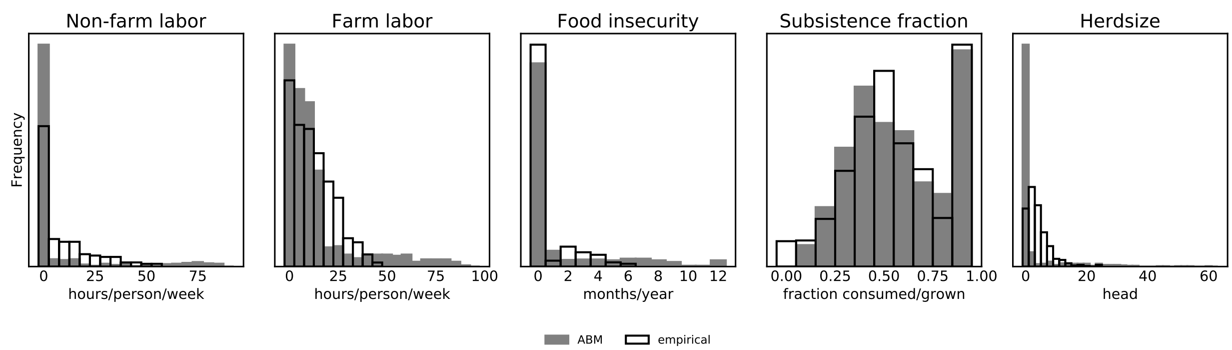

Figure C.3: ABM-empirical comparison for SP 2.

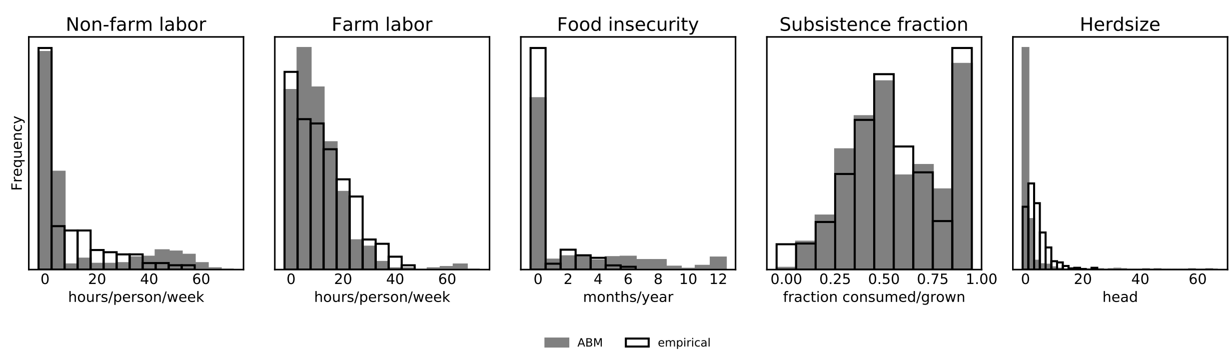

Figure C.4: ABM-empirical comparison for SP 3.

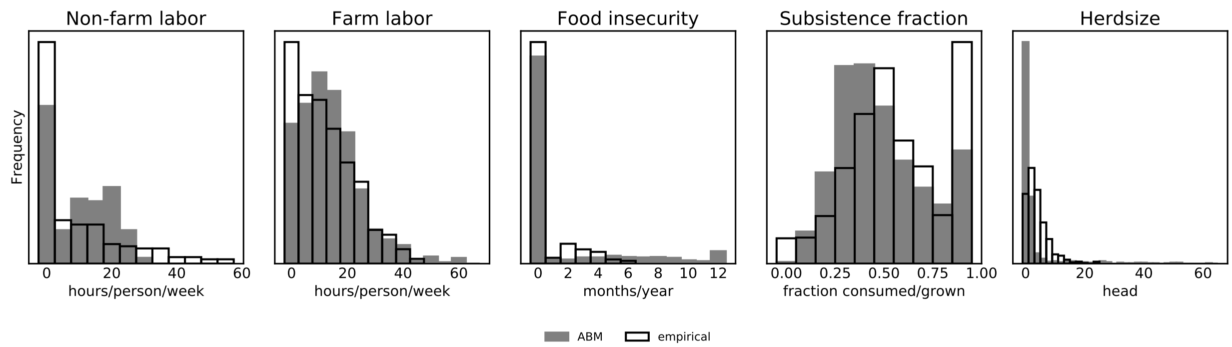

Figure C.5: ABM-empirical comparison for SP 4.

C.3 Resilience analysis with a 20% drought

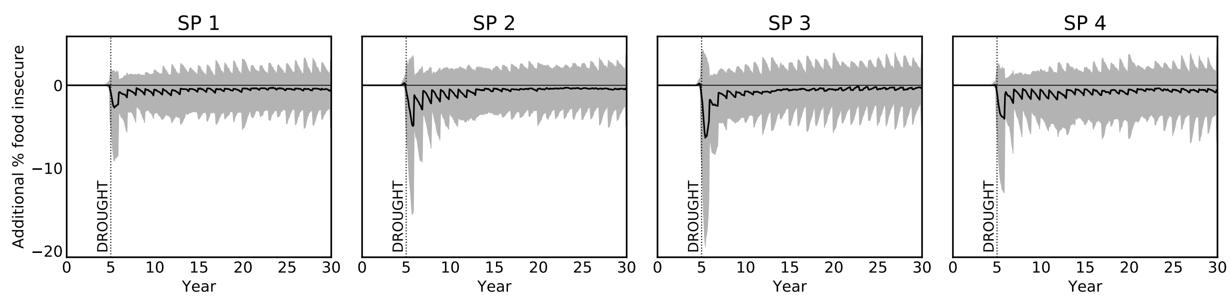

Figure C.6: Effect of a 20% drought on household food security.

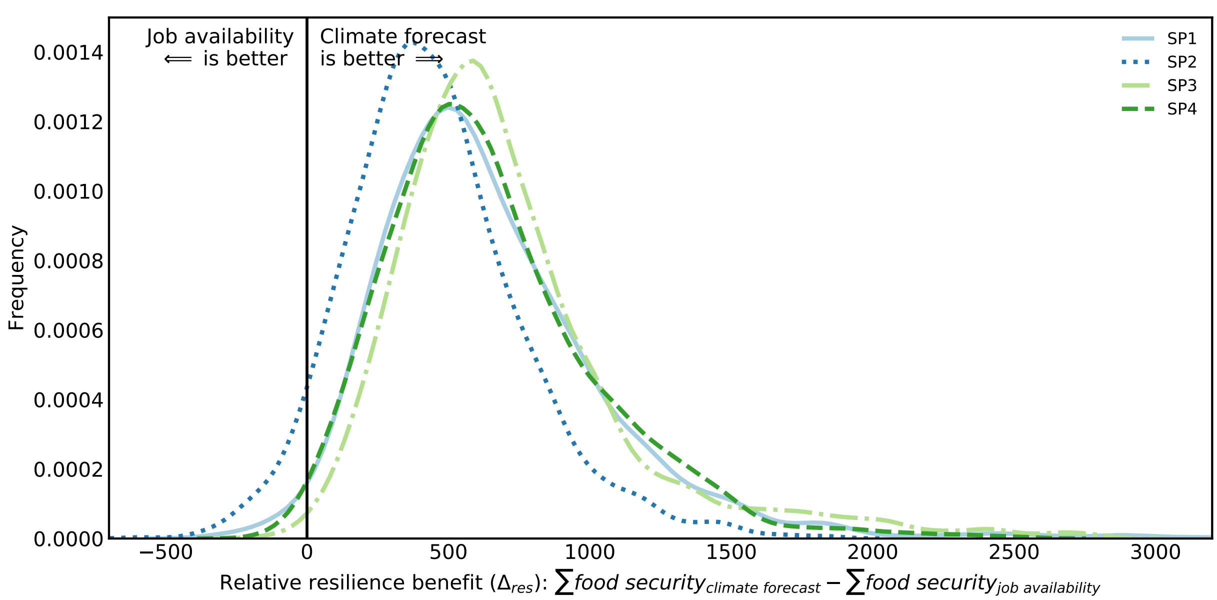

Figure C.7: Relative benefit of the two interventions under a 20% drought.

C.4 Case study experiments with nine SPs

C.4.1 Diverse model calibration

An experiment with nine SPs yielded similar results, but one of the SPs did not become feasible after 300 generations of the EAGA (Figure C.8). As a result, only eight parameterizations were selected for the resilience analysis.

The final parameterizations in this experiment are on average around five units apart in the normalized parameter space (Figure C.8B), which is slightly less diverse than the experiment with four SPs and indicates a crowding of the parameter space. However, there are now two models that contain the satisficing decision-making representation (SP1 and SP3; Figure C.9) and a greater overall diversity in some parameters such as the planting fraction.

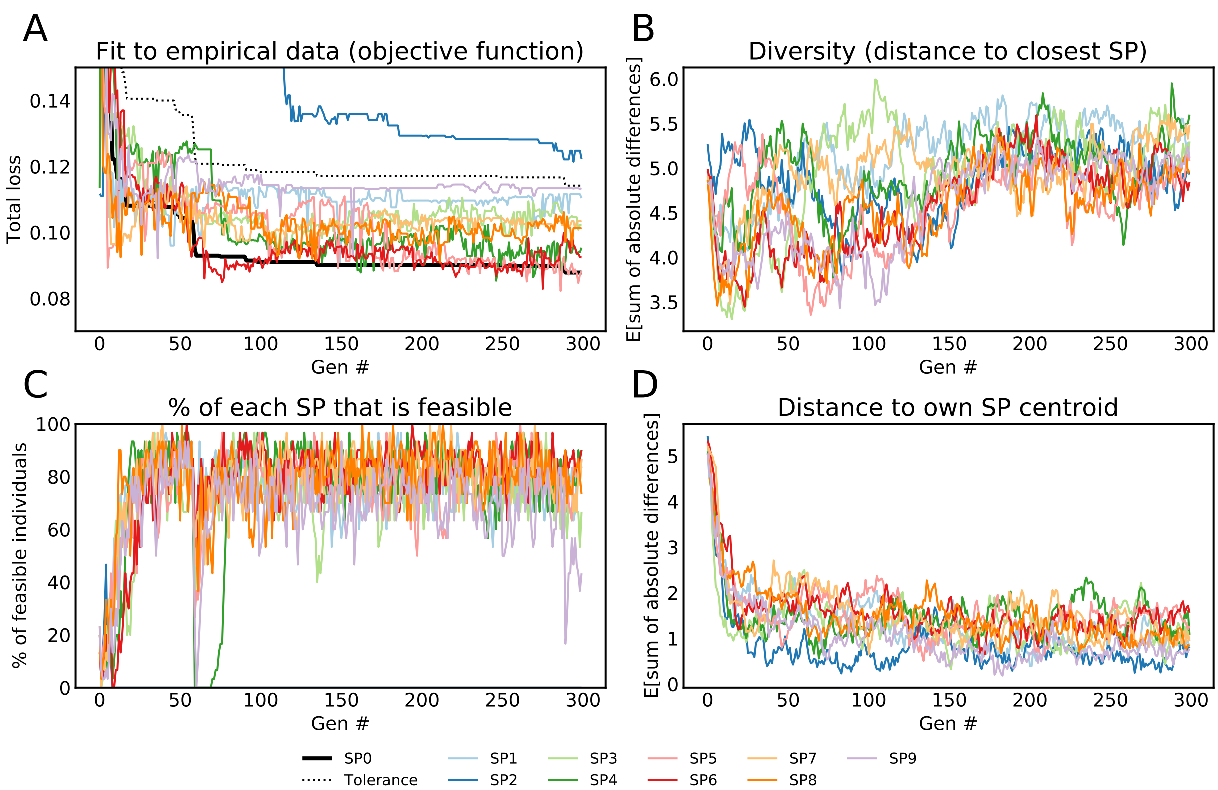

Figure C.8: EAGA convergence measures with nine SPs.

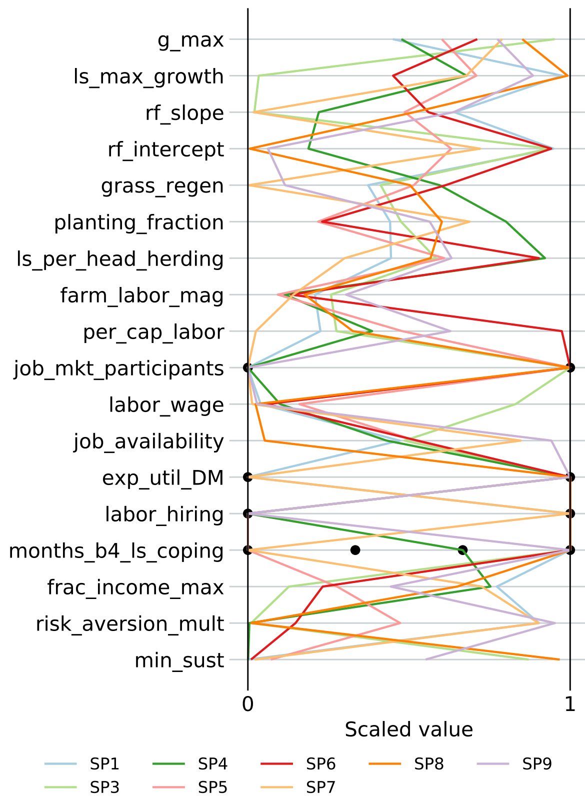

Figure C.9: Normalized parameter values with nine SPs. Note that there were no feasible solutions in SP2 so it is not displayed here.

C.4.2 Resilience analysis

The greater overall model diversity leads to a greater diversity in the effects of drought on household food security (Figure C.10). In particular, the maximum effect is not experienced until several years after the drought in SP3, and there is a permanently decreased level of food security in SP1.

The relative benefits of the interventions are also more variable over the eight selected models (Figure C.10). Five of these models show results similar to those in Figure 4.10 in the main body of the paper, while SPs 3, 4, and 9 show some evidence of a greater benefit in increased job availability. However, in seven of the eight retained parameterizations, provision of climate forecasts provides larger benefits to resilience than increased job availability.

SP3 displays the most discordant behavior of the SPs. The main way in which SP3 differs from the other SPs is in the livestock reproduction rate (ls_growth_rate; Figure C.9); SP3's lower livestock reproduction rate contributes to a slower recovery time in the wake of a drought (Figure C.10), as more livestock must be purchased to recover herd sizes and, in turn, food security. As a result of this, the additional financial capital provided by an increased job availability yields larger benefits to food security than the climate forecasts, which only provide financial capital indirectly through better agricultural management.

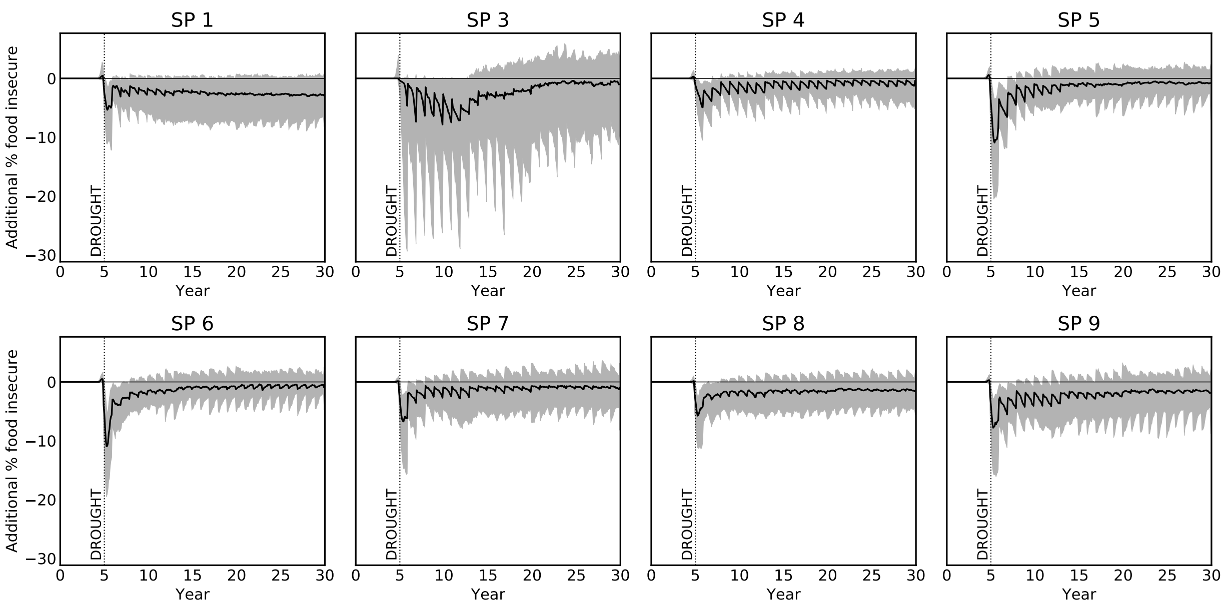

Figure C.10: Effect of a 50% drought on household food security with nine SPs. Note that no solutions in SP2 were feasible so it was not included in the resilience analysis.

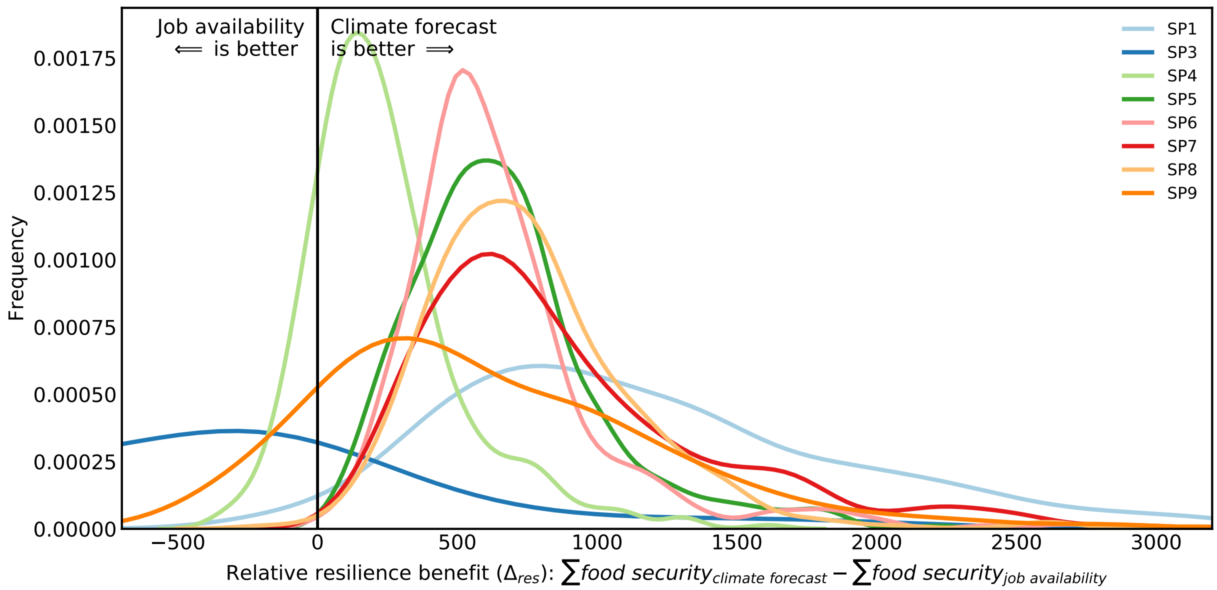

Figure C.11: Comparison of policy effects with nine SPs.

C.5 Sensitivity of selected model configurations

We conducted a univariate sensitivity analysis to assess the stability of the selected model configurations. To do so, we sequentially perturbed each of the continuous parameters from their calibrated values and assessed the effect of this on the fit to the empirical data. We did not include the model structural elements or the categorical parameter in this assessment. We note that a univariate sensitivity analysis does not capture potential dependencies between parameters (Lee et al., 2015), and, as such, we view these results in an exploratory manner.

The results reveal a wide variability in the stability of the models to parameter perturbations (Figure C.12). SP1 and SP2 are in general more stable parameterizations than SP3 and SP4 (i.e., their shaded bands in Figure C.12 are wider). Additional experimentation confirmed that the points at which the decline in fit increased above 20% were generally due to the model failing to generate one of the qualitative patterns, thus having a value of one added to its loss. SP3 and SP4 were each on the verge of not recreating one of these patterns, so small deviations in many of the parameters resulted in a dramatic decrease in calculated fit.

Some parameters (particularly rf_slope, $rf\_intercept$, and $risk\_aversion\_mult$) exert very little influence on the models' fit to the data. Given these results, these parameters could likely be excluded from the DMC process with little effect on the results. This kind of sensitivity analysis could inform the iterative development of a model by sequentially excluding the parameters that the model is not sensitive to.

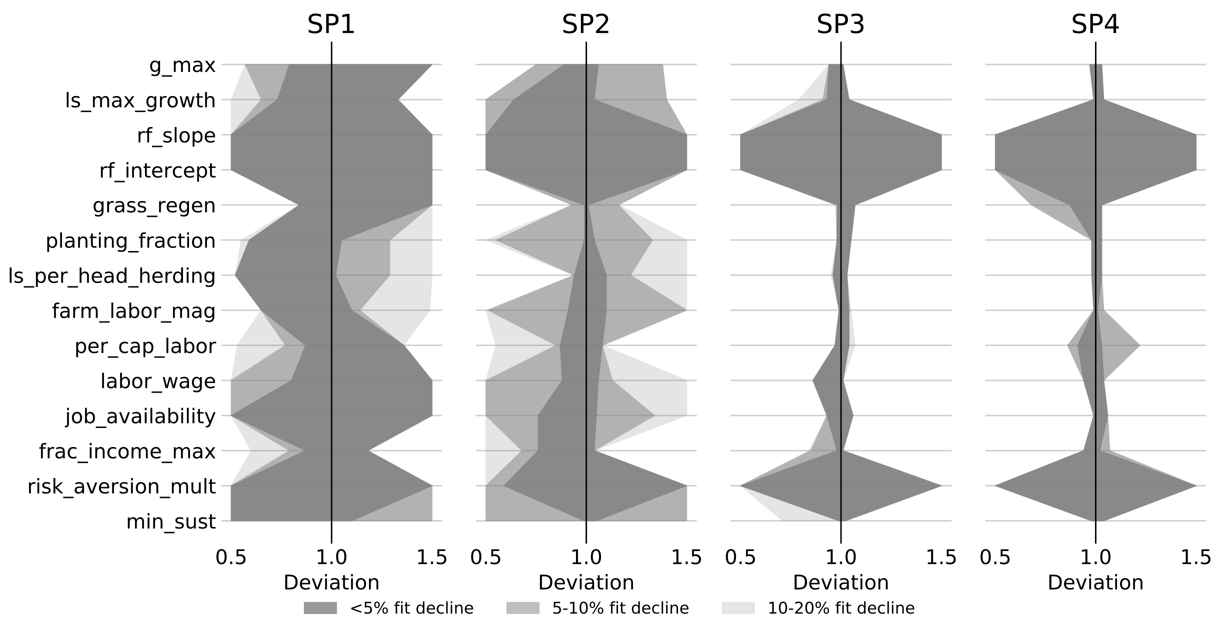

Figure C.12: The stability of each SP configuration. The shaded areas encompass the amount of deviation each parameter can undergo before the model's fit declines by the specified amount. Wider shaded bands indicate more stable configurations. A value of 0.5 on the horizontal axis represents a parameter being set to 50% of its scaled calibrated value.

Appendix D

Supplement to Ecological and Financial Strategies

D.1 Additional figures

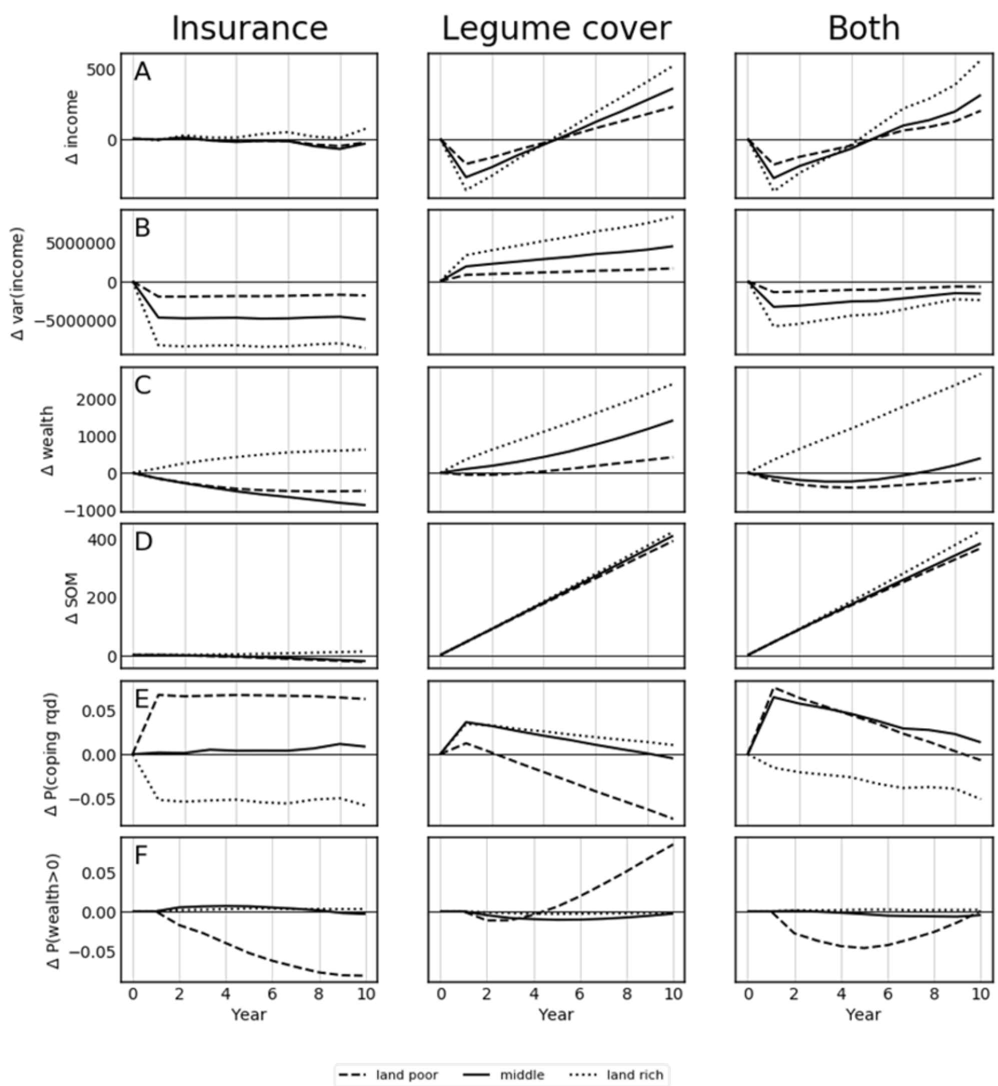

Figure D.1: Effects of the strategies on various model characteristics relative to the baseline scenario for each type of household. In all cases, the horizontal line at zero represents no change relative to the baseline model conditions. "Change in P(coping rqd)" refers to change in the probability that a household must sell their livestock at each time step. "Change in P(wealth ¿ 0)" refers to change in the probability that a household has positive wealth (i.e., livestock) at each time step.

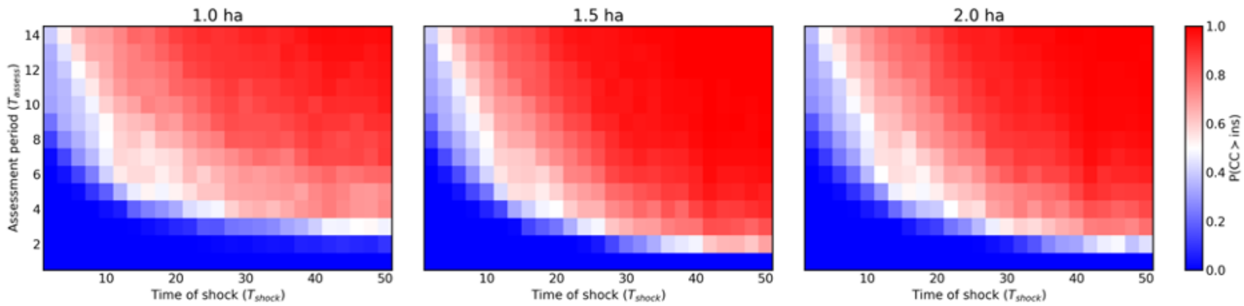

Figure D.2: Comparison of insurance and cover cropping on $P(CC \succ Ins)^{shock}$ for the three types of household, which differ solely in their land holdings. Land-poor households have 1 ha of land, middle households have 1.5 ha, and land-rich households have 2 ha.

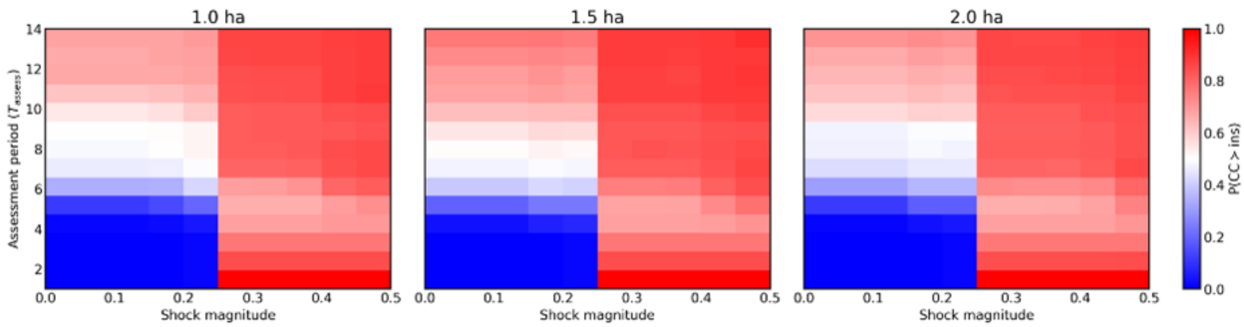

Figure D.3: Comparison of insurance and cover cropping on shock absorption as the magnitude of the drought is varied, with $T_{shock} = 10$. The vertical threshold at 0.25 represents the microinsurance climate index.

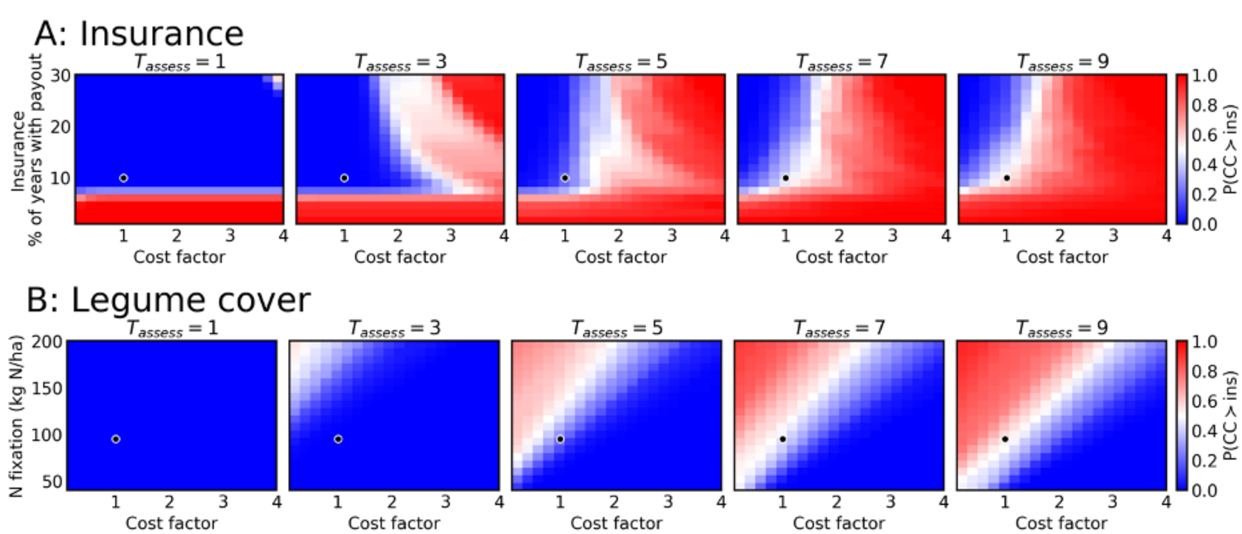

Figure D.4: Influence of strategy characteristics on the shock absorption comparison. The black dots represent the baseline settings used in other experiments. In all cases, we simulated a 0.2 magnitude shock with $T_{shock} = 10$ and averaged results over all household types. Results were qualitatively similar for each individual household type.

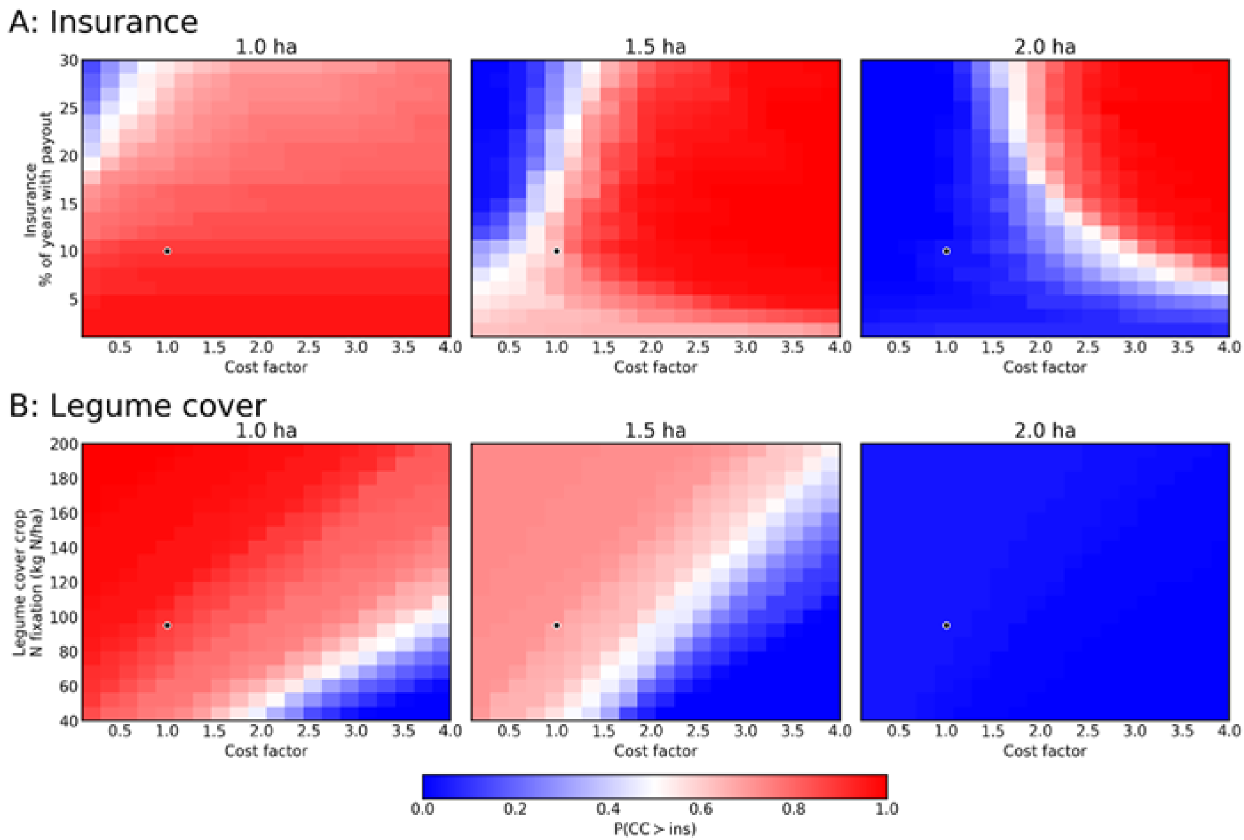

Figure D.5: Influence of strategy characteristics on the poverty reduction assessment for different household types with $T_{dev} = 20$. Note that poverty reduction measures households that have lost all their wealth; since the land-rich households (2 ha) very rarely lose their wealth even under baseline conditions (Figure 5.5), the stark differences seen in this assessment (right-most plots) for these households are not meaningful.

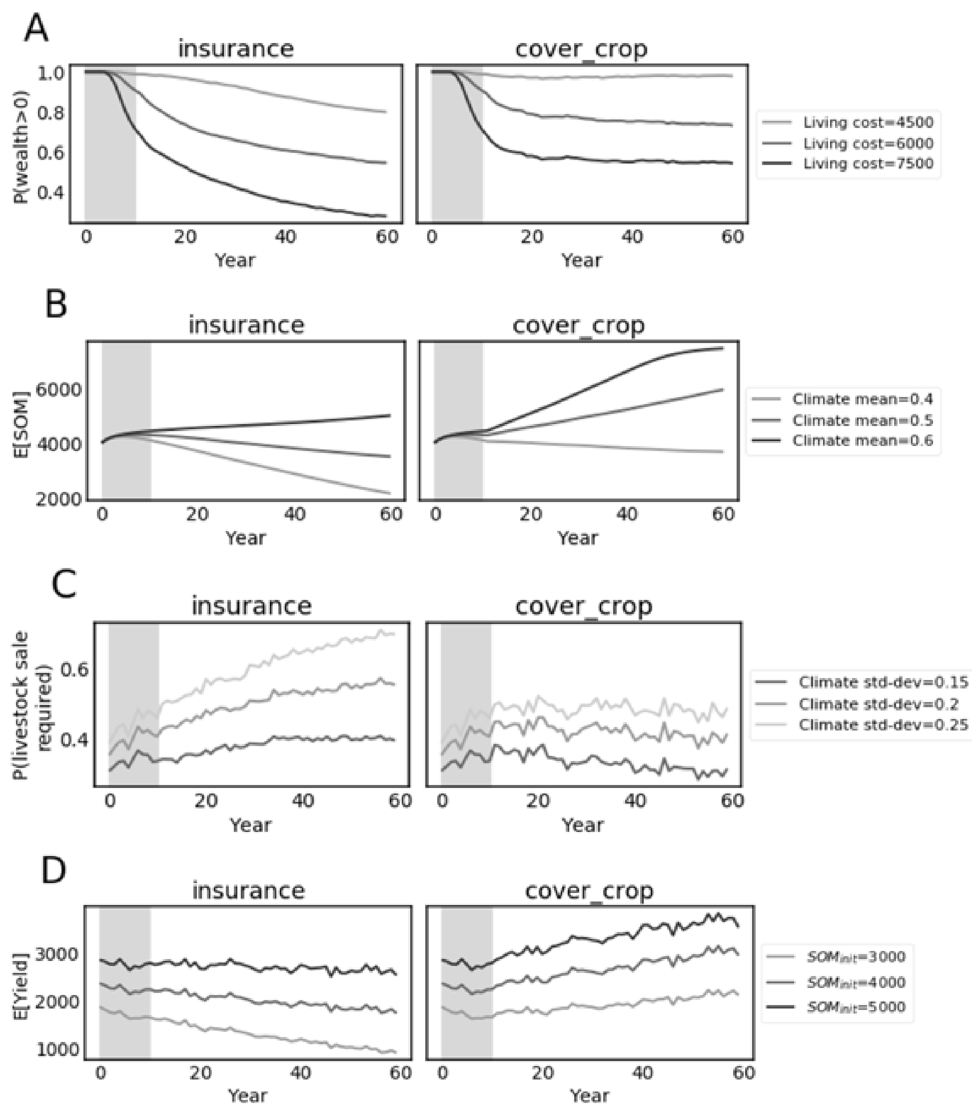

Figure D.6: Evolution of model metrics under different parameter settings. Plotted values represent averages over all household types. "P(livestock sale required)" represents the probability that *any* livestock sale is required, independent of the number.

D.2 Utility analysis

The focus in the main body of the paper centered primarily around the two measures of resilience: shock absorption and poverty reduction. Our results showed that—–predicated on the structure of the model and scenarios—–cover cropping reduces poverty by increasing income over time, while microinsurance effectively buffers income in the wake of a drought. However, other economic indicators may be relevant for households that are not as vulnerable to poverty (i.e., land-rich in our analysis). In particular, risk-averse households may be interested in reducing income variability in addition to increasing mean income. Hence, microinsurance may provide benefit to these types of household that our resilience analysis does not identify.

To formalize this benefit, we calculated an expected risk-averse utility on income over time under each scenario. We used an exponential utility function of the form $1 - exp(-X/R)$, where X represents income and R represents the household's risk tolerance. Figure D.7 shows that the utility of more risk-averse households (i.e., with lower risk tolerance) is more strongly benefited by insurance than cover cropping. Due to the delay in cover cropping's benefits on income, cover cropping leads to a short-term reduction in utility, which after 20-50 years increases to eventually exceed that of microinsurance. At lower levels of risk aversion (i.e., higher risk tolerance), the shape of the utility effects more closely mirrors that of expected income (Figure D.7). Hence, by reducing income variability (specifically, the downside income risk), microinsurance may be a more promising strategy for risk-averse households that are not in poverty or whose crop yields are not highly nutrient limited.

When both strategies are implemented together, the long-run utility exceeds that of both strategies in isolation, demonstrating a complementary effect on utility. However, due to the short-term financial tradeoffs associated with cover cropping, the shorter-term utility of both options together is lower than with microinsurance. Nevertheless, particularly for a risk-averse household, at no point does the combined utility decrease below the baseline condition. This demonstrates that, from a utility perspective, the welfare impacts of the short-term losses associated with cover cropping may be offset by the risk reduction offered by microinsurance.

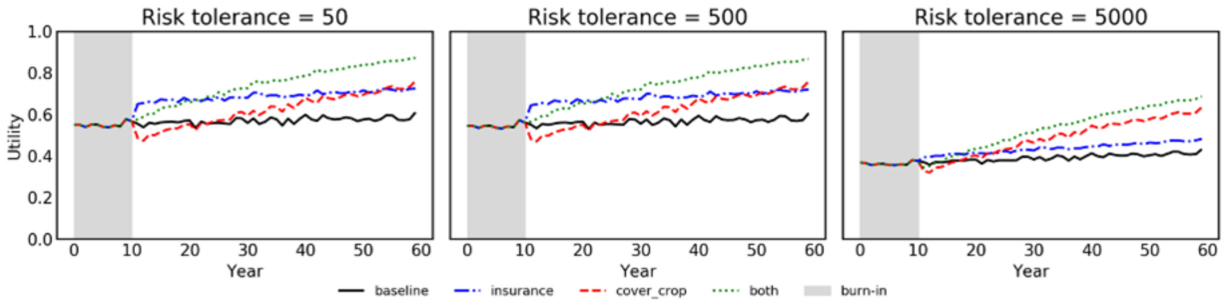

Figure D.7: Expected utility over time for a land-rich household under three levels of risk tolerance. Higher risk tolerance corresponds to lower risk aversion.

D.3 Synergies

The analysis in the main body of the article reveals a story of complementarity between microinsurance and cover cropping. Here, we examine whether the strategies, when implemented together, lead to *synergistic* effects. We conceptualize a synergy as a situation in which "the sum is greater than the parts". In this case, this represents:

$$Benefit_{both} > Benefit_{CC} + Benefit_{Ins} \tag{D.1}$$

where the "Benefit" is measured in the same way as shock absorption (Equation 5.3).

The results (Figure D.8) reveal that the modeled strategies exhibit synergies with respect to shock absorption in the long-term. In the short term, however, the combined effect is less than the sum of its parts. This is mainly explained by cover cropping's short-term detriment to shock absorption while soil organic matter (SOM) is being built. The long-term synergy is not surprising, given the structure of the model; each strategy operates through distinct mechanisms: cover cropping through the building of SOM and microinsurance through income stabilization. Each of these mechanisms enables the wealth-SOM feedback loop, consequently contributing to higher income. Due to this feedback, the combined effect of the strategies is heightened, and therefore synergistic.

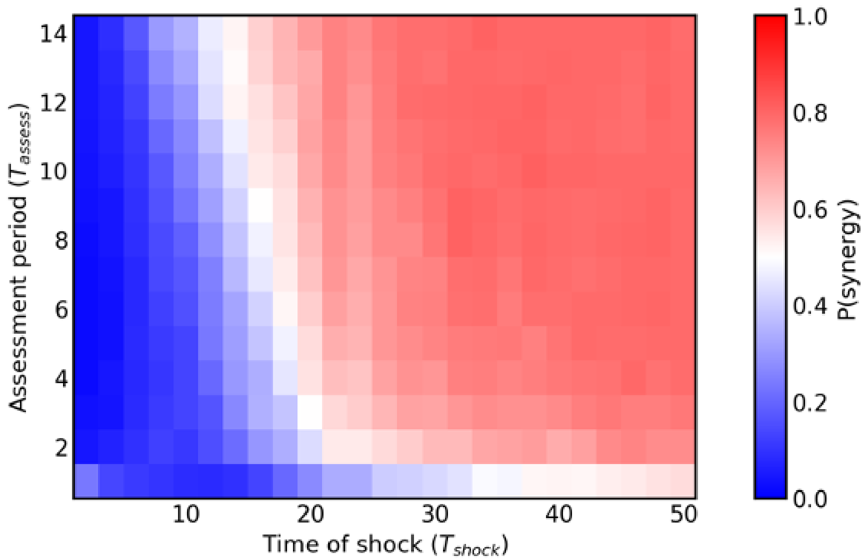

Figure D.8: Probability that both strategies together provide larger benefits than the sum of both strategies separately. This represents the outcomes for a "middle" household and a 0.2 magnitude drought.

D.4 Effect of microinsurance climate index

The microinsurance scheme is structured such that the insurance is "fair". For instance, if the insurance provides payouts in 5% of the years, the annual cost is 1/20$^{\text{th}}$ of the payout. Similarly, if the insurance provides payouts in 20% of the years, the annual cost is 1/5$^{\text{th}}$ of the payout. Thus, an insurance scheme with more frequent payouts entails higher premium costs. As a result, an insurance scheme that provides more regular payouts provides a lower *net* benefit to the household in a year in which the insurance is triggered. (Note that the strike rate affects the rainfall value at which the insurance is triggered.)

This characteristic results in a tradeoff in our model with respect to the microinsurance climate index (Figure D.9). Here, "climate condition" represents the annual realization of climate. The probability of a given climate condition occurring is influenced by the climate distribution (i.e., climatic context; here $\sim N(0.5, 0.2)$), but the outcomes in Figure D.9 under a given climate condition depend only on the climate condition itself.

For example, under the most extreme plotted climate condition (0.05), an insurance payout is received for all insurance indexes (strike rates). This payout is the same for all insurance indexes (5% insured, 10% insured, etc.). However, the cost of the premium is highest in the 30% insured case (i.e., 30% of the payout). This high premium means that, despite the payout being received, the household receives a lower net benefit in this year. As a result, the probability with which it must sell livestock is higher (0.55) than under an insurance scheme that provides less regular payouts (e.g., 0.10 probability under the 5% insurance index).

However, the higher insurance indexes (e.g., 30% insured) also provide payouts under less extreme drought conditions. For example, when the climate condition is 0.4, a payout is received under the 30% insurance index but not under any of the other assessed indexes. As a result, the probability with which livestock selling is required is lowest for the 30% insurance index under this climate condition.

Together, this represents a tradeoff in which insurance that provides more regular payouts offers protection under moderate climate conditions at the expense of vulnerability under more severe climate conditions, whereas insurance that provides less regular payouts protects against the severe climate conditions at the expense of vulnerability under more moderate conditions. Depending on the distribution of the climate condition (here, $\sim N(0.5, 0.2)$ truncated at 0 and 1), the net effect of this tradeoff will change as the probability of more and less extreme climate conditions shifts. In addition, farmer-level risk preferences may influence the aversion to different kinds of loss. Thus, the robust design of index-based microinsurance schemes in case study applications should consider the potential for this type of tradeoff.

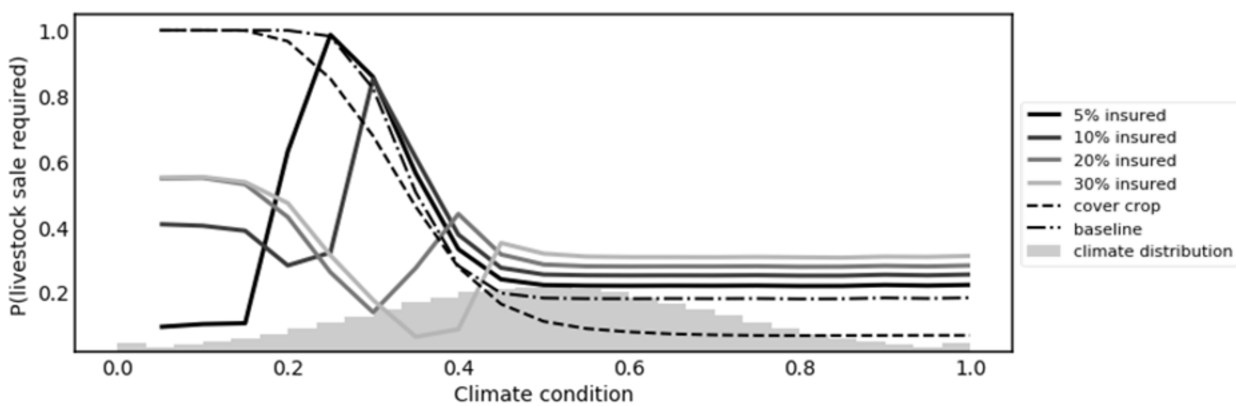

Figure D.9: The probability that livestock selling as a coping measure is required as a function of the annual climate condition in a simulation under regular climate variability ($\sim N(\mu = 0.5, \sigma = 0.2)$) and different insurance coverages. For example, a point (0.4,0.7) represents a case in which during a year with a climate condition at 0.4 (affecting crop production) there is a 70% chance that the household's annual income is insufficient to satisfy their consumption and they must sell livestock resources. 5% insured represents an index-based insurance in which a payment is received in 5% of years. This is for a land-poor household only.

D.5 Sensitivity analysis methodology and additional results

D.5.1 Parameter sampling

We conducted a global sensitivity analysis on the majority of the parameters of the model (see Table D.1 in the ODD+D description for the selected parameters). To generate perturbed parameter sets we employed the following procedure:

1. Generate a random deviation a_i for each of the P parameters ($\boldsymbol{a} = a_1, ..., a_P$), allowing the deviation to be 30% upwards or downwards: $\boldsymbol{a} \sim U(0.7, 1.3)^P$

2. Perturb each parameter from its baseline value $X_i(\boldsymbol{X} = X_1, ..., X_P)$ by this simulated value, giving a perturbed parameter set: $\boldsymbol{S'_r} = \boldsymbol{aX}$

3. Repeat this procedure 10,000 times, giving $\boldsymbol{S'} = \boldsymbol{S'_1}, ..., \boldsymbol{S'_{10000}}$. Here, we used latin hypercube sampling to increase the efficiency of the sampling of the parameter space.

D.5.2 Model evaluation

For each set of perturbed parameters $\boldsymbol{S'_r}$ calculate the Quantity of Interest (QoI), where the QoI takes two forms:

1. QoI_{shock} represents $P(CC \succ Ins)^{shock}$ in Experiment 1 (Table 5.1) with $T_{assess} = 5$ and $T_{shock} = 10$ and a 10% shock.

2. QoI_{pov} represents $P(CC \succ Ins)^{pov}$ in Experiment 2 (Table 5.1) with $T_{pov} = 50$.

The model evaluation procedure results in a "dataset" of sorts, where the independent variables are the parameters ($\boldsymbol{S'}$, with P columns and 10,000 rows) and the dependent variable is the quantity of interest (QoI_{pov} or QoI_{shock} of size 10,000).

D.5.3 Gradient-boosted regression forest

The goal of the sensitivity analysis is to assess how changes in the parameters affect the QoI. Hence, we are interested in exploring the function f in the relationship $QoI = f(\boldsymbol{S'})$. This function may be non-linear. We trained a gradient-boosted regression forest (GBRF) to yield a non-parametric representation of f. A GBRF consists of a set of simple regression trees that are fit in a stagewise manner, with each successive tree being fit to the residuals of the previous. GBRFs originated in the machine learning community, and generally exhibit a high predictive performance (Elith et al., 2008). We do not discuss this method in detail here and refer interested readers to Elith et al. (2008).

D.5.4 Assessing variable influence

We use partial dependence plots (PDPs)-—a common visualization technique for non-parametric models—to visualize the associations between changes in each parameter and the QoI, as assessed by the GBRF. Each point (x, y) on a partial dependence plot for parameter p_i represents the average prediction made by the GBRF (y value) if every instance of p_i is set to x, keeping all other parameters (p_{-i}) at their original values. The slope of the PDP gives an indication of both the magnitude and direction of influence of the parameter on the QoI. A PDP for a linear regression model would show a straight line representing the regression coefficient (β). To generate confidence bounds on our PDPs we bootstrap the "dataset" 100 times, each time re-training the GBRF and re-estimating the PDP. We also report a measure of variable importance using the scikit-learn package in Python (Pedregosa et al., 2011) (Figure D.10).

D.5.5 Supplemental results

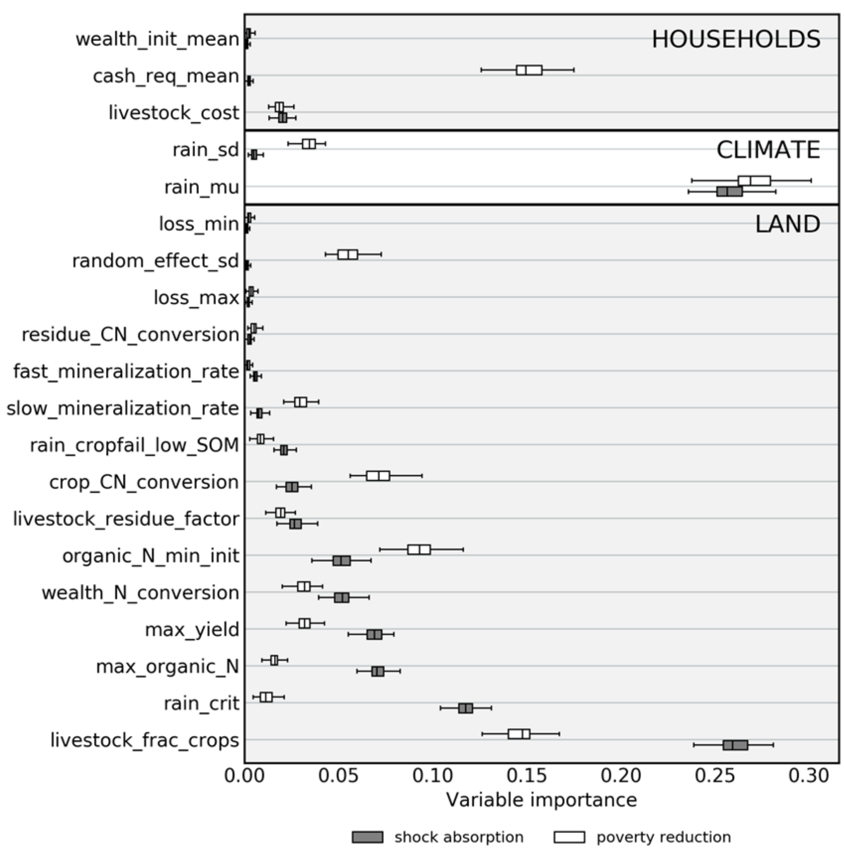

Figure D.10: Importance of different model parameters in the sensitivity analysis, as calculated by the GBRF. The "variable importance" measure is calculated by scikit-learn in Python and is a measure of the amount of variance that each variable explains.

D.6 Convergence analysis

The goal of the convergence analysis was to estimate how many replications of the model are required to generate model outputs that are not significantly influenced by stochasticity within the model. We refer to this number of replications as r^*. In our case, the quantity of interest is $P(CC \succ Ins)^{shock}$. We expect that this probability will vary considerably with both T_{shock} and T_{assess}. Hence, we choose $r^* = max(r^*_{T_{shock},T_{assess}})$, $\forall(T_{shock}, T_{assess})$ over $T_{shock} \in \{5, 10, 20\}$ and $T_{assess} \in \{1, 3, 5, 7, 9, 11, 13\}$.

Our approach for estimating each $r^*_{T_{shock},T_{assess}}$ was as follows:

1. Run a large number of model replications (1000).

2. Assume the estimated $P(CC \succ Ins)^{shock}$ over these replications ($\hat{X}_{1000}$) is the "true" value.

3. For each $r \in \{1, ..., 1000\}$, calculate the absolute error (AE) from the true value. For example, $AE_{50} = |\bar{X}_{1000} - \bar{X}_{50}|$, where $\bar{X}_{50}$ represents $P(CC \succ Ins)^{shock}$ calculated over the first 50 replications.

4. Choose r^* as the number of replications at which the absolute error in the estimated probability falls below 5%, i.e., $r^* = argmax_n(AE_n > 0.05)$.

The threshold of 5% was chosen as we do not require highly precise estimates of $P(CC \succ Ins)^{shock}$ for our assessment. We acknowledge that our approach is relatively ad-hoc and not formally statistically grounded. However, it captures the essence of what we desire: estimates of $P(CC \succ Ins)^{shock}$ that are robust to within-model stochasticity. We considered using the approach presented in Abreu and Ralha (2018), but the coefficient of variation (i.e., the standard deviation of $P(CC \succ Ins)^{shock}$ divided by the mean) is unstable with estimates near zero. Additionally, we considered the approach presented in Law (2008) (pg. 502), but because our model is not computationally intensive it was feasible to run a large number of simulations and calculate $\bar{X}_n \forall n$ and we adopted the approach described above.

The results indicate that $r^* = 188$ is sufficient (Figure D.11). To be conservative, we run the model at least 300 times for all experiments. For some figures we used a higher number of replications to improve visual clarity.

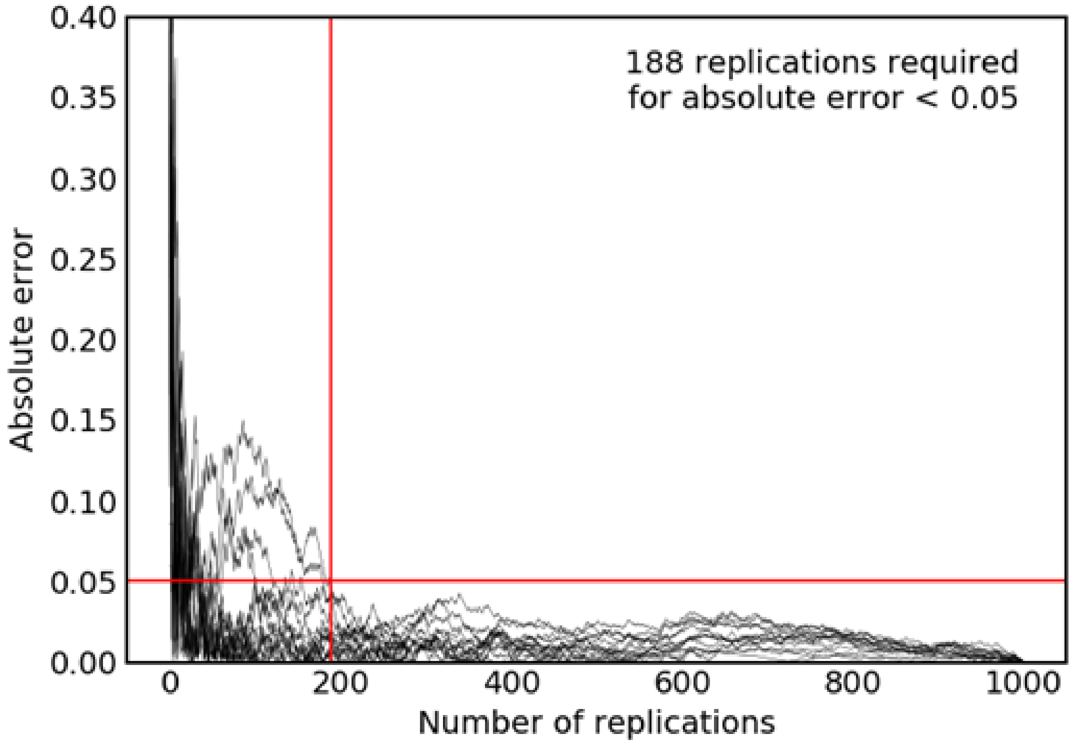

Figure D.11: Absolute error in the estimate of $P(CC \succ Ins)^{shock}$ as the number of model replications is increased. Each black line represents a unique (T_{shock}, T_{assess}). The red lines show the point at which the absolute error falls below 0.05 for all (T_{shock}, T_{assess}).

D.7 ODD+D model description

D.7.1 Overview

D.7.1.1 Purpose

The model was developed to investigate the short- and long-term resilience of a smallholder agricultural farming system and the effects of different household-level adaptation strategies on this resilience. It is intended to be used by researchers interested in exploring long-term dynamics of agricultural adaptation options. The model represents a mixed crop-livestock agricultural system, designed to be generally representative of a smallholder agricultural system in the Global South. Given the interest in exploring the general mechanisms through which different adaptation options affect resilience, the model is intentionally stylized and does not draw from empirical data to be representative of a specific location.

D.7.1.2 Entities, state variables, and scales

The model represents smallholder households that engage in agriculture and carry their wealth in the form of livestock. Each household is defined by a static land holding and has dynamic income and livestock holdings. Livestock are grazed on a combination of on-farm crop residues and an external rangeland, which is not explicitly modeled. The household's land has an evolving level of organic nutrients, which represent SOM and soil organic N together in a stylized manner. The model is spatially implicit, no environmental feedbacks beyond the household scale are represented, and households do not interact with each other.

D.7.1.3 Process overview and scheduling

The model operates at an annual time scale. Each year of the simulation involves calculation of: (1) soil nutrient flows; (2) crop yields; (3) household income; and (4) household wealth and coping measures (Figure D.12).

D.7.2 Design concepts

D.7.2.1 Theoretical and empirical background

The model represents soil nutrient dynamics in a stylized way. It models slow-evolving stocks of SOM and faster-acting pools of mineralized nutrients. Our representation is consistent with soil representations in biogeochemical models (Manzoni and Porporato, 2009) and is qualitatively comparable to other more complicated process-based models of soil nutrient dynamics used for

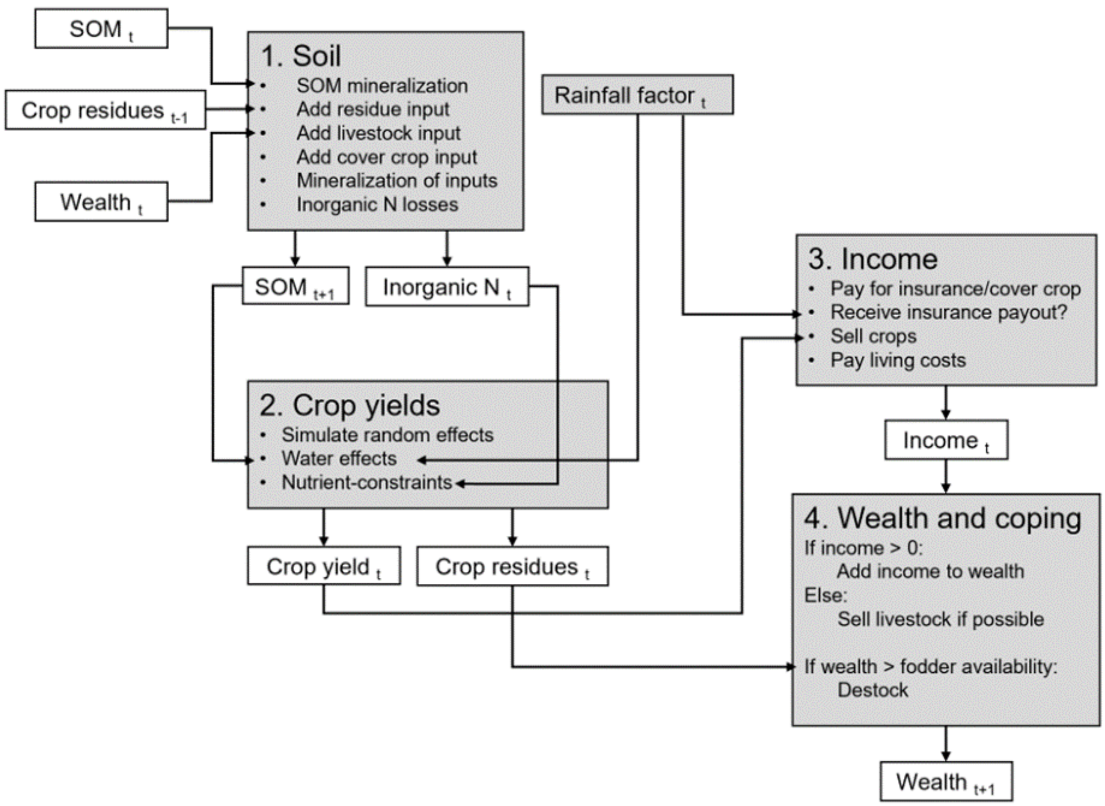

Figure D.12: Overview of annual simulation process.

agricultural applications (e.g., CENTURY (Metherell, 1993), DSSAT (Jones et al., 2003), and APSIM (Keating et al., 2003)).

Our crop yield model assumes that yields are influenced jointly by climate and nutrient availability. This representation generally follows Liebig's law of the minimum, which assumes that yields are influenced solely by the most constraining of these factors and plateau when each factor is above some threshold (Tittonell and Giller, 2013; Ferreira et al., 2017) (i.e., the crop can be water- or nutrient-limited). Similar representations are used in other more complicated process-based models of crop yield (e.g., CENTURY (Metherell, 1993), STICS (Brisson et al., 2003)) and in other simulation models (Grillot et al., 2018).

Together, our soil nutrient and crop yield representations exhibit the following qualitative characteristics:

1. Consistent cropping without replenishment of organic matter will slowly degrade soil quality and hence crop yields over time (Giller et al., 1997; Reeves, 1997; Bennett et al., 2012);

2. Soil quality can be maintained and built through organic inputs (e.g., manure or leguminous cover crops) (Giller et al., 1997; Drinkwater et al., 1998; Wittwer et al., 2017); and

3. Soil organic matter has benefits for drought sensitivity and nutrient losses (Drinkwater et al., 1998; Bommarco et al., 2013).

Household decision-making represents wealth accumulation and coping measures, and is modeled using a simple heuristic. This heuristic assumes that: (1) households store their wealth in the form of livestock and do not have cash savings; (2) livestock are sold if necessary to meet immediate cash needs (Bellemare and Barrett, 2006; Moyo and Swanepoel, 2010); and (3) total herd size is limited by feed availability (Valbuena et al., 2012; Assefa et al., 2013).

D.7.2.2 Individual decision making

The household makes two decisions related to their livestock wealth reserves, both of which are governed by simple heuristics. First, if the household's income in a given year is negative, they make up the deficit by drawing from their wealth reserves (a proxy for the selling of livestock). If wealth reserves are insufficient to make up the deficit, we assume that the household reduces their consumption. Both livestock selling and consumption reduction are considered as coping mechanisms. If, instead, their income is positive, they add this surplus to their wealth reserves (a proxy for the buying of livestock). This latter case is mediated by the second heuristic; if a household's livestock herd (i.e., wealth reserves) is larger than could be fed by their crop residues (assuming some percentage of their herd is grazed on common pastures), they are forced to destock these animals that cannot be fed. Given that wealth can only be held in the form of livestock-— i.e., we do not model financial resources—the household receives no monetary benefit for this destocking.

These heuristics are not influenced by any other factors and there are no notions of beliefs, memory, learning, adaptation, or social or cultural norms.

D.7.2.3 Learning

There is no notion of learning in the household's decision-making.

D.7.2.4 Individual sensing

Each year, the household observes its crop yields, residue production, and income, which influence the decision heuristics.

D.7.2.5 Individual prediction

The household does not predict future conditions.

D.7.2.6 Interaction

There are no interactions between households. Livestock are assumed to be partially grazed on common rangeland, which implies interactions with other households, but we do not explicitly model the rangeland dynamics, so this interaction is not endogenous to the model.

D.7.2.7 Collectives

The household does not form collectives.

D.7.2.8 Heterogeneity

The household is defined by its initial wealth reserves, initial soil quality, and land holdings. In our simulations, we consider only the implications of different levels of land holdings. Given that there are no interactions in our model, running the simulation for three households with heterogeneous land endowments is equivalent to running it three times separately with a single household.

D.7.2.9 Stochasticity

There are two sources of stochasticity in the model: (1) the generation of yearly climate conditions, which is constant across all households; and (2) a household-level random effect in the calculation of crop yields. The household-level effect conceptually represents other non-modeled factors that may influence crop yields, household-level (positive or negative) shocks, and household-level variability in the experience of the regional climate condition. Together, this requires us to simulate a set of hypothetical climate time series and, for each time series, run the model for a set of households that experience different random crop yield effects. Under the baseline model settings, the variability of the household-level effect is approximately half that of the region-level effect. The model therefore allows for considerable path dependencies introduced by household-level stochasticity.

D.7.2.10 Observation

Model outputs include yields, income, wealth, soil organic matter, and mineralized nutrients. These are observed at the household level at an annual basis.

D.7.2.11 Emergence

There exists a positive feedback loop, in which positive income enables accumulation of livestock (wealth reserves), providing additional soil organic matter, which in turn increases future crop yields and income. The ability for the household to experience this positive feedback cycle is mediated by their land endowment, initial soil organic matter, climate, and random yield effects. As such, household "trajectories" emerge as a combination of these random and non-random factors. Given the importance of stochasticity, there exists a considerable degree of path dependence in the model; a household that is *unlucky* one year (i.e., has a large, negative random effect in their crop yields) may be pushed into a downward spiral of decreasing livestock herds, soil organic matter, crop yields, and income. We investigate the possibility for household adaptation options (cover cropping and insurance) to influence these trajectories and hence contribute to different emergent outcomes.

D.7.3 Details

D.7.3.1 Implementation details

The model is implemented in Python 3.6. Code is available online at CoMSES.net.[1]

[1] https://www.comses.net/codebases/ee47544a-7eb0-4482-8967-42d6b0c05060/releases/1.0.0/

D.7.3.2 Initialization

The model is stylized and does not draw from any extensive empirical datasets. To initialize a single simulation, the climate time series is first generated, followed by a population of households with heterogeneous land endowments. Household initial wealth and soil organic matter levels are homogeneous and are specified by exogenous parameters (see section D.7.3.3). As stated above, a single model with multiple households is functionally no different to multiple models with a single household, but we do it in this way both for computational efficiency (through vectorization of calculations) and simpler management of random number seeds. Within an experiment, the random number seed is the only factor that is varied upon initialization.

D.7.3.3 Input data and parameterization

Model parameterization is achieved through a combination of information from literature and a pattern-oriented modeling calibration process. All model parameters are displayed in Table D.1. The calibration process is described in section D.7.3.5. Although we do not intend the model to be representative of any specific region or location, we chose to draw several of the parameters from Ethiopian data sources. Ethiopia's population is primarily engaged in smallholder agriculture—many in mixed crop-livestock systems—and thus Ethiopia serves as a relevant setting from which to draw stylized information. This enabled us to represent the relative scales of different model elements (e.g., maximum crop yields and crop selling prices) without requiring these values to be determined by the calibration process, thus reducing the dimensionality of the uncertain parameter set.

Additionally, although our representation of soil nutrient dynamics is stylized and we do not claim to realistically represent actual nutrient flows, we measure the SOM pool in units of kilograms of nitrogen per hectare (kg N/ha). This again allowed us to ground several parameters in empirically observed values (e.g., nitrogen-fixation of cover crops), reducing the number of uncertain parameters. However, we note that some values, particularly the C:N ratios, remain unrealistic in this model parameterization.

The derivation of several parameters requires some explanation:

- **Initial and maximum SOM:** In reality, baseline amounts of organic matter in a non-degraded soil are sufficient to provide nutrients for moderate levels of crop yield. To parameterize the initial SOM, we used information from other parameters to give a rough estimate of a reasonable value. Specifically, we assumed that the soil itself would initially be able to provide 4,000 kg/ha crop yield (approximately 2/3 of the maximum yield) in the absence of other inputs. Using the C:N ratio in the crop (50), this is equivalent to 80 kg N/ha of mineralized inorganic N that is produced solely through mineralization from SOM. With a mineralization rate of 0.02, this requires an initial SOM level of 4,000 kg/ha. We then chose the maximum SOM level to be double the initial SOM level.

- **Wealth to nitrogen conversion:** Using values from Newcombe (1987), we calculated that a cattle might produce 6,165 kg of fresh dung or, equivalently, 5,364 kg of dry matter per year. Assuming that 1.46% of the dry weight is nitrogen (also comparable to Lupwayi et al. (2000)), this equates to 78.3 kg N/cattle/year. Assuming a price of 3,000 birr (the Ethiopian currency) for a single animal, this is equivalent to 0.026 kg N/year/birr.

- **Land endowment:** In reality, smallholder land holdings vary by a larger degree than we represent in the model. However, we assume that each household – regardless of their land endowment and wealth – has the same annual living costs. In reality, land-rich households might have more household members, and consumption also generally increases with wealth. For simplicity in the analysis, our households vary over a single dimension (land endowment), so we do not incorporate such secondary effects and hence parameterize the variability in land endowment from only 1 to 2 ha. These values respectively correspond to the 47$^{\text{th}}$ and 75$^{\text{th}}$ quantiles of household landholdings in the Ethiopia 2015 LSMS data.

Table D.1: Parameter values and sources.

Parameter	**Symbol**	**Value**	**Unit**	**Source**	**Uncertain ¶**	**Sensitivity analysis**	**Description**
Simulation settings							
Number of households	N_A	200	-				Varied over simulation runs.
Random seed	s	0	-				
Households							
Land endowment	L	$\{1, 1.5, 2\}$	ha			x	Varied over households. See text in section D.7.3.3.
Initial wealth	W_0	36,165	birr		x	x	Proxy for livestock.
Cash requirement	CR	6,001	birr		x	x	Annual cash requirement for consumption.
Market							
Crop sell price	P_{crop}	2.17	birr/kg	FAO†			Mean 2015 price for maize in Addis Ababa.
Livestock price	P_{ls}	3,000	birr/head	CSA‡			Averae 2015 price.
Yields							

Parameter	Symbol	Value	Unit	Source	Unce-rtain ¶	Sens. anal-ysis	Description
Crop C:N	CN_{crop}	50	gC/gN	(Metherell, 1993)		x	Carbon to nitrogen ratio in harvested crop. Value loosely taken from the CENTURY model description.
Residue C:N	$CN_{residue}$	196	gC/gN		x	x	Carbon to nitrogen ratio in crop residue. In (Elias et al., 1998) this is approximately four times the ratio of the harvested crop.
Maximum yield	Y_{max}	6,590	kg/ha	LSMS§		x	95th percentile maize yield over Ethiopia in 2011, 2013, and 2015.
Climate upper threshold	C^{upper}	0.8	-	(Metherell, 1993)		x	Climate condition above which crop yields plateau
Climate lower threshold (low SOM)	C^{lower}_{low}	0.3	-		x	x	Climate condition below which crop failure occurs when SOM is zero.
Climate lower threshold (high SOM)	C^{lower}_{high}	0	-				Climate condition below which crop failure occurs when SOM is at its maximum.
Crop yield random effect	σ_y	0.3	-			x	Standard deviation of the crop yield random effect, simulated as $\sim N(1, 0.3)$
Residue loss factor	$l_{residue}$	10	%	(Assefa et al., 2013)			Percentage of crop residues not returned to the soil or fed to livestock
Residue multiplier	$mult$	2	-	(Bogale et al., 2008; Assefa et al., 2013)			Residue production per unit of harvested crop.
Soil							
SOM mineralization rate	k_{slow}	2	%/year	(Schmidt et al., 2011)		x	50-year turnover time of bulk SOM
Applied organic matter mineralization rate	k_{fast}	10	%/year		x	x	The percentage of applied organic matter (manure and/or crop residues) that mineralizes in the year of application.

Parameter	Symbol	Value	Unit	Source	Uncertain ¶	Sens. analysis	Description
Initial SOM	SOM_0	4,000	kg N/ha	-		x	See text in section D.7.3.3.
Maximum SOM	SOM_{max}	8,000	kg N/ha	-		x	See text in section D.7.3.3.
Maximum leaching rate	l_N^{max}	25	%	(Giller et al., 1997; Di and Cameron, 2002)	x	x	Rate of leaching of mineralized organic matter when SOM is zero.
Minimum leaching rate	l_N^{min}	5	%	(Di and Cameron, 2002)		x	Rate of leaching of mineralized organic matter when SOM is at its maximum.
Livestock							
Wealth:nitrogen ratio	SN_{conv}	0.0018	kg N/year/birr	-	x	x	0.026 kgN/year/birr is the derived value for comparison (see text in section D.7.3.3)
Percent crop grazing	$c_{residues}$	52	%	(Keftasa, 1988; Bediye et al., 2001)	x	x	Percentage of livestock food requirements that come from crop residues. The remainder comes from a non-modeled external rangeland.
Consumption requirement	cf	2,280	kg DM/TLU /year |	(Amsalu and Addisu, 2014)		x	We assume all residues are dry matter
Climate							
Mean	μ_c	0.5	-			x	
Standard deviation	σ_c	0.2	-			x	
Climate							
Climate percentile	Ins_{perc}	10	%				Climate threshold (percentile of cumulative distribution function) below which an insurance payout is received.

Parameter	Symbol	Value	Unit	Source	Uncertain ¶	Sens. analysis	Description
Payout magnitude	Ins_{payout}	1	-				Insurance payout relative to the expected yield. For example, if this is 1, the insurance payout will equal the income from an average year's yields (assuming no nutrient limitations on crop growth).
Cost factor	Ins_{cost}	1	-				Fairness of insurance. A value of 1 indicates an actuarially fair policy, where the annual cost is equivalent to the expected annual benefit.
Cover cropping							
Nitrogen fixation	CC_{N_fix}	95	kg N/ha	(Büchi et al., 2015; Wittwer et al., 2017; Couëdel et al., 2018)			Maximum value with no water limitation.
Cost factor	CC_{cost}	1	-				Annual cost of cover cropping relative to the cost of insurance.

¶The values displayed for the uncertain parameters were calibrated using the pattern-oriented modeling process (section D.7.3.5)

†http://www.fao.org/giews/food-prices/tool/public/

‡CSA = Ethiopian Central Statistical Agency. Source = annual retail price sheets.

§LSMS = Living Standards Measurement Study

| DM = dry matter, TLU = tropical livestock unit

D.7.3.4 Sub-models

Soil nutrients The model contains two main pools of soil nutrients: organic and mineralized. The states of these pools are measured in kg N/ha. Each year, a portion of the organic pool of nutrients (SOM) mineralizes according to a linear decay process. Organic nutrients applied to the soil (manure and crop residues; N_{added}) also are partially mineralized in the year of application (with a linear rate constant larger than that of the SOM), with the non-mineralized component added to the bulk SOM. We do not differentiate between the addition of "organic matter" and "nitrogen" and use a single variable to retain simplicity.

$$N_{mineralized}^{SOM} = k_{slow} * SOM_t \tag{D.2}$$

$$N_{mineralized}^{added} = k_{fast} * N_{added} \tag{D.3}$$

$$N_{mineralized}^{total} = N_{mineralized}^{SOM} + N_{mineralized}^{added} \tag{D.4}$$

$$SOM_{t+1} = (SOM_t - N_{mineralized}^{SOM}) + (N_{added} - N_{mineralized}^{added}) \tag{D.5}$$

After mineralization, a percentage of the mineralized nutrients is leached from the system. Higher levels of SOM contribute to lower leaching rates (Drinkwater et al., 1998). Specifically, we assume a maximum leaching rate with no SOM (l_N^{max}) and a minimum leaching rate when SOM is at its maximum (l_N^{min}), with a linear interpolation between these two points (see Table D.1 for parameter values).

Mineral N that remains after leaching is assumed to be fully available to the crop. If this is higher than the crop's N requirements, any excess mineral N is assumed to be lost from the system via leaching (i.e., the mineral nutrient pool is reset each year).

This nutrient balance is partial and we do not model soil erosion (Cobo et al., 2010), yet the loss pathways that we include represent the largest magnitude pathways in mixed cropping-livestock systems (Tittonell et al., 2006). However, in its stylization, our representation of soil nutrient dynamics contains a number of simplifying assumptions, namely: (1) no endogenous or dynamic representation of C:N ratios, (2) a single soil layer, (3) a single pool of organic nutrients with a single mineralization rate, (4) no explicit modeling of soil microbial biomass or other labile SOM pools, (5) no climate dependence in nutrient mineralization or leaching, (6) no nutrient dependence (e.g., N-limitations) in mineralization, (7) no differentiation between ammonium and nitrate as forms of inorganic N, and (8) no atmospheric losses of N through denitrification. Despite these assumptions, we believe that our representation provides a reasonable first-level approximation of more complicated soil dynamics and requires far less parameterization.

Climate Climate is represented through a single value, which is drawn each year from a normal distribution (parameters in Table D.1) that is bounded between 0 and 1. This value does not represent a specific physical climate characteristic (e.g., rainfall), but a stylized notion of the "climate

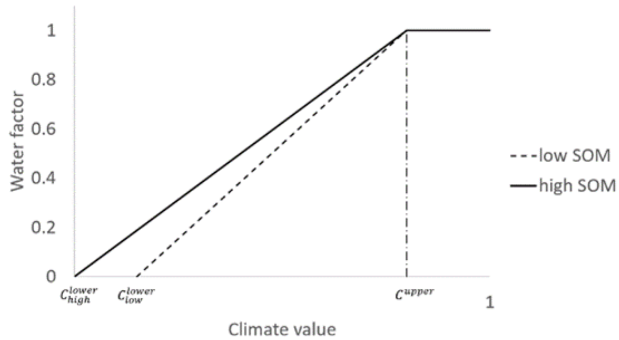

Figure D.13: Effect of climate on crop yields.

condition". Under baseline conditions, the simulated climate values interact with the model solely through crop yields. Under the insurance scenario, payouts are received in years in which the climate condition is below the insurance index value, which is defined as some percentile of the cumulative distribution of the climate condition (i.e., a 10% index represents the 10th percentile of the cumulative distribution). With cover cropping, the climate condition also affects cover crop nitrogen fixation. The climate value is qualitatively similar to the outputs of process-based methods that calculate ratios of actual evapotranspiration to potential evapotranspiration (e.g., applications of the FAO crop water requirements methodology (FAO, 1984; Block et al., 2008) and the CENTURY model (Metherell, 1993)), but requires far less parameterization.

Crop yields Crop yields can be reduced from a maximum potential value (Y_{max}) through water and/or nutrient limitations (Tittonell and Giller, 2013). First, we calculate a water factor, C_{water}, with $0 \leq C_{water} \leq 1$. It is assumed that (see Figure D.13): (1) if the climate value is greater than C^{upper} (0.8 in the parameterized model), then $C_{water} = 1$; (2) there is a critical climate value (≥ 0) at which $C_{water} = 0$; (3) higher levels of SOM lead to higher drought tolerance and hence a lower critical climate value; and (4) C_{water} scales linearly between the critical value and C^{upper}. The maximum water-constrained yield (Y_w) is then assumed to be:

$$Y_w = C_{water} * Y_{max} \tag{D.6}$$

Second, we determine the maximum attainable nutrient-constrained crop yield (Y_N) given the available mineral N in the soil ($N^{total}_{mineralized}$):

$$Y_N = \frac{N^{total}_{mineralized}}{\frac{1}{CN_{crop}} + \frac{mult}{CN_{residue}}} \tag{D.7}$$

This represents a partitioning of the $N^{total}_{mineralized}$ between the N in the harvested crop (adjusted by CN_{crop}) and the crop residues (adjusted by $CN_{residue}$ and multiplied by $mult$).

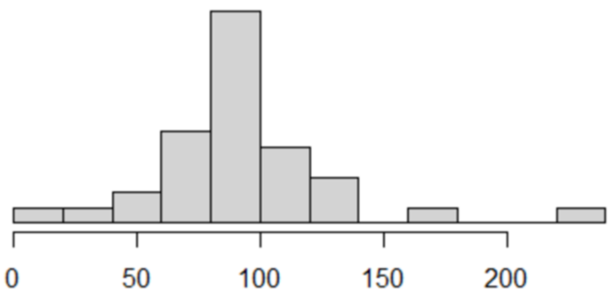

Figure D.14: Distribution of cover crop N fixation (kg N/ha) in temperate climates reported in Badgley et al. 2007. The median value is 95 kg N/ha.

The actual yield (Y^{obs}) is then calculated as:

$$Y^{obs} = min(Y_W, Y_N) * \epsilon \tag{D.8}$$

where $\epsilon \sim N(1, \sigma_y^2)$ is a household-level stochastic effect with σ_y given in Table D.1.

In this stylized crop yield model, we omit or simplify several processes that are included in more detailed process-based crop yield models, for example: (1) our one-dimensional representation of the effects of climate proxies any non-linearities in relationships between climate and yield as well as potential interactions between rainfall and temperature; (2) we do not model solar irradiation and growth of leaf area; and (3) we do not model the partitioning of growth between above- and below-ground biomass. Given the modular nature of our yield model, additional reduction factors could be added (e.g., see (Schreinemachers et al., 2007)) or more sophisticated process-based calculations could replace the existing calculations of water and nutrient limitations. However, this increased complication would require a greater amount of data and calibration, as well as reduce transparency in how specific inputs and structures mechanistically influence yields and the broader model dynamics.

Cover crop N_2 fixation As with vegetable crops, cover crops' biomass generation, and thereby their soil organic matter contributions, is also constrained by rainfall (Ewansiha and Singh, 2006). We assume that the N fixed by the cover crop follows the same water response function as vegetable crop yields (i.e., Figure D.13). Thus, in a year with no rainfall, no N is fixed. We set the default upper bound on N_2 fixation as 95 kg N/ha (Figure D.14; (Badgley et al., 2007)).

D.7.3.5 Pattern-oriented modeling (POM)

Description We use latin hypercube sampling to generate 100,000 potential parameter sets, where each parameter is drawn uniformly from the ranges in Table D.2. For each potential parameterization we run the model 10 times (to encompass climate variability) for a population of

Table D.2: Parameters included in the POM calibration.

	Parameter	Symbol	Minimum	Maximum	Notes
1	Households: initial wealth	W_0	5,000	50,000	
2	Households: annual cash requirement	CR	5,000	30,000	Median annual expenditure in 2015 LSMS is 17,261 birr.
3	Yields: climate lower threshold (low SOM)	C_{low}^{lower}	0	0.5	
4	Yields: residue C:N	$CN_{residue}$	25	200	Bounding the crop C:N ratio
5	Livestock: percent crop grazing	$c_{residues}$	0.5	1	Livestock are often grazed primarily on crop residue (Keftasa, 1988; Bediye et al., 2001).
6	Livestock: wealth:nitrogen conversion	WN_{conv}	0.01	0.05	Bounding the empirically-derived value (section D.7.3.3).
7	Soil: applied organic matter mineralization rate	k_{fast}	0.05	0.95	Must be faster than the SOM mineralization.
8	Soil: maximum leaching rate	l_N^{max}	0.05	0.95	

Table D.3: Patterns included in the POM calibration.

	Pattern	Requirements
1	Divergent household wealth trajectories	(a) All land-rich households finish the simulation with positive wealth AND (b) All land-poor households finish the simulation with no wealth AND (c) 20%-80% of the middle households finish the simulation with positive wealth.
2	Households can recover from shocks	There is at least one middle household that: (a) Has no wealth at some point during the simulation AND (b) Has positive wealth at the end of the simulation.
3	No saturation of SOM	There are no households consistently at the maximum level of SOM throughout the last 10 years of the simulation.
4	Some households can build SOM	At least 10% of households finish the simulation with a higher SOM than the initial value

100 households (to encompass variability induced by the random yield effect) for a period of 100 years. We choose only 10 model replications here due to computational reasons.

We assess whether each simulation generates a set of qualitative "patterns" (Table D.3). These patterns collectively represent desired model behavior under baseline simulation conditions. To evaluate a potential parameter set we: (1) measure which patterns are generated in each simulation, (2) calculate the probability that each pattern is generated over the 10 replications, and (3) sum these averages over all patterns.

Results Of the 100,000 parameter sets, three generated on average 3.2 of the four patterns (Figure D.15). We retained one of these parameterizations for the analysis presented in this paper. Experimentation with the other two parameterizations yielded qualitatively similar results that do not affect the conclusions drawn in this paper.

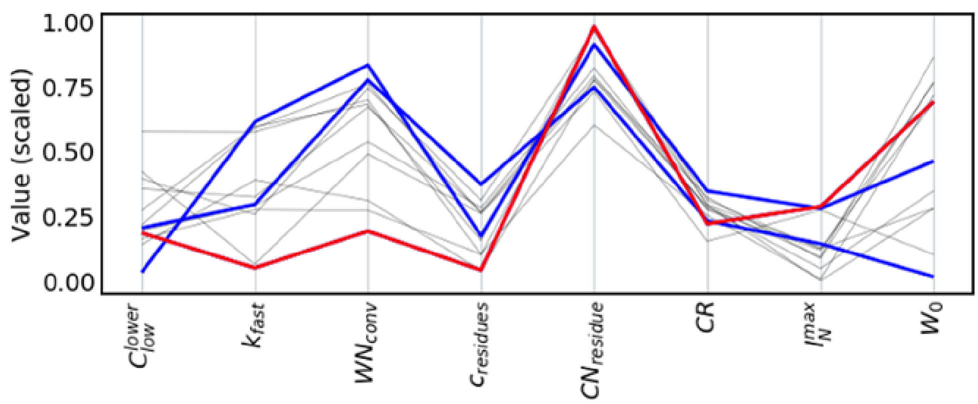

Figure D.15: Scaled parameter values of the resultant POM parameterizations. The red line represents the selected parameterization. Blue lines represent the other parameterizations that reproduced the same number of patterns. Grey lines show parameterizations that were within 20% of the best.

Appendix E

Supplement to Large-scale Land Acquisitions

E.1 Additional outputs

E.1.1 Illustrative simulation output

To illustrate the general dynamics of the calibrated ABM, we first present the model's outputs under a representative simulation (Figure E.1). Under the *Displacement* scenario, by decreasing the overall amount of land cultivated by the agents (Figure E.1 A1), the LSLA reduces the agents' crop production (Figure E.1 B1) and increases the amount of time spent in off-farm employment (Figure E.1 D1). Yet the increased off-farm employment is insufficient to curb the impacts of land loss on food security and livestock holdings (Figure E.1 E1 and F1). Conversely, the CF scenarios contribute to higher population-level food security (Figure E.1 F1), which progressively increases over time in the CFchoice scenario as more agents join the CF scheme (Figure E.1 A1).

The experience of any single agent is more dynamic than the regional outcomes. For example, a selected land-poor agent, which originally owned 0.75 ha of land and lost 0.25 ha to the LSLA, is food insecure in most years under baseline conditions and consistently food insecure with the LSLA (Figure E.1 F2). The CF arrangements, through increases in crop productivity, increase the agent's food security (Figure E.1 F2) and in the long run also increase its livestock holdings (Figure E.1 E2). At times, the agent allocates all of its land to the contract farming scheme in the CF_{choice} scenario (Figure E.1 A2), but this fluctuates between years due to the agent's evolving beliefs about the returns to contract farming. This agent uses non-farm employment as a coping mechanism and hence the time allocated to non-farm activities is higher in years and scenarios with lower crop production—when the extra income is needed most (Figure E.1 D2). A selected land-rich agent, in contrast, originally owns 1.75 ha of land and loses 0.75 ha to the LSLA (Figure E.1 A3). It is able to achieve higher levels of crop production and livestock holdings, and under baseline conditions it rarely allocates labor to non-farm employment and is rarely food insecure.

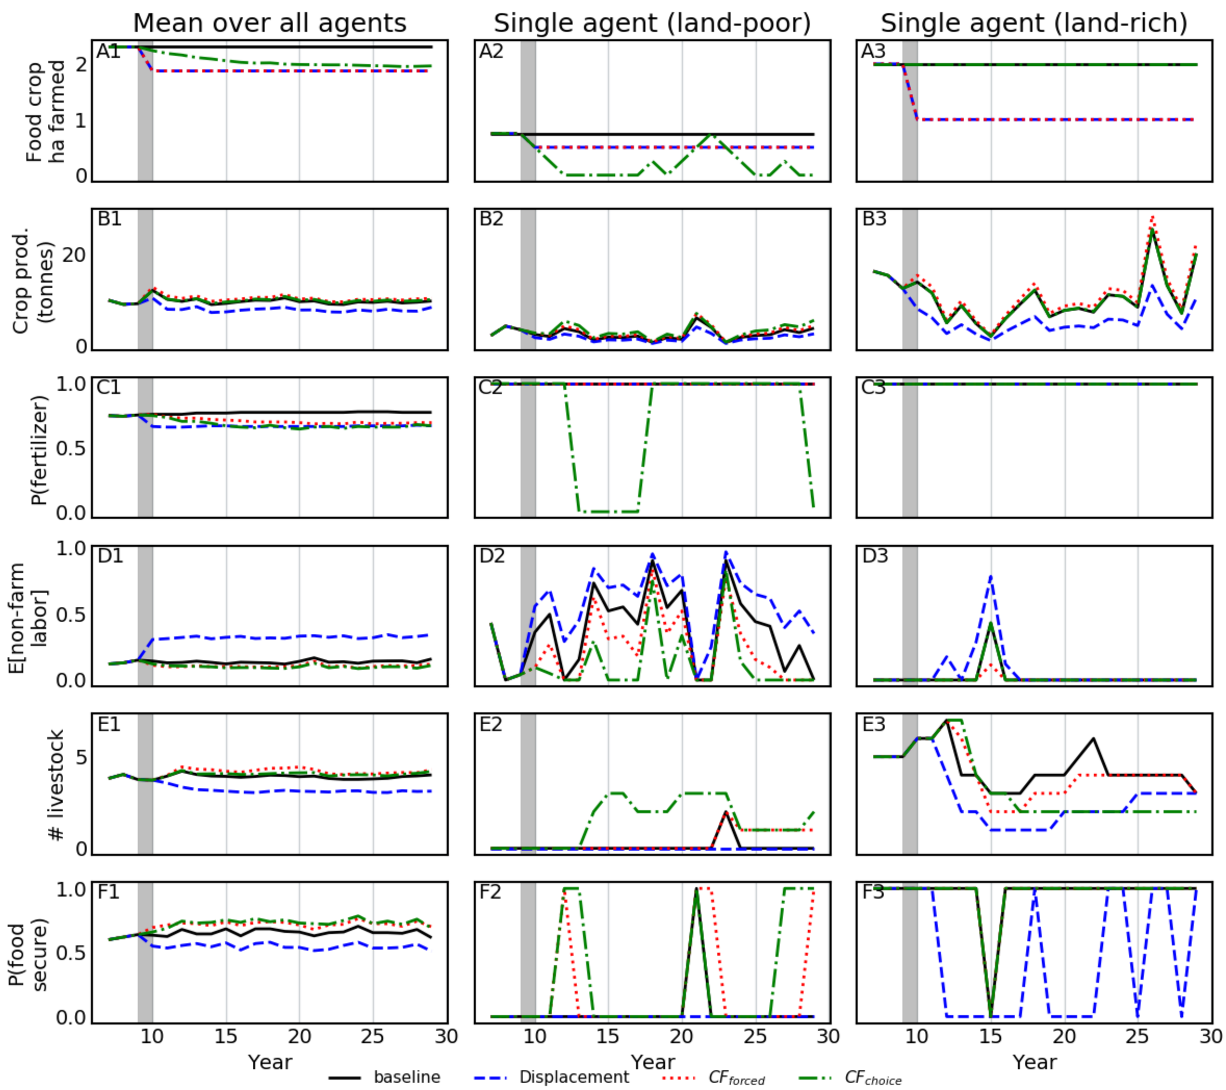

Figure E.1: Illustrative ABM output showing selected variables under baseline conditions and the three LSLA/contract farming (CF) scenarios. Vertical grey lines denote the LSLA/CF implementation year. The left-hand panel shows outputs averaged over the agent population. The middle and right-hand panels each show the outputs for a single agent. Note that crop production (B) only includes smallholder production and not that produced within the LSLA.

E.2 Supplemental results

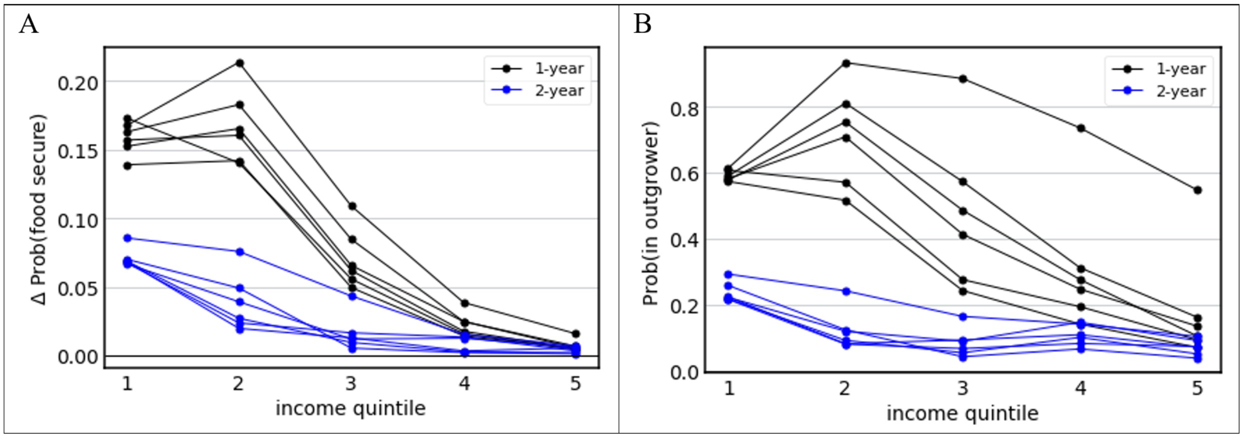

Figure E.2: Distribution of effects under the CF_{choice} scenario. (A) shows the effects on food security. (B) shows the probability of joining the contract farming scheme. Each line plots the average response from a different model calibration. Results are disaggregated by the cash crop harvest frequency.

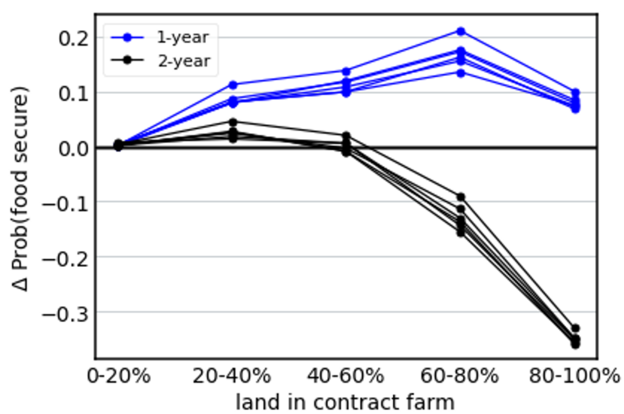

Figure E.3: Distribution of effects for households forced to participate in contract farming (CF_{forced}) based on the fraction of their overall land in the contract farming scheme. The outcomes are disaggregated by the cash crop's harvest frequency (1-year and 2-year). Each line represents the mean response under a different model calibration.

E.3 Convergence Analysis

We conducted a convergence analysis to calculate the required number of model replications such that our model outputs are not significantly biased by stochasticity within the ABM. To do this, we examined how the regional productivity and food security outcomes depend on the number of model replications, under a range of model conditions. Specifically, we ran a set of simulations under default settings for the four LSLA/CF scenarios (i.e., baseline, *Displacement*, CF_{forced}, and CF_{join}) for all four sites and all six model calibrations. This resulted in a total of 96 simulation conditions. For each condition, we ran a large number of replications (300) and assumed that the output values estimated over these replications approximate the asymptotic mean values, i.e., $\hat{Y}_{300} \approx \bar{Y}$. Then, we calculated the absolute relative error (ARE) for all replication sizes, r, less than 300:

$$ARE_r = \frac{\left|\hat{Y}_r - \bar{Y}\right|}{\bar{Y}} \tag{E.1}$$

Finally, we calculated the maximum value of r at which the ARE is larger than 0.05 (i.e., 5%) for either outcome. This yields a value of 30 (Figure E.4). We therefore run 30 replications for all experiments in the main body of the paper.

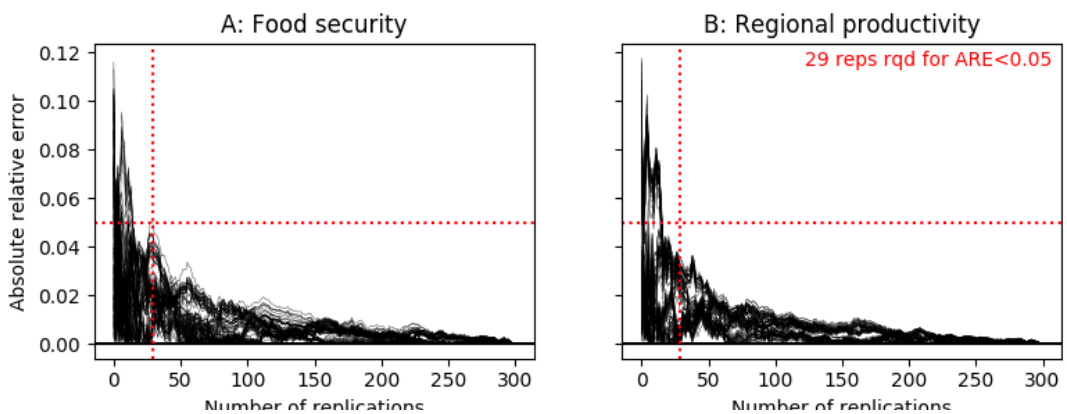

Figure E.4: Absolute relative error (ARE) as a function of the number of model replications. Each line plots the response under a different simulation condition (i.e., site-calibration-LSLA combination). 30 replications are required to achieve an ARE of less than 5% in both outcomes across all simulation conditions.

E.4 Model calibration

E.4.1 Procedure

We used a genetic algorithm to calibrate parameter values that could not be specified from available empirical data sources. There were 14 parameters included within the model calibration. These parameters, including their bounding ranges, are listed along with all other model parameters in Table E.3 and Table E.4 in the ODD+D description (section E.6).

The calibration procedure aimed to identify the set of parameter values from within the bounding ranges that leads to ABM outputs best matching a set of "patterns" derived from the household survey dataset. Although the household survey was cross-sectional, some questions asked households to recall information from before the LSLA, providing a proxy for pre-LSLA conditions. We used these pre-LSLA proxies to initialize the ABM, then the post-LSLA data (collected in 2019) to create the calibration patterns.

We took two approaches to reduce the chances of overfitting our model to the available data. First, we aimed to find a single set of region-level parameters that could generate the patterns across all four sites, thereby considerably reducing the dimensionality of the unknown parameter set (calibrating to each of the four sites separately would have required four times as many parameters). Second, in order to acknowledge the potential for equifinality (i.e., multiple plausible descriptions of the calibration data), we used the genetic algorithm to identify six sets of parameter values that similarly match the empirical data and are as different from each other as possible (Williams et al., 2020b). We retained all six parameterizations for the main simulation experiments.

We developed a set of patterns to span the range of modeled livelihood activities, including the distribution of these throughout the smallholder population (Table A 5 1). For each set of potential region-level model parameters, we ran the ABM for all four sites, using the site-level input data. For each pattern in each site, we calculated a measure of "loss" that represents the level of disagreement between the ABM outputs and the empirical data. For distributional patterns, the loss measured the mean squared difference in the height of the empirical and model-generated cumulative distribution functions:

$$Loss_{distrib} = \int_{-\infty}^{\infty} \Big(F_{model}(y) - F_{data}(y)\Big)^2 dy \tag{E.2}$$

where $F()$ is the cumulative distribution function, which we discretized into equal-sized bins. For probabilistic patterns, the loss measured the squared difference between the modeled and empirical values:

$$Loss_{prob} = (model - data)^2 \tag{E.3}$$

Both loss measures are bounded between zero and one. To calculate a total discrepancy for a region-level parameter set, we summed the loss over all (five) patterns and all (four) sites.

Table E.1: Fitting patterns

	Pattern	Description	Discrepancy measure
1	Livestock herd size (TLU)	Total livestock holdings.	Distributional
2	Non-farm labor (fraction of total)	Fraction of total productive household labor allocated to off-farm employment activities	Distributional
3	Crop yield (kg/ha)	Total reported crop production per cultivated area.	Distributional
4	Food security (probability)	Probability that a household experienced no difficulty in meeting household food needs throughout the past year.	Probability
5	Fertilizer use (probability)	Probability that a household applied inorganic fertilizer to their land in the past year.	Probability

E.4.2 Results

E.4.2.1 Comparison to empirical data

The resulting six model calibrations demonstrate relatively low levels of bias, with all loss values falling below 0.08 (Figure E.5). Similar patterns of bias exist across the sites and fitting patterns. The models have relatively low bias in the non-farm labor distribution, crop yield, and fertilizer use. There is a moderate bias for livestock holdings in OR4, as well as in the food security outcomes in sites OR1 and OR3. Examining the model-derived and empirical patterns (Figure E.6) suggests a tradeoff between the food security outcomes in these two sites in the model calibration: the model overpredicts food security in OR3 and underpredicts in OR1.

Calibrating separate parameters for each site would have led to a lower bias in these outcomes, but we consider the observed levels of bias to be acceptable for the intended application, which seeks to evaluate generalized impacts of LSLA configurations across all four sites in the region. We therefore retain the region-level calibration in order to reduce the model variance (i.e., overfitting to the empirical data).

E.4.2.2 Calibrated parameter values

Some calibrated parameter values are highly consistent between the six models (e.g., the non-food cost, risk tolerance, and mean soil fertility) (Figure E.7), suggesting that the model outcomes are sensitive to these parameters. Other parameters are present over their entire range in the calibrated models (for example, the three parameters pertaining to livestock). There are two potential explanations for this. First, it is possible that there is a degree of equifinality, in that there are multiple distinct explanations of the empirical data. Alternatively, it is possible that the model is not sen-

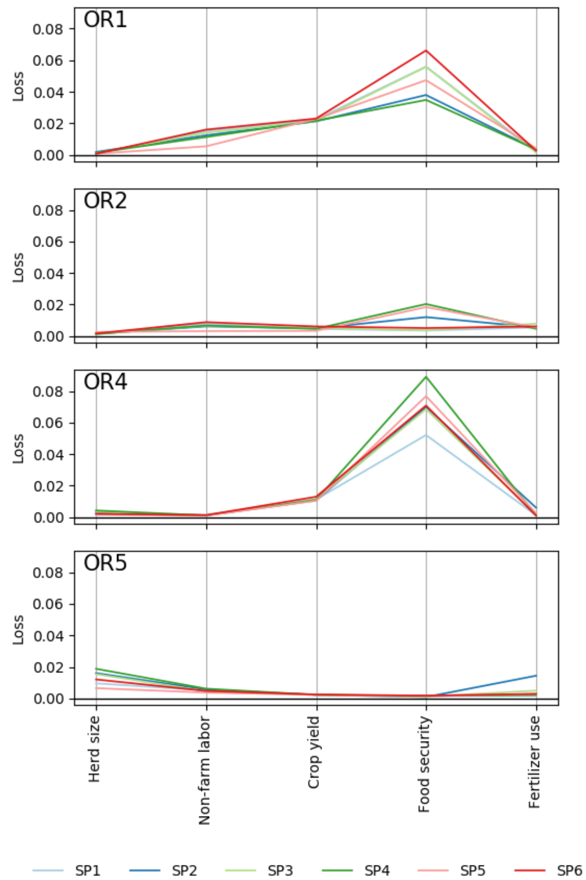

Figure E.5: Final loss values for each fitting pattern and site. Each line ("SP"=sub-population) represents the results for a different model calibration.

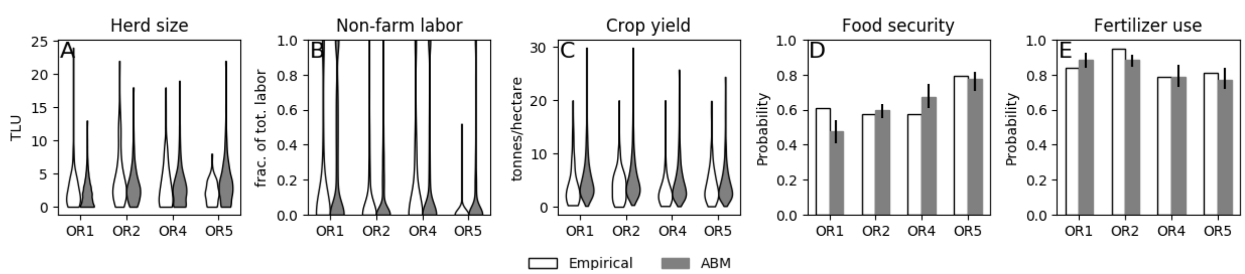

Figure E.6: Comparison of surveyed household livelihood characteristics to the calibrated ABM outputs for each site (OR=Oromiya) under a single model calibration. Bias across all six calibrations is relatively similar (Figure E.5) so we present only one here. ABM outputs above the 99th percentile are excluded from A, B, and C for visual clarity. The uncertainty bands in D and E represent the 95% prediction interval over 20 model replications. The data used to create the empirical patterns are summarized in Table E.1.

sitive to these parameters and so they have little effect on the loss calculations. As the different calibrations do not lead to structurally different results in the main analysis within this paper, the latter explanation (i.e., non-sensitivity) is the most likely. Nevertheless, in either case, presenting the results to the six calibrations achieves a higher level of robustness.

Figure E.7: Calibrated parameter values. Each line represents a different model calibration. Parameter ranges are bounded by the values given in Table E.3 and Table E.4 in the ODD+D description (section E.6).

E.5 Model validation experiment

E.5.1 Overview

We ran an additional experiment in which we isolate the effect of the LSLAs on smallholder food security. We compare the ABM-derived effects to empirical estimates using the household survey data. We disaggregate the effect both by site and by household displacement status. As we did not calibrate the model to match the empirical LSLA effects, this experiment serves as a form of model validation.

E.5.2 ABM effect estimation

For each site and each model calibration, we ran two simulations: control and treatment. The control simulation modeled the baseline subsistence conditions (i.e., no LSLA) and the treatment simulation included the LSLAs with displacement. For each agent, we calculated the difference in food security between these two simulations (i.e., what is the additional probability the agent is food (in)secure with the LSLA). We report the mean, 2.5%, and 97.5% agent-level effects.

E.5.3 Empirical effect estimation

To enable more robust empirical estimates of the LSLAs' effects on smallholder livelihoods, each "Treatment" site (shown in Figure 6.1 in the main manuscript and used to initialize the ABM) was matched with a nearby "Control" site (not pictured) that had not experienced an LSLA but had similar pre-LSLA socio-environmental characteristics. More details are given in (Williams et al.). The same household survey was given to approximately 100 households in each Control site. This Treatment-Control design enables us to conduct a two-stage matching procedure that controls for both site-level selection bias and household-level characteristics.

To control for the potential effects of unbalanced household-level characteristics between the Treatment and Control sites, we used covariate matching to match each Treatment household to a similar Control household (with replacement). We matched on a set of characteristics (Table E.2) that hypothetically affect both household food security and the LSLAs' placement within the landscape. Wherever possible, we drew from data recalling conditions before the LSLA.

When calculating site-level effects, we calculated the average difference in household food security between the matched Treatment-Control household pairs (i.e., what is the additional probability a Treatment household experienced food insecurity relative to its matched Control household). When calculating the overall effect for displaced households, we selected households within the Treatment sites that reported reductions in their landholdings, matched each of these to a single household within the corresponding Control site, then calculated the average difference in food security, pooled over all sites. In both cases, we report the estimated mean effect and 95% confidence interval on the estimate.

We note that this household-level matching is intended to give a general sense of the empirical patterns. We did not run comprehensive tests to verify the robustness of our results and so do not claim that they are free from statistical bias.

Table E.2: Covariates included within the household-level matching. All pre-LSLA variables were collected post-LSLA but asked households to recall pre-LSLA conditions.

Variable	Time	Unit
Total land cultivated	Pre	ha
Household size	Post	people
Livestock holdings	Pre	TLU
Use of inorganic fertilizer	Pre	Binary
Non-farm income	Pre	Binary
Access to electricity	Pre	Binary
Distance to forest	Pre	km

E.5.4 Results

The empirical estimates of the LSLAs' effects on food security reveal two main trends (Figure E.8): (1) the site-level effects are weak and inconsistent (significant decrease in food security in OR1, significant increase in OR3,[1] and non-significant outcomes in OR2 and OR4); and (2) households that lost land to the LSLA experienced substantially reduced food security. Together, these empirical results suggest a form of "selective marginalization" within the sites (Oberlack et al., 2016): some households were strongly affected and others not at all. These trends, combined with the difficulty we had in mechanistically explaining these effects with the available empirical data, provide motivation for our pooled analysis across sites (i.e., synthetic site "ORX") and our focus on land loss as the primary mechanism.

The ABM-derived effects roughly agree with the site-level empirical patterns (Figure E.8A) and reasonably match the magnitude of food security impacts for households losing land (Figure E.8B). In reality, LSLAs affect smallholder communities through a variety of additional mechanisms, and displaced households may receive compensation or adapt in ways that we do not model. We therefore do not expect to perfectly recreate the empirically estimated effects and rely on theoretical evidence supporting the structures within our ABM——in conjunction with the acceptable levels of fit to the empirical data——to justify its suitability for the intended application.

[1] Additional examination of the empirical data revealed a considerable site-level difference in the climatic conditions and crop yields between the OR3 Treatment and Control sites. This empirical effect should therefore be treated with caution.

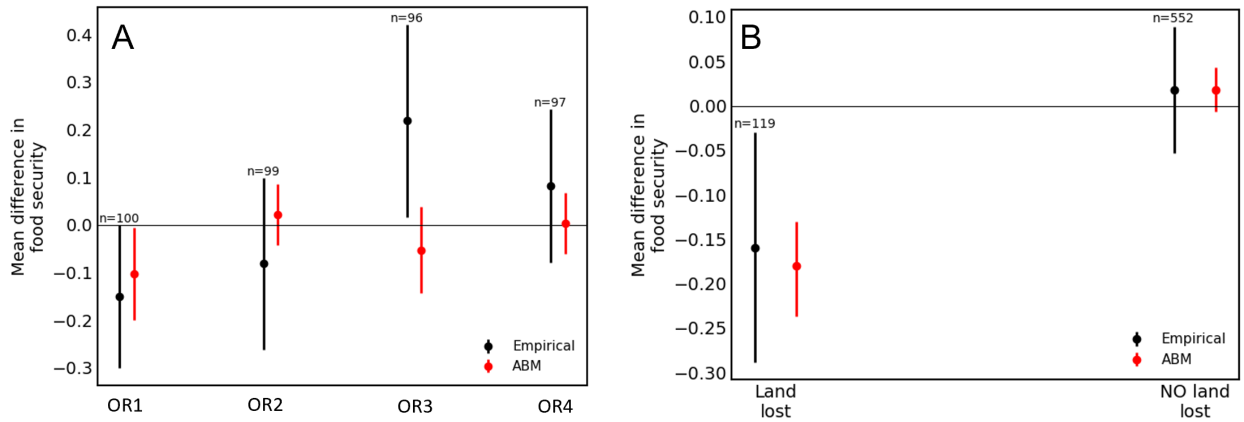

Figure E.8: Comparison of ABM-estimated and empirical LSLA effects. Probability that the LSLA increases household food security, disaggregated by (A) site and (B) land loss pooled over the sites. The empirical estimates represent the average difference between household food security in the Treatment sites and comparable households in the Control sites (selected using covariate matching). The ABM estimates represent the average and 2.5%/97.5% agent-level differences over 30 simulations and six calibrated parameter sets.

E.6 ODD+D model description

Here we provide an overview of the ABM in the ODD+D (Overview, Design Concepts, Details, and Decisions) format (Müller et al., 2013), including suggested modifications provided by Grimm et al. (2020).

E.6.1 Overview

E.6.1.1 Purpose and patterns

What is the purpose of the study? The model was designed to understand how alternative configurations of large-scale land acquisitions (LSLAs) and contract farming (CF) schemes may affect household food security and regional productivity within mixed cropping-livestock smallholder systems.

The model does not aim to make predictions of food security or productivity or to generate results directly relevant for informing policy in a specific socio-environmental context. Rather, the model is intended as an exploratory tool for examining how smallholder farmers may respond to changes in land access and intensification opportunities, as well as the potential implications of these changes for food security throughout a heterogeneous population.

For whom is the model designed? Researchers interested in development within smallholder agricultural systems.

What patterns are used as criteria for evaluating the model's suitability for its purpose? We evaluate our model by its ability to produce several patterns of livelihood characteristics, drawn from household survey data collected within four LSLA-affected regions of Ethiopia.

The patterns were selected to fulfil the following criteria:

- To provide information encompassing the range of dominant livelihood activities within these regions, as well as the outcomes of food security and productivity that we are interested in.
- To leverage the questions included in the household survey to inform modeled processes and decision-making alternatives.
- To represent the distribution of livelihood characteristics within the heterogeneous population, which is important given our focus on food (in)security.

The patterns include:

- Food shortages. Regional probability of a household experiencing difficulty meeting their food needs at some point within the past year.
- Crop yield (kg/ha).
- Livestock holdings (tropical livestock units (TLU)).
- Fertilizer application. Regional probability of applying inorganic fertilizer to agricultural land.
- Non-farm labor. Fraction of total productive labor allocated to non-farm income-generating activities.

As the patterns represent quantitative information, we cannot evaluate whether a given model configuration does or does not reproduce a given pattern. Instead, we used a genetic algorithm to identify model parameterizations that best match the empirical patterns, in conjunction with an iterative process of model structure development and refinement. Further details are provided in Appendix D.7.3.5.

E.6.1.2 Entities, state variables and scales

What kinds of entities are in the model?

- Agents. Each agent represents a single smallholder household. There are N agents (N=200 for the experiments in the main body of the manuscript).
- Livestock. Each agent manages a livestock herd.
- Agricultural land. Each agent manages a set of "fields".
- Common land. Used for grazing of livestock and represented as a single region-level stock.
- Model environment. This encompasses the exogenous drivers of the model:
    - LSLA/CF scenarios

- Climate
- Crop prices

Rationale: The above entities were included to encompass the primary livelihood activities (farming, livestock rearing, non-farm employment) and exogenous stressors (climate, prices) identified within the household survey data. We did not include forest-based aspects of smallholder livelihoods (e.g., firewood collection, hunting and gathering) as these are not dominant within the Oromiya region. We do not model the LSLA and CF schemes as agents, as we are interested in understanding the effects of given LSLA/CF arrangements on smallholder systems, rather than understanding drivers of LSLA/CF location or feedbacks from smallholders to the LSLA/CF decision-makers.

By what attributes (i.e. state variables and parameters) are these entities characterised?
Definitions:

- Static (/dynamic) – does (/does not) evolve throughout a simulation.
- Homogeneous (/heterogeneous) – constant (/variable) over the agent population.

Agents:

- Landholding (static, heterogeneous). Each agent manages a fixed, discrete number of fields. Agents do not buy, sell, or rent land, but it may be lost under conditions of LSLA.
- Household size (static, heterogeneous). Determines labor availability and food demand.
- Network (static, heterogeneous). Set of other agents with which experiences are shared and crop residues can be shared.
- Livestock herd size (dynamic, heterogeneous). See "livestock" below.
- Risk tolerance (static, heterogeneous).
- Beliefs (dynamic, heterogeneous). Beliefs are represented probabilistically and updated each year using Bayesian methods. See section 1.2.10.
- Behavior (dynamic, heterogeneous): o Fertilizer application (yes/no) o Livestock stocking (yes/no) o Salary employment allocation (no change, increase, decrease) o Contract farming land allocation (if relevant) (no change, increase, decrease)
- Salary employment (dynamic, heterogeneous). May be less than agents' allocation if regional demand is greater than supply.
- Wage employment (dynamic, heterogeneous). Used as a coping mechanism.
- Food security (dynamic, heterogeneous).

Livestock:

- Owner id (static, heterogeneous). Reference to the agent owning the livestock herd.
- Herd size (dynamic, heterogeneous). Represents large livestock and measured using integer values.

Agricultural land:

- Soil fertility (static, heterogeneous). Affects nutrients available for crop growth. We assume that the soil fertility is homogeneous over each agent's fields (i.e., soil fertility is represented at the agent-level, not the field-level).
- Fertilizer application (dynamic, heterogeneous).
- Crop yield (dynamic, heterogeneous).

Common land:

- Demand. Region-level sum of livestock requiring grazing on common land (dynamic, NA).
- Availability (hectares) (static, NA). As we do not endogenously model indirect land-use change, this is constant throughout the simulation. However, LSLAs and CF schemes can reduce common land availability from their time of implementation.

Model environment:

- Climate: rainfall amount (dynamic, homogeneous).
- Market: selling price for subsistence and cash crops (dynamic, homogeneous).
- Employment markets: region-level availability of salaried and wage employment (static, NA). This is static except under LSLA scenarios that provide employment.

What are the exogenous factors / drivers of the model?

- Climate. This affects crop yields.
- Crop prices for subsistence (/food) and cash (/non-food) crops.
- Employment markets. There is a limited availability of off-farm employment, which is modeled in two separate pools: salary and wage labor. The availability of employment within each market is static throughout the simulation but may be affected by the LSLA.
- LSLAs and CF. These are the primary exogenous drivers in our analysis. A LSLA or CF scheme is implemented at a pre-specified year within the simulation. We model three distinct LSLA/CF scenarios:
  - (1) LSLA with displacement. Agricultural and common land located within the LSLA are lost and converted to cash crop production (managed by the firm). Additional salary-based employment is created, depending on the implementation extent and employment intensity.
  - (2) forced CF. Common land located within the LSLA is lost and converted to cash crop production (managed by the firm). Agents retain agricultural land located within the LSLA but are forced to produce cash crops for the firm.
  - (3) opt-in CF. No land change occurs. All agents now have the ability to contribute land to the contract farming scheme. Each year agents can choose to increase, decrease, or make no change to their previous level of contribution.

If applicable, how is space included in the model? The model is spatially implicit.

- Agents are initialized in a 1D "space" (i.e., agent 1, 2, ..., N).
- Agents occupy heterogeneous areas of land that are conceptually contained within a 2D plane, however the land is not mapped into a 2D plane.
- All agents are assumed to have equal access to the reservoir of common land as well as agricultural and employment markets.
- Agents' networks are comprised of adjacent agents in both directions within the 1D agent "space".
- Agents are probabilistically affected by the LSLA.

What are the temporal and spatial resolutions and extents of the model?

- The temporal resolution is one year.
- In the main body of the paper, experiments are run for 30 years, conceptually representing the period of 2000-2029.
- Each field is 0.25 ha.
- The overall spatial extent is dependent on the number of agents, the agricultural land held by these agents, and the amount of common land. For example, households in OR2 each hold an average of 1.75 ha and agricultural land comprises 65% of the overall site area. Thus, with 200 agents there is 350 ha of agricultural land and 189 ha of common land, making a total of 539 ha, which is approximately 23 km x 23 km.

Rationale: We selected a one-year temporal resolution as it represents the frequency at which primary agricultural decisions are made and a 0.25 ha spatial resolution as it represents the resolution at which relevant agricultural management decisions are made. We selected a 30-year simulation period to ensure that our results encompass the long-term effects of LSLAs and CF schemes. A longer temporal extent would make modeling of processes such as demographic change and environmental degradation more critical. Although these processes are unquestionably important over a 30-year time period, we exclude them so as to focus on the direct effects of LSLA/CF on smallholder livelihoods.

E.6.1.3 Process overview and scheduling

What entity does what, and in what order? A single simulation step proceeds as follows:

1. Update the model environment:
    - The regional *climate condition*. The climate condition is simulated from a normal distribution truncated between 0 and 1. There is no correlation between years.
    - The regional *crop prices* for subsistence and cash crops. There is a correlation (ρ_{price})

between years for each crop type, such that $P_t = \rho_{price} P_{t-1} + (1 - \rho_{price})x$ where P denotes price, $x \sim N(\mu_{price}, \sigma^2_{price})$, and μ_{price} is defined separately for subsistence and cash crops. There is no correlation between the two crop prices in each year. Parameters are given in Table E.4.

2. If LSLA/CF implementation year: implement the LSLA/CF (see section E.6.4.3). This can affect the agents' *landholdings*, common land *availability*, and the regional *employment availability*.
3. Agents make annual livelihood decisions (*fertilizer use*, *livestock stocking*, *salary labor allocation*, *contract farming land allocation*). See section E.6.4.4.
4. The market allocates salaried employment among the candidate agents to determine their realized *salary employment*. See section E.6.4.5.
5. The model environment calculates each agent's *crop yields*. See section E.6.4.6.
6. Agents consume and purchase food using their available livelihood sources. This determines their *food security*. See section E.6.4.7.
7. Update livestock and rangeland (see section E.6.4.8):
    - Livestock consume crop residues from their owners' land
    - Livestock (if necessary) consume crop residues remaining on their neighbors' land
    - Livestock (if necessary) consume fodder from the common land
    - If fodder on common land is insufficient, destocking is apportioned between agents using the common land
    - Livestock reproduce, determining the final *herd size*
8. Agents engage in coping measures: *wage-based employment*, livestock selling (affecting herd size), and consumption reduction. See section E.6.4.9.
9. Agents update their *beliefs*. See section E.6.4.10.

Rationale: For all processes in which there is competition between the agents (employment allocation, livestock grazing on communal rangeland), the agent order is randomized at each time step to avoid artefacts of execution order. The loss of land to the LSLA is also randomly allocated between the agents. Refer to the "Stochasticity" section.

The model scheduling outlined above describes the chronological order of activities in smallholder agricultural systems (e.g., households make agricultural management decisions under uncertainty about crop yields, and production decisions are made asynchronously with consumption decisions (Sadoulet and De Janvry, 1995)). However, to simplify the scheduling, we assume that crop yields within year t are also consumed within year t (i.e., food consumption (step 6 above) occurs after food production (step 5)). This is somewhat discordant with reality, in which throughout a year, households consume crop yields from the previous harvest. This simplification should have no important implications for model dynamics.

E.6.2 Design concepts

E.6.2.1 Theoretical and Empirical Background

Which general concepts, theories or hypotheses are underlying the model's design at the system level or at the level(s) of the submodel(s) (apart from the decision model)? What is the link to complexity and the purpose of the model? The model is predicated on the notion that household food security is a complex, emergent outcome in smallholder agricultural systems. That is, food security emerges as a result of multiscale interactions between: socio-environmental context (land availability, soil fertility, climate); household-level attributes, decision-making, and interactions; and top-down institutional structures (LSLA/CF). Many types of models attend to these concepts (Müller et al., 2020).

Within the model, we measure food security using a version of a "food availability ratio" (Frelat et al., 2016), in which a food secure household has sufficient staple food available to meet their food needs. We do not represent a more comprehensive measure of food security that has empirical analogues (e.g., the household dietary diversity score (HDDS)) (Headey and Ecker, 2012; Nicholson et al., 2019), as this would considerably complicate the model and could not be validated against the available household survey data, which did not ask sufficient questions about dietary diversity. Yet, our model incorporates (to some extent) the availability, access, and stability pillars of the FAO's widely-employed definition of food security (FAO, 2008).

The other primary model output is agricultural production. Crop yields are calculated using the "yield gap" concept (Tittonell and Giller, 2013; Ferreira et al., 2017), which assumes that yield is influenced by the most constraining factor on production. We consider the effects of water, nutrient, and labor availability. Further details are given in section E.6.4.5. Although our implementation is stylized and not calibrated to field-level empirical measurements of yield under alternative climate and management conditions, it is conceptually similar to more complicated process-based yield models (e.g., CENTURY (Metherell, 1993), STICS (Brisson et al., 2003)) and through its simplicity enables more transparent mapping of relations between inputs and outputs.

On what assumptions is/are the agents' decision model(s) based? Agents are assumed to be boundedly rational actors that maximize a utility function using subjective, imperfect beliefs. The utility function considers the household's net income, subject to satisfaction of their food needs and expenditures.

Why is /are certain decision model(s) chosen? Unfortunately, the available household survey data did not provide sufficient information regarding decision-making processes for us to specify an empirically informed decision-making model. We therefore chose this approach as it does not require the specification of a large number of parameters (e.g., thresholds in satisficing or reference points for prospect theory).

If the model / submodel (e.g. the decision model) is based on empirical data, where do the data come from? Empirical data used to parameterize and calibrate the model were drawn from household surveys conducted by our research team in four regions of Oromiya, Ethiopia. Each of these regions was targeted by an LSLA.

At which level of aggregation were the data available? Household-level

E.6.2.2 Individual Decision Making

What are the subjects and objects of the decision-making? On which level of aggregation is decision-making modelled? Are multiple levels of decision making included? Decision-making is modeled at the household level. Households (the subjects) make decisions about their labor allocation, land management, and livestock management (the objects). Specifically, households choose between all feasible combinations of: fertilizer purchase (Y/N), invest savings in stocking livestock herd (Y/N), and off-farm salary employment (no change, decrease, increase). Under the CF_{join} scenario, households also make decisions about allocation of their land to the CF scheme (no change, decrease, increase).

What is the basic rationality behind agent decision-making in the model? Do agents pursue an explicit objective or have other success criteria? Agents pursue an explicit objective: they maximize the utility function $U(X_i) = 1 - exp(-X_i/R)$, where R is a risk tolerance and X represents anticipated net cash availability under decision option i. We note that concerns have been raised regarding economic optimization in agent-based models (Groeneveld et al., 2017; Schlüter et al., 2017). Although our utility function operates on a financial measure (i.e., net income), this represents the cash available after attempting to satisfy food and expenditure needs through own production and purchase from the market. Thus, food availability is implicit to the utility function. This is consistent with descriptions of household objectives in smallholder agricultural systems, where food provisioning and risk reduction are critical considerations in household decision-making (Demissie and Legesse, 2013).

How do agents make their decisions? Agents consider a finite set of decision options and select the one that (1) is feasible with respect to cash, land, and labor availability and (2) maximizes risk averse utility.

Do the agents adapt their behaviour to changing endogenous and exogenous state variables? And if yes, how? The decision-making process does not change throughout the simulation, however agents' beliefs about climate, prices, and the availability of off-farm employment evolve throughout the simulation, based on observation of endogenous (employment allocation) and exogenous (climate, prices) state variables. This affects the anticipated returns to different livelihood activities and therefore can affect agent behavior.

Do social norms or cultural values play a role in the decision-making process? No.

Do spatial aspects play a role in the decision process? No. The model is spatially implicit. However, agents observe their neighbors' outcomes (in a hypothetical 1D space), which affects the updating of beliefs.

Do temporal aspects play a role in the decision process? In the baseline model implementation, decisions are evaluated over a single-year time horizon. The only instance in which a longer horizon is considered is when cultivating a crop with a two-year harvest period (as part of the CF scheme). Here, agents consider the outcomes over two years, applying a time discounting rate to outcomes from the second year.

To which extent and how is uncertainty included in the agents' decision rules? Agent beliefs are represented probabilistically. In evaluating each decision option, agents consider a range of potential outcomes (e.g., a range of potential realized climate conditions), conditional on their probabilistic beliefs. This leads to a distribution of anticipated utility outcomes under each decision option. Agents then calculate the expectation of these utility values for each decision option. See further details in section E.6.4.4.

E.6.2.3 Learning

Is individual learning included in the decision process? How do individuals change their decision rules over time as consequence of their experience? No.

Is collective learning implemented in the model? No.

E.6.2.4 Individual sensing

What endogenous and exogenous state variables are individuals assumed to sense and consider in their decisions? Is the sensing process erroneous? Each year, agents observe the following:

- Climate: effect of climate on their own and their neighbors' crop yields
- Off-farm employment: their success (only if they sought it) and their neighbors' success (if they sought it)
- Regional crop prices
- Contract breaching: (if relevant) whether the firm honored their contract farming arrangement and their neighbors' arrangements.

These observations are not erroneous and are used to update agent beliefs.

What state variables of which other individuals can an individual perceive? Is the sensing process erroneous? Effect of climate on yields, off-farm employment success, contract breaching (as above).

What is the spatial scale of sensing? Household (agent) level.

Are the mechanisms by which agents obtain information modelled explicitly, or are individuals simply assumed to know these variables? No mechanisms are modeled.

Are the costs for cognition and the costs for gathering information explicitly included in the model? No.

E.6.2.5 Individual prediction

Which data do the agents use to predict future conditions? Their probabilistic beliefs (see section E.6.4.10).

What internal models are agents assumed to use to estimate future conditions or consequences of their decisions? For each decision option and realization of their probabilistic beliefs, agents estimate food and cash availability (steps 4-6 in section E.6.1.3) to calculate the utility of the option. This "internal model" is the same as the "external model," but is evaluated using the subjective, imperfect beliefs rather than the realized values.

Might agents be erroneous in the prediction process, and how is it implemented? No. Agents evaluate decision options under uncertainty, but it is not erroneous.

E.6.2.6 Interaction

Are interactions among agents and entities assumed as direct or indirect?

- Direct: information sharing through observations; grazing livestock on neighbors' leftover fodder.
- Indirect: agents compete in the employment market and for the limited availability of fodder in the communal rangeland.

On what do the interactions depend? All agents have equal access to employment markets and communal rangeland. Beliefs and livestock fodder are shared between neighbors.

If the interactions involve communication, how are such communications represented? Communication is not modeled.

If a coordination network exists, how does it affect the agent behaviour? Is the structure of the network imposed or emergent? N/A

E.6.2.7 Collectives

Do the individuals form or belong to aggregations that affect and are affected by the individuals? Are these aggregations imposed by the modeller or do they emerge during the simulation? No.

How are collectives represented? N/A

E.6.3 Heterogeneity

Are the agents heterogeneous? If yes, which state variables and/or processes differ between the agents? Yes. Agents are heterogeneous in their landholding, household size, network connections (neighbors), livestock herd size, risk tolerance, and beliefs.

Are the agents heterogeneous in their decision-making? If yes, which decision models or decision objects differ between the agents? Agents have heterogeneous risk tolerances and beliefs, but utilize the same decision-making model.

E.6.3.1 Stochasticity

What processes (including initialisation) are modelled by assuming they are random or partly random? All stochasticity is controlled using random number seeds. Multiple replications are run to produce distributions over this (within-model) stochasticity. Variability between experimental settings is reduced by using common random numbers across experiments.

Initialization:

- Generation of time series for climate, crop prices, and random perturbations for crop yields
- Sub-sampling of N agents from the surveyed households
- Assignment of agent-level risk tolerances and soil fertility

Simulation:

- Distribution of land loss between the agents in the LSLA simulations. There are two stochastic processes: (1) the order of agents from which land is sequentially taken and (2) the amount of each agent's land that is taken.
- Sampling over uncertain beliefs in decision-making.
- Allocation of non-farm employment between candidate agents
- Livestock reproduction
- Agent order for: sharing of leftover crop residues for livestock, destocking of livestock if communal rangeland is in deficit
- Contract breaching

E.6.3.2 Observation

What data are collected from the ABM for testing, understanding and analysing it, and how and when are they collected?

- Testing: to verify the model structure and implementation, we examined a range of model dynamics over time, including agent decisions, food security, crop yields, coping measures, and non-farm employment, as well as the effect of various LSLA and CF arrangements on these. We examined these both as averages/distributions over the population and by looking at individual agent trajectories.
- Calibration: to match with the timing of the household surveys, we ran simulations from 2000-2019 and collected the following household-level variables in the final year of the simulation: livestock herd sizes, fertilizer decisions, non-farm employment, food security, and crop yields. See Appendix E.4.
- Analysis: the primary outputs include household-level food shortages and household-level crop production. In the main body of the paper, we aggregate these over the agents and take the average over the simulation years following the LSLA/CF implementation.

What key results, outputs or characteristics of the model are emerging from the individuals? (Emergence) We consider food security to be an emergent outcome from the processes and interactions within the model. Some agents are predisposed to experience more or less food security (e.g., through their landholdings, which determines productive ability), yet transient or chronic food security can emerge from other interactions and stochasticity within the model, for example receiving a salaried job or being required to destock a livestock. Food security can be examined at the regional level (i.e., what percent of agents are food secure in year t) or disaggregated along dimensions of agent heterogeneity (e.g., how does food security compare between agents that did and did not lose land to the LSLA).

The second primary model output is regional productivity. In addition to direct effects resulting from the exogenous characteristics of the LSLA/CF scheme and (static) distributions of soil fertility, regional crop production is driven by agent-level fertilizer decisions, as well as agent-level decisions about contract farming participation (if relevant). These decisions are influenced by agent attributes (e.g., the financial ability to apply fertilizer), beliefs (e.g., the expected returns to applying fertilizer), and interactions (e.g., observation of contract breaches with neighboring agents). Each of these factors is dynamic and heterogeneous within the simulation.

E.6.3.3 Implementation Details

How has the model been implemented? Python 3.

Is the model accessible, and if so where? All data inputs and model and analysis code will be made publicly available on CoMSES.net upon acceptance of the publication.

E.6.3.4 Initialisation

What is the initial state of the model world, i.e. at time t=0 of a simulation run? Agents: The agents' initial state variables—landholding, household size, livestock herd, fertilizer use, and salary labor—are drawn from the household survey data, using the available recall questions. For the model calibration, we run a different model experiment for each site (i.e., OR1, OR2, OR3, OR4). For the experiments in the main body of the paper, we run a single synthetic site (ORX) by pooling the household data across all four sites. All variables are drawn jointly to account for dependencies in the data.

In all experiments, we initialize 200 agents. This number was chosen so that for each simulation there is a subsampling from the overall empirical household population (approximately 400 households, 100 from each site). Thus, for each random number seed, the population of households is different. We chose this approach to represent uncertainty more fully in household characteristics between simulations, which will result in wider prediction intervals.

For the agent state variables not drawn from the empirical data:

- Risk tolerance is drawn from a normal distribution (see parameters in Table E.4)
- Each agent's network is comprised of the m closest agents (m is an even number). For example, if m=4, agent i's network includes agents i-2, i-1, i+1, i+2.

- Beliefs are initially homogeneous over the agents and are set using the input parameters. This is described in more detail in section E.6.4.1.

Agricultural land: Soil fertility is drawn from a lognormal distribution (see parameters in Table E.4).

Common land: Remote sensing data was used to derive estimates of pre-LSLA common-land availability.

Is the initialisation always the same, or is it allowed to vary among simulations? The initialization process is always the same. However, the assignment of some state variables is probabilistic. The purpose of the model analysis is to understand the effects of LSLA/CF on model outcomes, so we desire to reduce the impact of this stochasticity. We do so by running the model multiple times for each LSLA/CF scenario using different random number seeds.

Are the initial values chosen arbitrarily or based on data? As described above, some initial state variables are derived from the household survey data. The initialized model therefore approximates the conditions in Oromiya, Ethiopia, but could be adapted in future work with survey data from different regions that are dominated by similar livelihood activities.

Variables not contained within this data were chosen as reasonable values as well as through the model calibration procedure. For the initialization of agent beliefs (which comprises the majority of arbitrary parameters), we reduce the impact of the initialization by (1) using a moderately weak prior, so agents' beliefs quickly adjust based on their experiences and (2) only recording the outputs after the LSLA/CF scheme is implemented (i.e., effectively a $\sim$10-year burn-in period).

E.6.3.5 Input data

Does the model use input from external sources such as data files or other models to represent processes that change over time? The only input data that drives model processes over time is the empirically estimated amount of common land within the LSLA, which is used to update common land availability when the LSLA is implemented.

E.6.4 Model and sub-model details

E.6.4.1 Additional initialization details

Table E.3 shows the parameters used in model initialization. Agent beliefs are represented probabilistically. Beliefs about non-negative, continuous quantities (rainfall and crop prices) are represented using normal distributions—i.e., $\sim N(\mu, \sigma^2)$. Their initialization requires several parameters. The use of these parameters is described in more detail in section E.6.4.10. Beliefs about probabilities (the probability of receiving salary labor and the probability of the firm honoring the contract) are represented using beta distributions—–i.e., $\sim Beta(\alpha, \beta)$. Agent beliefs for these quantities are initialized with a strength (γ) and an expectation (E), such that:

$$\alpha = E * \gamma \tag{E.4}$$

$$\beta = \alpha/\mu - \alpha \tag{E.5}$$

The second equation satisfies the property of a beta distribution that:

$$E = \frac{\alpha}{\alpha + \beta} \tag{E.6}$$

Together, the equations also mean that $\gamma = \alpha + \beta$, a measure of the strength of knowledge underlying the beta distribution. With a higher strength (γ), new observations have a weaker effect on the overall belief.

Table E.3: Values used in ABM initialization.

Entity	Parameter name	Symbol	Description	Value
Agent	Network N	$c_{network}$	Number of agents in network	4
	Risk tolerance: mean	μ_{risk}	Mean value of normal distribution	(75,7500)‡
	Risk tolerance: coefficient of variation	CV_{risk}	Ratio of mean and standard deviation in the normal distribution	0.1
Agent (beliefs)	Binomial strength	γ	Strength of prior beliefs on binomially distributed quantities	10
	Binomial expectation	E	Expected prior belief on binomially distributed quantities	
			- Receiving salary labor	0.2
			- Firm honors contract	1
	Normal mean	$E[\mu_0]$	Prior expectation on the mean	
			- Rainfall	μ_{rain}†
			- Crop prices	μ_{price}†
	Normal variance	$E[\sigma_0^2]$	Prior expected variance	
			- Rainfall	0.25
			- Crop prices	0.5
		n_0	Prior strength on the mean	1
		α_0	Prior strength on the variance	1
Land	Soil fertility		Conceptually represents log(kg Nitrogen/ha). Affects the nutrients available for crop growth.	
	- Mean	μ_{soil}		(7,9)‡
	- Standard deviation	σ_{soil}		(0,0.6)‡

‡These parameters are set by the calibration process. Given that we calibrate multiple models, these parameters assume multiple values. The values within parentheses indicate the range of values the parameter can take. See details in Appendix E.4.

Entity	Parameter name	Symbol	Description	Value

†The expectations for these beliefs are accurate—–i.e., they are equal to the actual expected values used to simulate rainfall and crop prices (see Table E.4).

E.6.4.2 Parameter values

Table E.4: ABM parameter values.

Parameter	Symbol	Value	Unit	Source	Description/notes
Model					
Number of agents	N_A	200			
Time extent	T	30	years		
Random seed	$seed$	0			Varied upon initialization
Agents					
Living cost	$\$_{nonfood}$	(10,5000)‡	birr/year/ person ¶		Annual non-food expenditure requirement.
Food requirement	req_{food}	206	kg/person/ year	(Williams et al., 2020a)	Represents 18 kg of staple crop consumption per person per month.
Number of years for smoothing	N_{yr_smooth}	3	years		Agents can choose to voluntarily destock from their herds at the end of the year if they do not anticipate they will be able to support them (see section E.6.4.8). When making this consideration, they recall the fodder availability over the previous 3 years.
Coping threshold	c_{coping}	(0,1)‡			Fraction of food deficit before agents engage in coping measures. If zero, then agents engage in coping measures with any amount of food shortage. If one, then agents can perfectly reduce their food consumption.
Agricultural labor requirement	req_{ag_lbr}	1	person/ ha		Agricultural labor requirements do not preclude agents from farming their land, but can lead to reduced yields when labor availability is low relative to land cultivated and livestock held (section E.6.4.6).
Livestock labor requirement	req_{ls_lbr}	0.2	person/ head		Similar to above, livestock labor requirements do not preclude agents from holding livestock. Livestock herd sizes are constrained by fodder availability.
Discount rate	$c_{discount}$	0.586		(Holden et al., 1998)	Only applied when producing crops with a two-year harvest period.
Number of simulations in utility calculations	N_{sim}	10			Within the decision-making process, we simulate a sampling over agents' probabilistic beliefs. This parameter denotes the number of samples that are taken. See details in section E.6.4.4. We chose a small number (10) for both computational reasons and so that there is some variability/stochasticity between model replications in the belief sampling.
Land					
Field size	$c_{fieldsize}$	0.25	Ha	HH survey	Minimum increment in household survey data.

Parameter	Symbol	Value	Unit	Source	Description/notes
Fertilizer application rate	c_{fert_app}	100	kg Nitrogen / ha	HH survey	Approximately median value for Oromiyan households in household survey data.
Soil organic matter mineralization rate	$c_{mineralize}$	0.02		(Schmidt et al., 2011)	50-year turnover rate of bulk soil organic matter (SOM)
Minimum nutrient loss rate	$c_{N\_lost\_min}$	0.05			Minimum fraction of nutrients lost through leaching (with high soil fertility)
Maximum nutrient loss rate	$c_{N\_lost\_max}$	0.5			Maximum fraction of nutrients lost through leaching (with low soil fertility)
Crop yields					
Maximum yield	c_{yield_max}	20,000	kg/Ha	HH survey	Tail from the distribution of maize yields in the household survey data.
Critical rainfall value	c_{rain_crit}	0.8			Rainfall value below which water begins to limit crop yield. Although the representation of water constraints is different, 0.8 is used in the CENTURY model (Metherell, 1993).
Random effect std. dev.	σ_{yield_error}	(0,0.6)‡			Standard deviation for the normally-distributed error term in crop yields.
Crop residue multiplier	$c_{residue\_mult}$	2		(Bogale et al., 2008; Assefa et al., 2013)	Residue production per unit of harvested crop.
Crop residues lost	$c_{residue\_lost}$	0.1		(Assefa et al., 2013)	Fraction of crop residues lost from system
Nitrogen composition					
- Residues	$N_{residue}$	0.37	% of total yield	(Elias et al., 1998)	
- Crop	N_{crop}	1.2	% of total yield	(Elias et al., 1998)	
Market					
Crop price					
- Mean	μ_{price}	6	Birr/kg	HH survey	Median selling price for maize reported within the household survey data
- Std. dev.	σ_{price}	0.5	Birr/kg		
- Inter-annual correlation	ρ_{price}	0.7			
Markup for buying food	c_{markup}	1.2			Ratio of the buying price to selling price. This proxies the effect of transaction costs.
Fixed farming cost	$\$_{farm}$	100	birr/ha		Represents all miscellaneous costs associated with farming. Insufficient information was available in the survey to empirically estimate.
Livestock cost	$\$_{livestock}$	3,000	birr/ head	CSA†	
Fertilizer cost	$\$_{fertilizer}$	13.2	birr/ kg	LSMS	Median value from the 2015 Living Standards Measurement Study (LSMS) in Ethiopia.
Labor allocation increment					

Parameter	Symbol	Value	Unit	Source	Description/notes
- Salary	$incr_{salary}$	0.5	person-years		Increment at which salary labor can be allocated. 0.5 represents one person working half of their time.
- Wage	$incr_{wage}$	0.005	person-years		Increment at which wage labor can be allocated. 0.005 = 1/200, conceptually representing 1 day.
Returns to non-farm work					
- Salary	$\$_{salary}$	(1e4, 2e4)‡	birr/ person/ year		
- Wage	$\$_{wage}$	(1e4, 2e4)‡	birr/ person/ year		
Regional job availability					
- Salary	$avail_{salary}$	(0,0.1)‡	person-years/ year/ agent		
- Wage	$avail_{wage}$	(0,0.5)‡	person-years/ year/ agent		
Climate					
Rainfall					
- Mean	μ_{rain}	(0.4,0.8)‡			
- Std. dev.	σ_{rain}	0.2			
Rangeland					
Fodder availability	$avail_{fodder}$	(750,4000)‡	kg/ha		
Livestock					
Birth rate	c_{reprod}	(0,0.5)‡			Probability that a livestock head reproduces each year.
Consumption	c_{ls_cons}	(2000,4000)‡	kg/ head/ year		
Income	$\$_{livestock\_earn}$	175	birr/ head/ year		Derived from (Redda, 2002): 240-480 birr/year with a local cow. Assume mid-point of ∼350 birr/year and 50% female animals, which gives 175 birr/head/year.

‡These parameters are set by the calibration process. Given that we calibrate multiple models, these parameters assume multiple values. The values within parentheses indicate the range of values the parameter can take. See details in Appendix E.4.

¶Birr is the Ethiopian currency.

†CSA = CSA = Ethiopian Central Statistical Agency; http://www.csa.gov.et/monthly-retail-price

E.6.4.3 LSLA/CF implementation

Change in model parameters and processes In all cases, the LSLA/CF is implemented at the beginning of a pre-specified year, based on the best available information (OR1: 2012, OR2: 2012, OR3: 2003, OR4: 2008, ORX: 2010). The event permanently affects some parameters and state variables, and in some cases opens up new livelihood opportunities (i.e., affects processes within the model). The parameters associated with the LSLA/CF implementation are shown in Table 6.3

in the main manuscript. Table E.5 describes the changes associated with each modeled LSLA/CF scenario.

Table E.5: Effects of LSLA/CF scenarios on model state variables and processes.

Geographic scenario	**Parameter / process affected**	**Details**
$Displacement$	Increase salary labor supply	Conceptually represents employment on the plantation-style farm. Employment addition (jobs) = employment rate (jobs/ha) * LSLA area (ha) * implementation fraction
	Reduce common land availability	Common land lost (ha) = LSLA area (ha) * fraction in common land
	Displace agricultural land	Agricultural land lost (ha) = LSLA area (ha) * (1 – fraction in common land) To apportion this between the agents: • Generate a random ordering for the N agents: $random\_index$ • Generate a list of N values between 0 and 1: $random\_fraction$ • Set $land\_taken \leftarrow 0$ ha • Set $r \leftarrow 1$ • while $land\_taken <$ agricultural land lost: – Take $random\_fraction$ of agent $random\_index[r]$'s land – Add this to $land\_taken$ – Set $r \leftarrow r + 1$
	Grow plantation crops	• This process operates every year after the LSLA's implementation. • Set the soil fertility within the LSLA as the median over all agents' soil fertility. • Assume fertilizer is applied at the cash crop rate, which is the baseline subsistence rate multiplied by the "intensification" parameter (e.g., 1.5), shown in Table 6.3 in the main manuscript. • Calculate crop yield using the same process as for agent crop yields, described in section E.6.4.6. • Crop production = crop yield * LSLA area * implementation fraction
CF_{forced}	Reduce common land availability	As in $Displacement$ above
	Assign agents to contract farming scheme	Use the same process as "Displace agricultural land" above, but instead of each agent losing this amount of land, this denotes the amount of land they must contribute to the contract farming scheme.
	Farm cash crops	Note: The way in which cash crops interact with agents' income and food consumption is described in section E.6.4.7. Here we summarize the key differences between farming of subsistence crops and cash crops. • Agents apply fertilizer at the cash crop rate. • Cash crop production cannot be consumed and must be sold to the market at the cash crop price. • Agents know this price with certainty at the beginning of the year (i.e., during their decision-making process). This represents the price guaranteed by the firm.
	Grow plantation crops	Same as "Grow plantation crops" above, but only operates on the area of LSLA within common land.

Geographic scenario	Parameter / process affected	Details
CF_{choice}	Add contract farming to decision options	Complement the set of agents' decision options with the contract farming decisions: $\{$no change, increase, decrease$\}$. The magnitude of the increase/decrease is specified by the "land requirement" parameter in Table 6.3 in the main manuscript (e.g., 0.25 ha – equivalent to one field).
	Add possibility for contract breaching	If contract breaching is enabled: • Activate agents' belief about the probability of the firm honoring the contract. • Within the utility calculations in the decision-making: include agents' uncertainty about the outcome of the contract farming. More details are included in section E.6.4.4. • After crop yields are calculated: probabilistically simulate contract breaches. • If a contract is breached, the agent loses some cash crop production (specified by "production losses" in Table 6.3 in the main manuscript) and must sell their remaining cash crop production at the subsistence crop price (i.e., without the market premium).

Attributing household-level land losses In developing the model, we experimented with several different methods for attributing the LSLA-induced land losses between the agents. There were two available potential sources of empirical information for this: the household survey data and the remotely sensed LULC data. The household survey data provided estimates at the household-level about changes in cultivated land area, in both the Treatment and paired Control sites. Although these data most closely approximate the empirical conditions, we decided not to use them. The primary factor motivating this choice was that households in both Treatment and Control sites reported both increases and decreases in land cultivation. It was difficult to empirically identify the LSLA-induced component of these changes, making it difficult to identify the LSLA-induced effect within the ABM. Because we do not endogenously model decisions around land cultivation (e.g., land rental), we opted for an arrangement that more directly isolates the effects of the LSLAs on smallholder agricultural land.

To do this, we used the site-level LULC data. These data provide site-level estimates of LSLA-induced losses of smallholder agricultural land (see Figure 6.1 and Table 6.1 in the main manuscript). We experimented with two methods to attribute these site-level changes between the agents. In the first option—–"random_percent" described in Table E.5—–the order of the agents is randomized and each agent sequentially loses a random percentage of their land until the site-level changes are satisfied (Figure E.9B). In the second option—–"random_field"—–fields are randomly selected from the landscape until the site-level changes are satisfied (Figure E.9C). The first option generates a higher degree of differentiation, whereby many agents are unaffected but some agents lose large amounts of land. This better approximated the distribution of land changes in the household survey data (Figure E.9A), so we chose to use this as the default method.

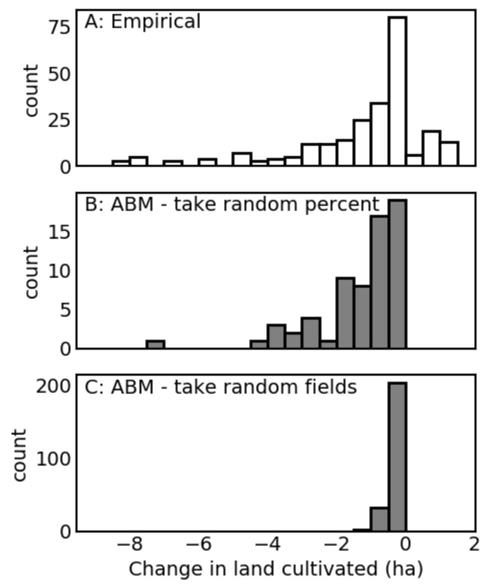

Figure E.9: Changes in cultivated land area for households reporting change (i.e., excluding households with no change). (A) shows the values reported by the surveyed households, with data pooled across sites. (B) and (C) respectively show the distribution across households in the ABM using the "random_percent" and "random_field" algorithms.

E.6.4.4 Agent decision-making

Decision options Under baseline conditions, decision options include all (maximum 12 (2*2*3)) feasible combinations of:

- Fertilizer purchase
  - Yes: Apply inorganic fertilizer at a rate of c_{fert_app} (Table E.4) to all land; or
  - No: Do not apply inorganic fertilizer.
- Livestock stocking
  - Yes: Purchase livestock with any leftover cash from the previous year; or
  - No: Do not purchase any livestock.
- Off-farm salary employment:
  - No change: allocate the same amount of labor as the previous year; or
  - Increase: increase allocation by an increment of $incr\_salary$ (Table E.4); or
  - Decrease: decrease allocation by an increment of $incr\_salary$.

Under the CF_choice arrangement, the decision set is expanded to all include all (maximum 36) feasible combinations of:

- Contract farming:
  - No change: allocate the same amount of land as the previous year; or

– Increase: increase the amount of land contracted by one increment (given in Table 6.3 in the main manuscript); or
– Decrease: decrease the amount of land contracted by one increment.

Incorporating uncertainty in beliefs Agents anticipate the returns they might receive under each decision option, given their uncertain beliefs about climate, prices, and non-farm employment allocation. The beliefs are represented using two different distributions (see section E.6.4.1): normal (climate and prices) and binomial (non-farm employment allocation and, if relevant, contract breaching). To explicitly represent uncertainty over all beliefs——without relying on an analytical, multivariate probability distribution——we conduct a sampling over the individual distributions.

For the normally distributed quantities, we achieve this through the following steps:

1. Draw N_sim values ($N\_sim = 10$) from a standard normal distribution: $Z_a \sim N(0, 1)^{N_{sim}}$
2. Scale to a sequence of real-valued quantities using agent a's beliefs: $X_a = E[\mu]_a + E[\sigma]_a * Z_a$

For the binomially distributed quantities, we achieve this through the following steps:

1. Draw N_{sim} values from a uniform distribution: $Z_a \sim U(0, 1)^{N_{sim}}$
2. Convert to a sequence of success/failure using agent a's beliefs: $X_a = \mathbb{1}(Z_a \leq E_a)$, where $\mathbb{1}$ is an indicator function and E_a is the expected value of agent a's belief.

Evaluation procedure To evaluate each option——in order to calculate a utility——agents run an internal process that simulates their livelihood if they were to choose this option, using the N_{sim} artificial realizations of climate, prices, and employment. The process contains the steps:

1. Calculate the start-of-year cost of the option, including non-food expenditure, livestock investment, and costs for farming and fertilizer application.
2. Estimate income from livestock and salary employment, using the N_{sim} realizations of employment allocation.
3. Estimate crop yields, using the N_{sim} realizations of climate and the known fertilizer and crop residue applications.
4. If contract farming: Estimate income from cash crops, using the N_{sim} realizations of crop yields and the cash crop price guaranteed by the contracting firm.
5. If contract breaching is active: Reduce cash crop production and income, using the N_{sim} realizations of contract honoring/breaches.
6. If subsistence crop production exceeds food consumption requirements: Sell excess subsistence crop production using the N_{sim} realizations of subsistence crop prices.
7. Else: Purchase food to fill the consumption deficit using the N_{sim} realizations of subsistence crop prices.

8. Return the N_{sim} values of net cash availability, accounting for the value of livestock. Here, each value could be positive or negative, respectively representing cash surplus and deficit. A cash deficit indicates that coping measures will be required, but these are not simulated here.

Finally, agent a calculates the expected utility of the decision option as:

$$EU_a = \frac{1}{N_{sim}} \sum_{i=1}^{N_{sim}} 1 - exp\Big(-\frac{cash_{i,a}}{R_a}\Big) \tag{E.7}$$

where R is the risk tolerance.

Feasibility Each decision option is subject to a feasibility check. An option is infeasible if any of the following conditions are met:

1. A labor allocation is negative (e.g., if attempting to decrease salary employment below zero);
2. A non-farm labor allocation is larger than the household's labor capacity;[2]
3. Total land cultivation is larger than the household's land availability (e.g., if attempting to increase contract farming land); or
4. For a two-year cash crop only: the two-year crop is being abandoned half-way through its growth period.

Selection procedure Agents select the feasible option with the highest utility.

E.6.4.5 Salary employment allocation

Agents compete for a finite supply of salaried employment. Except in simulations where the LSLA provides additional salaried employment (Table E.5), employment availability remains constant throughout the simulation. We assume that there is some inertia in salaried employment: agents retain salaried employment between years. Salaried jobs are allocated using the following procedure:

- For agents previously with salaried employment and choosing to retain previous allocation:
  - Agents retain jobs.
- For agents previously with salaried employment and reducing their previous allocation:
  - Add these jobs back to the market.

[2]Note: we only incorporate labor constraints in the allocation of non-farm labor, and not in the allocation of labor to farming or livestock rearing (i.e., labor availability does not directly constrain the ability to cultivate crops or hold livestock). This is because some households in the survey data led to infeasible allocations of labor (i.e., higher labor use than labor availability) under the conditions within the initialization data. Labor allocated to farming and livestock therefore only affects households through reductions in their crop yield (see section E.6.4.6).

- For agents seeking to increase allocation:
  - If total new demand $<$ jobs available:
    - All agents receive jobs.
  - Else:
    - Randomize the order of agents.
    - Allocate jobs in this order at the $incr\_salary$ rate (Table E.4) until none remain.

E.6.4.6 Crop yields

Crop yields are calculated using a stylized, process-based model, slightly adapted from (Williams et al., 2021). In contrast to other more complicated process-based yield models (e.g., DSSAT), this yield model does not attempt to predict site-level yields. Rather, our yield model has two aims: (1) to be simple, so as to retain interpretability; and (2) to recreate the distribution of crop yields reported by the surveyed households. Hence, we retain a simple structure and include the empirical yield distribution as an ABM fitting pattern (Table E.1).

All agents grow a single staple food crop, which conceptually represents maize——the dominant crop in the modeled region. Crop yields are calculated using the yield gap concept (Tittonell and Giller, 2013), in which yields are constrained by the most limiting factor to crop growth. We proxy the effects of water availability, nutrient availability, and labor intensity. We describe these effects below.

Water availability Under baseline conditions, water is provided exclusively by rainfall (i.e., no irrigation) and is homogeneous each year across all agents. The water limitations are calculated as follows (using parameters from Table E.4):

1. Sample the annual rainfall value:

$$rain \sim N(\mu_{rain}, \sigma^2_{rain}) \tag{E.8}$$

2. Convert to a water factor:

$$\begin{Bmatrix} 1 & if rain \geq c_{rain\_crit} \\ 1 - \frac{c_{rain\_crit} - rain}{c_{rain\_crit}} & else \end{Bmatrix} \tag{E.9}$$

In the LSLA/CF experiments involving irrigation, $c_{water} = 1$ in all years for irrigated land. The water factor provides a linear scaling on maximum yield:

$$Y^{water} = c_{water} * c_{yield\_max} \tag{E.10}$$

Nutrient availability We represent the effects of nutrient availability using a generic "nutrient" and a partial nutrient balance. To ground parameter values in empirical data, we measure this in units of kg Nitrogen (N)/ha. N is an important nutrient for crop growth and is frequently a limiting factor for yields in smallholder agricultural systems (Giller et al., 1997). Crops can only uptake inorganic forms of N, which are provided by inorganic fertilizer, mineralization of soil organic matter, and mineralization of applied organic nutrients. Agents have heterogeneous levels of soil fertility that provide inorganic N for crop growth each year through a linear decay process. We do not model soil degradation and so soil fertility is held static throughout the simulation. A portion of this organic matter mineralizes within the year to provide nutrients for crop growth. The total amount of inorganic nutrients available for agent a is given by:

$$N_a^{total} = N_a^{fertilizer} + k_{slow} * soil\_fertility_a + k_{fast} * N_a^{crop\_residue} \tag{E.11}$$

where

- $N_a^{fertilizer} = c_{fert\_app}$ if the agent has chosen to apply fertilizer and 0 otherwise
- k_{slow} and k_{fast} are rate constants given in Table E.4
- $N_a^{crop\_residue} = c_{residue\_mult} * (1 - c_{residue\_lost} * crop\_production_{a,t-1})$. That is, applied crop residues are a scaled value of the previous year's crop yields.

Next, some fraction of N_a^{total} is lost through leaching.

$$N_a^{available} = N_a^{total} * (1 - loss\_rate_a) \tag{E.12}$$

where higher soil fertility contributes to lower leaching rates (Drinkwater et al., 1998):

$$loss\_rate_a = c_{N\_lost\_min} + \frac{max\_fertility - soil\_fertility_a}{max\_fertility} * \left(c_{N\_lost\_max} - c_{N\_lost\_min}\right) \tag{E.13}$$

where $max\_fertility$ denotes the highest soil fertility over all agents.

The maximum possible crop yield, given these nutrients, is then calculated as:

$$Y_a^{nutrient} = \frac{N_a^{available}}{N_{crop} + c_{residue\_mult} * N_{residue}} \tag{E.14}$$

This represents a partitioning of the available N between the residues and harvested crop.

Labor intensity We proxy an effect of limited labor availability on crop yields. This affects agents with small household sizes relative to their landholdings, as well as agents allocating a lot of labor to non-farm activities, who must negotiate a tradeoff with crop yields. The labor effect is calculated as:

$$labor\_effect_a = 1 - exp\Big(- \frac{household\_size_a}{labor\_allocated_a} \Big) \tag{E.15}$$

where $labor\_allocated$ is the sum of all labor allocated to farming, livestock, and non-farm employment.

Yield calculation These three reduction factors are then combined to calculate the crop yield:

$$Yield_a = min(Y^{water}, Y_a^{nutrient}) * labor\_effect_a * \epsilon \tag{E.16}$$

where $\epsilon \sim N(1, \sigma^2_{yield\_error})$ is a household-level stochastic effect.

E.6.4.7 Income, food consumption, and food security

This process simulates the agents' satisfaction of food and expenditure requirements through their set of livelihood sources. It allows for self-consumption of food crop production as well as buying and selling of food from the market. The process is similar to the internal evaluation procedure conducted by agents in their decision-making (see section E.6.4.4), but it utilizes realized values of crop yields, employment, and prices, rather than the agents' beliefs. The process runs as follows.

Income balance The income balance accounts for the agents' non-food expenditures and earnings:

$$net\_income = herd\_size * \$_{livestock\_earn} + salary\_labor * \$_{salary} - \$_{nonfood} - \$_{farm\_total} \tag{E.17}$$

where all cost parameters are as specified in Table E.4 and $\$_{farm_total}$ includes the fixed costs of farming (at rate $\$_{farm}$ in Table E.4) and the costs for fertilizer (at rate $\$_{fertilizer} * c_{fert\_app}$). It is possible that this income balance is negative (i.e., the agent is in debt). In this income balance, we assume that surplus income cannot be carried from the previous year (i.e., agents do not have access to cash savings accounts). We made this model design decision to prevent the emergence of "runaway" households that, if earning a net-positive income in an average year, accumulate progressively more and more cash over time. In reality, richer households in such contexts likely spend a large proportion of their excess income (e.g., on weddings or celebrations) and thus savings is not as high. As we do not model such behavior, we simply assume that any surplus cash at the end of the time step is spent.

Crop production First, agents sell any cash crops (produced through contract farming) at the prevailing cash crop market rate. If contract breaching is included, agents with whom the firm breaches lose a fraction of their cash crop production (specified in Table 6.3 in the main manuscript) and sell at the market price for subsistence crops. Next, agents account for their food crop production, with the following priority ordering:

1. Sell food crops to absolve any remaining debts from the income balance (i.e., to repay required costs)
2. Satisfy own food consumption needs (i.e., self-consume production)
3. Sell any remaining production.

Food purchase If a food consumption deficit remains, agents then purchase food from the market at the prevailing market rate with an added markup (c_{markup} in Table E.4) that proxies the effects of transaction costs. It is possible that agents have insufficient cash to do this——i.e., a food deficit remains.

Food shortages If a food deficit remains after this process—i.e., the agent has insufficient production and income sources to procure the required food and non-food quantities—they are classified as "food insecure". This is a binary measure that is modeled at the household-level within the ABM. It incorporates to some extent the FAO's food security dimensions of availability, access, and stability (FAO, 2008). However, given that we model production and consumption of a single cereal food item, our food security measure does not imply that a household consumes adequate nutrients to meet dietary needs or is able to utilize these nutrients within their bodies. Additionally, we do not model the distribution of food between different household members.

E.6.4.8 Livestock consumption, reproduction, and stocking

Consumption Livestock consumption requirements are met through three pathways. First, each agent's livestock are grazed on the agent's own crop residues. The volume of available crop residue is calculated using a multiplication factor ($c_{residue\_mult}$ in Table E.4) and loss factor ($c_{residue\_lost}$ in Table E.4) on crop production within the same year. Second, agents unable to meet their livestock's consumption needs through their own crop residues graze their livestock on their neighbors' leftover residues (if any remain). To do so, the order of agents is randomized and each agent sequentially grazes livestock on their neighbors' remaining residues.[3] Third, remaining livestock consumption shortfalls are met by the communal rangeland. The communal rangeland has a limited availability of fodder, dictated by its area and biomass density ($avail_{fodder}$ in Table E.4). If this reserve is insufficient to meet all livestock needs, the livestock that cannot be fed must be destocked. Destocking is allocated randomly between the livestock that are grazed on the communal

[3] In reality, smallholder households have an incentive to let other livestock graze on their land due to its benefits to soil quality (e.g., through livestock manure).

rangeland. Agents do not earn money for these lost livestock.

Reproduction Each animal has a probability of reproducing in a given year (c_{reprod} in Table E.4). Each livestock herd's reproduction is simulated using sampling from a binomial distribution.

Destocking If, after these processes, an agent holds more livestock than they anticipate to be able to feed in the subsequent year, they destock from their herd. Agents use their previous levels of fodder availability (on-farm and off-farm) to estimate the maximum number of livestock they will be able to support. Any animals destocked here are sold at the market rate ($\$_{livestock}$ in Table E.4).

E.6.4.9 Agent coping mechanisms

Agents that cannot meet their food or cash needs engage in the following three coping mechanisms.

1. Reduce food consumption: Agents have a coping threshold (c_{coping} in Table E.4), which represents the degree to which they are able to reduce their food consumption before engaging in other coping measures. Agents that have food deficits reduce their food consumption up to this coping threshold.
2. Casual labor: Agents with remaining food or cash deficits then attempt to find casual, wage-based non-farm employment. There is a limited availability of wage-based jobs at the regional level ($avail_{wage}$). If the total demand exceeds the supply, jobs are allocated randomly between the agents seeking employment. The allocation is conceptually made on a daily basis (e.g., an agent might seek 20 days of work but only receive 12), using the $incr_{wage}$ parameter in Table E.4.
3. Sell livestock: If a food or cash deficit remains after the above two coping measures, agents then sell the required number of livestock to make up this shortfall. Agents that have insufficient livestock to do so sell the maximum number possible and are absolved of their debt at the end of the year (i.e., shortfalls cannot carry between simulation time steps).

E.6.4.10 Agent belief updates

Agents' beliefs are represented probabilistically (see section E.6.4.1 for initialization details) and are updated with Bayesian methods using agents' own experiences and observation of their neighbors' experiences. For simplicity, we assume that agents treat their neighbors' experiences with equal weight to their own experiences.

Normally distributed beliefs We assume that crop prices and climate (symbolically: X) follow a normal distribution with unknown mean and variance:

$$X \sim N(\mu, \tau) \tag{E.18}$$

where $\tau = 1/\sigma^2$ is the precision. The agent has beliefs on the μ and τ parameters as follows:

$$\mu \sim N(\mu_0, n_0\tau) \tag{E.19}$$

$$\tau \sim Ga(\alpha, \beta) \tag{E.20}$$

where n_0 is the prior strength on the mean. Using an observation, x, these beliefs are updated using:[4]

$$\mu_1 = \frac{n * x * n_0 * \mu_0}{n_0 + n} \tag{E.21}$$

$$n_1 = n_0 + n \tag{E.22}$$

$$\alpha_1 = \alpha_0 + \frac{n}{1} \tag{E.23}$$

and when $n = 1$ (i.e., for a single new observation)

$$\beta_1 = \beta_0 + 1/2 * \left(x - \mu_0\right)^2 \left(1 + \frac{n * n_0}{2 * (n + n_0)}\right) \tag{E.24}$$

When using their beliefs to make predictions of future conditions, agents evaluate the expectation of μ and τ:

$$E[\mu] = \mu_0 \tag{E.25}$$

$$E[\tau] = \frac{1}{E[\sigma^2]} = \alpha/\beta \tag{E.26}$$

Beta distributed beliefs Beliefs that follow a beta distribution—probabilities of receiving non-farm employment and firm honoring contract—are updated using the (more standard) beta-binomial conjugate prior. Given n successes out of m tries, α and β are updated as:[5]

$$\alpha_1 = \alpha_0 + n \tag{E.27}$$

[4]Simplified from Lemma 12 in https://people.eecs.berkeley.edu/~jordan/courses/260-spring10/lectures/lecture5.pdf

[5]Note that this refers to different α and β than for the normally distributed beliefs

$$\beta_1 = \beta_0 + m - n \tag{E.28}$$

Again, when using their beliefs to make predictions of future conditions, agents use the expected value of the Beta distribution:

$$E[X] = \frac{\alpha}{\alpha + \beta} \tag{E.29}$$

Bibliography

C. G. Abreu and C. G. Ralha. An empirical workflow to integrate uncertainty and sensitivity analysis to evaluate agent-based simulation outputs. *Environmental Modelling & Software*, 107: 281–297, Sept. 2018.

A. Adiga, S. Chu, S. Eubank, C. J. Kuhlman, B. Lewis, A. Marathe, M. Marathe, E. K. Nordberg, S. Swarup, A. Vullikanti, and M. L. Wilson. Disparities in spread and control of influenza in slums of Delhi: findings from an agent-based modelling study. *BMJ open*, 8(1):e017353, Jan. 2018.

A. Agrawal, D. G. Brown, and J. A. Sullivan. Are Global Land Grabs Ticking Socio-environmental Bombs or Just Inefficient Investments? *One Earth*, 1(2):159–162, Oct. 2019.

N. K. Ajami, Q. Duan, and S. Sorooshian. An integrated hydrologic Bayesian multimodel combination framework: Confronting input, parameter, and model structural uncertainty in hydrologic prediction. *Water resources research*, 43(1), 2007.

D. Ali, K. Deininger, and A. Harris. Does Large Farm Establishment Create Benefits for Neighboring Smallholders? Evidence from Ethiopia. *Land economics*, 95(1):71–90, Feb. 2019.

S. Alobo Loison. Rural Livelihood Diversification in Sub-Saharan Africa: A Literature Review. *The journal of development studies*, 51(9):1125–1138, Sept. 2015.

M. L. Amadou, G. B. Villamor, and N. Kyei-Baffour. Simulating agricultural land-use adaptation decisions to climate change: An empirical agent-based modelling in northern Ghana. *Agricultural systems*, 166:196–209, Oct. 2018.

A. Amare and B. Simane. Determinants of smallholder farmers' decision to adopt adaptation options to climate change and variability in the Muger Sub basin of the Upper Blue Nile basin of Ethiopia. *Agriculture & Food Security*, 6(1):64, Dec. 2017.

T. Amsalu and S. Addisu. Assessment of Grazing Land and Livestock Feed Balance in Gummara-Rib Watershed, Ethiopia. *Current Agriculture Research Journal*, 2(2):114–122, 2014.

L. An. Modeling human decisions in coupled human and natural systems: Review of agent-based models. *Ecological modelling*, 229:25–36, Mar. 2012.

C. R. Anderson, J. Bruil, M. J. Chappell, C. Kiss, and M. P. Pimbert. From Transition to Domains of Transformation: Getting to Sustainable and Just Food Systems through Agroecology. *Sustainability: Science Practice and Policy*, 11(19):5272, Sept. 2019.

A. Angassa and G. Oba. Relating long-term rainfall variability to cattle population dynamics in communal rangelands and a government ranch in southern Ethiopia. *Agricultural systems*, 94 (3):715–725, June 2007.

J. Angwin, J. Larson, S. Mattu, and L. Kirchner. Machine bias. *ProPublica, May*, 23(2016): 139–159, 2016.

I. G. K. Ansah, C. Gardebroek, and R. Ihle. Resilience and household food security: a review of concepts, methodological approaches and empirical evidence. *Food Security*, 11(6):1187–1203, Dec. 2019.

W. Anseeuw and G. M. Baldinelli. Uneven Ground: Land Inequality at the Heart of Unequal Societies. Technical report, International Land Coalition, 2020.

C. Arndt, R. Benfica, F. Tarp, J. Thurlow, and R. Uaiene. Biofuels, poverty, and growth: a computable general equilibrium analysis of Mozambique. *Environment and Development Economics*, 15(1):81–105, Feb. 2010.

J. C. Arnott, R. J. Neuenfeldt, and M. C. Lemos. Co-producing science for sustainability: can funding change knowledge use? *Global environmental change: human and policy dimensions*, 2020.

S. R. Arnstein. A Ladder Of Citizen Participation. *Journal of the American Institute of Planners*, 35(4):216–224, July 1969.

D. Assefa, A. Nurfeta, and S. Banerjee. Assessment of feed resource availability and livestock production constraints in selected Kebeles of Adami Tullu Jiddo Kombolcha District, Ethiopia. *African Journal of Agricultural Research*, 8(29):4067–4073, 2013.

A. H. Auchincloss, R. L. Riolo, D. G. Brown, J. Cook, and A. V. Diez Roux. An agent-based model of income inequalities in diet in the context of residential segregation. *American journal of preventive medicine*, 40(3):303–311, Mar. 2011.

J. Augusiak, P. J. Van den Brink, and V. Grimm. Merging validation and evaluation of ecological models to 'evaludation': A review of terminology and a practical approach. *Ecological modelling*, 280:117–128, May 2014.

R. Axtell, R. Axelrod, J. M. Epstein, and M. D. Cohen. Aligning simulation models: A case study and results. *Computational & Mathematical Organization Theory*, 1(2):123–141, 1996.

R. L. Axtell and J. M. Epstein. Agent-based modeling: Understanding our creations. *The Bulletin of the Santa Fe Institute*, 9(2):28–32, 1994.

F. N. Bachewe, G. Berhane, B. Minten, and A. S. Taffesse. *Non-farm income and labor markets in rural Ethiopia*. Intl Food Policy Res Inst, June 2016.

C. Badgley, J. Moghtader, E. Quintero, E. Zakem, M. Jahi Chappell, K. Avilés-Vázquez, A. Samulon, and I. Perfecto. Organic agriculture and the global food supply. *Renewable Agriculture and Food Systems*, 22(2):86–108, June 2007.

J. Baffes, V. Kshirsagar, and D. Mitchell. What Drives Local Food Prices? Evidence from the Tanzanian Maize Market. *The World Bank economic review*, 2017.

C. Barnaud and A. Van Paassen. Equity, power games, and legitimacy: dilemmas of participatory natural resource management. *Ecology and Society*, 18(2), 2013.

J. Barnett and S. O'Neill. Maladaptation. *Global environmental change: human and policy dimensions*, 20(2):211–213, May 2010.

H. Baroud, J. E. Ramirez-Marquez, K. Barker, and C. M. Rocco. Stochastic measures of network resilience: applications to waterway commodity flows. *Risk analysis: an official publication of the Society for Risk Analysis*, 34(7):1317–1335, July 2014.

C. B. Barrett. Rural poverty dynamics: development policy implications. *Agricultural economics*, 32(s1):45–60, Jan. 2005.

C. B. Barrett. Measuring food insecurity. *Science*, 327(5967):825–828, Feb. 2010.

C. B. Barrett and L. E. M. Bevis. The self-reinforcing feedback between low soil fertility and chronic poverty. *Nature geoscience*, 8:907, Nov. 2015.

C. B. Barrett and M. R. Carter. The Economics of Poverty Traps and Persistent Poverty: Empirical and Policy Implications. *The journal of development studies*, 49(7):976–990, July 2013.

C. B. Barrett and M. A. Constas. Toward a theory of resilience for international development applications. *Proceedings of the National Academy of Sciences of the United States of America*, 111(40):14625–14630, Oct. 2014.

C. B. Barrett, T. Reardon, and P. Webb. Nonfarm income diversification and household livelihood strategies in rural Africa: concepts, dynamics, and policy implications. *Food policy*, 26(4): 315–331, 2001.

C. B. Barrett, B. J. Barnett, M. R. Carter, S. Chantarat, J. W. Hansen, A. G. Mude, D. Osgood, J. R. Skees, C. G. Turvey, and M. N. Ward. Poverty Traps and Climate Risk: Limitations and Opportunities of Index-Based Risk Financing. Technical Report 07-02, IRI, Sept. 2007.

C. B. Barrett, M. R. Carter, and C. P. Timmer. A Century-Long Perspective on Agricultural Development. *American journal of agricultural economics*, 92(2):447–468, Apr. 2010.

C. B. Barrett, M. E. Bachke, M. F. Bellemare, H. C. Michelson, S. Narayanan, and T. F. Walker. Smallholder Participation in Contract Farming: Comparative Evidence from Five Countries. *World development*, 40(4):715–730, Apr. 2012.

P. Baumgartner, J. von Braun, D. Abebaw, and M. Müller. Impacts of Large-scale Land Investments on Income, Prices, and Employment: Empirical Analyses in Ethiopia. *World development*, 72: 175–190, Aug. 2015.

S. Baumgärtner and M. F. Quaas. Managing increasing environmental risks through agrobiodiversity and agrienvironmental policies. *Agricultural economics*, 2010.

J. Bayala, G. W. Sileshi, R. Coe, A. Kalinganire, Z. Tchoundjeu, F. Sinclair, and D. Garrity. Cereal yield response to conservation agriculture practices in drylands of West Africa: A quantitative synthesis. *Journal of arid environments*, 78:13–25, Mar. 2012.

A. A. Beal Cohen, P. R. Cohen, and G. Kiker. Modeling Gender Inequity in Household Decision-Making. In *Social, Cultural, and Behavioral Modeling*, pages 31–38. Springer International Publishing, 2019.

M. W. Beck, O. Quast, and K. Pfliegner. Ecosystem-based Adaptation and Insurance: Success, Challenges and Opportunities. Technical report, InsuResilience Global Partnership, 2019.

S. Bediye, G. Assefa, A. Tedla, and D. Fekadu. Present status and future direction in feed resources and nutrition research targeted for wheat based crop-livestock production system in Ethiopia. CIMMYT, 2001.

A. R. Bell, P. S. Ward, and M. A. A. Shah. Increased water charges improve efficiency and equity in an irrigation system. *Ecology and Society*, 21(3), 2016.

M. F. Bellemare. Contract farming: opportunity cost and trade-offs. *Agricultural economics*, 49 (3):279–288, 2018.

M. F. Bellemare and C. B. Barrett. An Ordered Tobit Model of Market Participation: Evidence from Kenya and Ethiopia. *American journal of agricultural economics*, 88(2):324–337, 2006.

M. F. Bellemare and J. R. Bloem. Does contract farming improve welfare? A review. *World development*, 112:259–271, Dec. 2018.

M. F. Bellemare and S. Lim. In All Shapes and Colors: Varieties of Contract Farming. *Applied Economic Perspectives and Policy*, 40(3):379–401, Sept. 2018.

A. J. Bennett, G. D. Bending, D. Chandler, S. Hilton, and P. Mills. Meeting the demand for crop production: the challenge of yield decline in crops grown in short rotations. *Biological reviews of the Cambridge Philosophical Society*, 87(1):52–71, Feb. 2012.

T. Berger and P. Schreinemachers. Creating agents and landscapes for multiagent systems from random samples. *Ecology and Society*, 11(2), 2006.

T. Berger and C. Troost. Agent-based Modelling of Climate Adaptation and Mitigation Options in Agriculture. *Journal of Agricultural Economics*, 65(2):323–348, 2014.

T. Berger, C. Troost, T. Wossen, E. Latynskiy, K. Tesfaye, and S. Gbegbelegbe. Can smallholder farmers adapt to climate variability, and how effective are policy interventions? Agent-based simulation results for Ethiopia. *Agricultural economics*, 48(6):693–706, Nov. 2017.

B. Bernard and A. Lux. How to feed the world sustainably: an overview of the discourse on agroecology and sustainable intensification. *Regional Environmental Change*, 2017.

D. Berre, T. Diarisso, N. Andrieu, C. Le Page, and M. Corbeels. Biomass flows in an agro-pastoral village in West-Africa: Who benefits from crop residue mulching? *Agricultural systems*, 187: 102981, Feb. 2021.

K. Beven. A manifesto for the equifinality thesis. *Journal of Hydrology*, 320(1):18–36, Mar. 2006.

K. Beven and J. Freer. Equifinality, data assimilation, and uncertainty estimation in mechanistic modelling of complex environmental systems using the GLUE methodology. *Journal of Hydrology*, 249(1):11–29, Aug. 2001.

F. Beyene and M. Muche. Determinants of food security among rural households of Central Ethiopia: an empirical analysis. *Quarterly Journal of International Agriculture*, 49(4):299, 2010.

P. Biommel, M.-P. Bonnet, E. Coudel, E. Haentjens, C. N. Kraus, G. Melo, S. Nasuti, and C. Le Page. Livelihoods of Local Communities in an Amazonian Floodplain Coping with Global Changes: From Role-Playing Games to Hybrid Simulations to Involve Local Stakeholders in Participatory Foresight Study at Territorial Level. In *International Congress on Environmental Modelling and Software*, 2016.

P. Bitterman and D. A. Bennett. Leveraging Coupled Agent-Based Models to Explore the Resilience of Tightly-Coupled Land Use Systems. In *Agent-Based Models and Complexity Science in the Age of Geospatial Big Data*, pages 17–30. Springer International Publishing, 2018.

H. Blanco-Canqui, M. M. Claassen, and D. R. Presley. Summer Cover Crops Fix Nitrogen, Increase Crop Yield, and Improve Soil–Crop Relationships. *Agronomy journal*, 104:137–147, 2012.

S. Blazkova and K. Beven. A limits of acceptability approach to model evaluation and uncertainty estimation in flood frequency estimation by continuous simulation: Skalka catchment, Czech Republic. *Water resources research*, 45(12), 2009.

P. J. Block, K. Strzepek, M. W. Rosegrant, and X. Diao. Impacts of considering climate variability on investment decisions in Ethiopia. *Agricultural economics*, 39(2):171–181, Sept. 2008.

A. Bogale, K. Hagedorn, and B. Korf. Determinants of poverty in rural Ethiopia. *Quarterly Journal of International Agriculture*, 44(2):101–120, 2005.

S. Bogale, S. Melaku, and A. Yami. Matching Livestock Systems with Available Feed Resources in the Bale Highlands of Ethiopia. *Outlook on agriculture*, 37(2):105–110, June 2008.

R. Bommarco, D. Kleijn, and S. G. Potts. Ecological intensification: harnessing ecosystem services for food security. *Trends in ecology & evolution*, 28(4):230–238, Apr. 2013.

S. Borras and J. Franco. Towards a broader view of the politics of global land grab: rethinking land issues, reframing resistance. *Initiatives in Critical Agrarian Studies Working Paper Series*, 1:1–39, 2010.

P. Bottazzi, D. Crespo, L. O. Bangura, and S. Rist. Evaluating the livelihood impacts of a large-scale agricultural investment: Lessons from the case of a biofuel production company in northern Sierra Leone. *Land use policy*, 73:128–137, Apr. 2018.

E. D. Brill, S.-Y. Chang, and L. D. Hopkins. Modeling to Generate Alternatives: The HSJ Approach and an Illustration Using a Problem in Land Use Planning. *Management science*, 28(3): 221–235, Mar. 1982.

N. Brisson, C. Gary, E. Justes, R. Roche, B. Mary, D. Ripoche, D. Zimmer, J. Sierra, P. Bertuzzi, P. Burger, F. Bussière, Y. M. Cabidoche, P. Cellier, P. Debaeke, J. P. Gaudillère, C. Hénault, F. Maraux, B. Seguin, and H. Sinoquet. An overview of the crop model stics. *European journal of agronomy: the journal of the European Society for Agronomy*, 18(3):309–332, Jan. 2003.

D. G. Brown, S. Page, R. Riolo, M. Zellner, and W. Rand. Path dependence and the validation of agent-based spatial models of land use. *International journal of geographical information science: IJGIS*, 19(2):153–174, 2005.

D. G. Brown, P. H. Verburg, R. G. Pontius, and M. D. Lange. Opportunities to improve impact, integration, and evaluation of land change models. *Current Opinion in Environmental Sustainability*, 5(5):452–457, Oct. 2013.

M. E. Brown, E. R. Carr, K. L. Grace, K. Wiebe, C. C. Funk, W. Attavanich, P. Backlund, and L. Buja. Do markets and trade help or hurt the global food system adapt to climate change? *Food policy*, 68:154–159, Apr. 2017.

E. Bryan, T. T. Deressa, G. A. Gbetibouo, and C. Ringler. Adaptation to climate change in Ethiopia and South Africa: options and constraints. *Environmental science & policy*, 12(4):413–426, 2009.

B. P. Bryant and R. J. Lempert. Thinking inside the box: A participatory, computer-assisted approach to scenario discovery. *Technological forecasting and social change*, 77(1):34–49, Jan. 2010.

C. Béné, R. G. Wood, A. Newsham, and M. Davies. Resilience: New Utopia or New Tyranny? Reflection about the Potentials and Limits of the Concept of Resilience in Relation to Vulnerability Reduction Programmes. *IDS Working Papers*, 2012(405):1–61, Sept. 2012.

L. Büchi, C.-A. Gebhard, F. Liebisch, S. Sinaj, H. Ramseier, and R. Charles. Accumulation of biologically fixed nitrogen by legumes cultivated as cover crops in Switzerland. *Plant and soil*, 393(1):163–175, Aug. 2015.

CAADP. Nutrition country paper, Ethiopia. In *East and Central Africa Regional CAADP Nutrition Program Development Workshop*, 2013.

J. F. Cabell and M. Oelofse. An Indicator Framework for Assessing Agroecosystem Resilience. *Ecology and Society*, 17(1), 2012.

B. Calvez and G. Hutzler. Automatic Tuning of Agent-Based Models Using Genetic Algorithms. In *Multi-Agent-Based Simulation VI*, pages 41–57. Springer Berlin Heidelberg, 2006.

H. E. Campbell, Y. Kim, and A. M. Eckerd. *Rethinking Environmental Justice in Sustainable Cities: Insights from Agent-Based Modeling*. Routledge, May 2015.

A. Cardaci. Inequality, household debt and financial instability: An agent-based perspective. *Journal of economic behavior & organization*, 149:434–458, May 2018.

S. Carpenter, B. Walker, J. M. Anderies, and N. Abel. From metaphor to measurement: resilience of what to what? *Ecosystems*, 4(8):765–781, 2001.

M. R. Carter, S. A. Janzen, and Q. Stoeffler. Can insurance help manage climate risk and food insecurity? Evidence from the pastoral regions of East Africa. In *Climate Smart Agriculture*, pages 201–225. Springer, Cham, 2018.

W. Chamberlain and W. Anseeuw. Inclusive businesses in agriculture: Defining the concept and its complex and evolving partnership structures in the field. *Land use policy*, 83:308–322, Apr. 2019.

P. P. Chandanabhumma and S. Narasimhan. Towards health equity and social justice: an applied framework of decolonization in health promotion. *Health promotion international*, 35(4):831–840, Aug. 2020.

S. Chantarat, A. G. Mude, C. B. Barrett, and C. G. Turvey. Welfare Impacts of Index Insurance in the Presence of a Poverty Trap. *World development*, 94:119–138, June 2017.

A. V. Chayanov. *A.V. Chayanov on the Theory of Peasant Economy*. Manchester University Press, 1986.

M. Chica, J. Barranquero, T. Kajdanowicz, S. Damas, and Cordón. Multimodal optimization: An effective framework for model calibration. *Information sciences*, 375:79–97, Jan. 2017.

T. Choi and P. J. Robertson. Deliberation and Decision in Collaborative Governance: A Simulation of Approaches to Mitigate Power Imbalance. *Journal of Public Administration Research and Theory*, 24(2):495–518, Feb. 2013.

E. K. Chu and C. E. B. Cannon. Equity, inclusion, and justice as criteria for decision-making on climate adaptation in cities. *Current Opinion in Environmental Sustainability*, 51:85–94, Aug. 2021.

W. C. Clark and A. G. Harley. Sustainability Science: Toward a Synthesis. *Annual review of environment and resources*, 45(1):331–386, Oct. 2020.

D. C. Clay, D. Molla, and D. Habtewold. Food aid targeting in Ethiopia: A study of who needs it and who gets it. *Food policy*, 24(4):391–409, Aug. 1999.

J. G. Cobo, G. Dercon, and G. Cadisch. Nutrient balances in African land use systems across different spatial scales: A review of approaches, challenges and progress. *Agriculture, ecosystems & environment*, 136(1):1–15, Feb. 2010.

S. Cole, X. Giné, and J. Vickery. How Does Risk Management Influence Production Decisions? Evidence from a Field Experiment. *The review of financial studies*, 30(6):1935–1970, June 2017.

R. Conte and M. Paolucci. On agent-based modeling and computational social science. *Frontiers in psychology*, 5:668, July 2014.

S. Corbett-Davies, E. Pierson, A. Feller, S. Goel, and A. Huq. Algorithmic Decision Making and the Cost of Fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '17, pages 797–806, New York, NY, USA, Aug. 2017. Association for Computing Machinery.

A. Couëdel, L. Alletto, H. Tribouillois, and Justes. Cover crop crucifer-legume mixtures provide effective nitrate catch crop and nitrogen green manure ecosystem services. *Agriculture, ecosystems & environment*, 254:50–59, Feb. 2018.

R. Cretney. Resilience for whom? Emerging critical geographies of socio-ecological resilience. *Geography Compass*, 8(9):627–640, 2014.

E. Cruz Cortés and D. Ghosh. An Invitation to System-wide Algorithmic Fairness. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, AIES '20, pages 235–241, New York, NY, USA, Feb. 2020. Association for Computing Machinery.

A.-S. Crépin and T. Lindahl. Grazing Games: Sharing Common Property Resources with Complex Dynamics. *Environmental & Resource Economics*, 44(1):29–46, Sept. 2009.

CSA. LSMS—Integrated Surveys on Agriculture Ethiopia Socioeconomic Survey (ESS), 2017.

S. L. Cutter. Resilience to what? Resilience for whom? *The Geographical journal*, 182(2): 110–113, 2016.

M. Dainese, E. A. Martin, M. A. Aizen, M. Albrecht, I. Bartomeus, R. Bommarco, L. G. Carvalheiro, R. Chaplin-Kramer, V. Gagic, L. A. Garibaldi, J. Ghazoul, H. Grab, M. Jonsson, D. S. Karp, C. M. Kennedy, D. Kleijn, C. Kremen, D. A. Landis, D. K. Letourneau, L. Marini, K. Poveda, R. Rader, H. G. Smith, T. Tscharntke, G. K. S. Andersson, I. Badenhausser, S. Baensch, A. D. M. Bezerra, F. J. J. A. Bianchi, V. Boreux, V. Bretagnolle, B. Caballero-Lopez, P. Cavigliasso, A. Ćetković, N. P. Chacoff, A. Classen, S. Cusser, F. D. da Silva E Silva, G. A. de Groot, J. H. Dudenhöffer, J. Ekroos, T. Fijen, P. Franck, B. M. Freitas, M. P. D. Garratt, C. Gratton, J. Hipólito, A. Holzschuh, L. Hunt, A. L. Iverson, S. Jha, T. Keasar, T. N. Kim, M. Kishinevsky, B. K. Klatt, A.-M. Klein, K. M. Krewenka, S. Krishnan, A. E. Larsen, C. Lavigne, H. Liere, B. Maas, R. E. Mallinger, E. Martinez Pachon, A. Martínez-Salinas, T. D. Meehan, M. G. E. Mitchell, G. A. R. Molina, M. Nesper, L. Nilsson, M. E. O'Rourke, M. K. Peters, M. Plećaš, S. G. Potts, D. d. L. Ramos, J. A. Rosenheim, M. Rundlöf, A. Rusch, A. Sáez, J. Scheper, M. Schleuning, J. M. Schmack, A. R. Sciligo, C. Seymour, D. A. Stanley, R. Stewart, J. C. Stout, L. Sutter, M. B. Takada, H. Taki, G. Tamburini, M. Tschumi, B. F. Viana, C. Westphal, B. K. Willcox, S. D. Wratten, A. Yoshioka, C. Zaragoza-Trello, W. Zhang, Y. Zou, and I. Steffan-Dewenter. A global synthesis reveals biodiversity-mediated benefits for crop production. *Science advances*, 5(10):eaax0121, Oct. 2019.

O. De Schutter. The green rush: the global race for farmland and the rights of land users. *Harv. Int'l LJ*, 52:503, 2011a.

O. De Schutter. How not to think of land-grabbing: three critiques of large-scale investments in farmland. *The Journal of peasant studies*, 38(2):249–279, Mar. 2011b.

A. de Vos, R. Biggs, and R. Preiser. Methods for understanding social-ecological systems: a review of place-based studies. *Ecology and Society*, 24(4), Nov. 2019.

T. E. de Wildt, E. J. L. Chappin, G. van de Kaa, P. M. Herder, and I. R. van de Poel. Conflicted by decarbonisation: Five types of conflict at the nexus of capabilities and decentralised energy systems identified with an agent-based model. *Energy Research & Social Science*, 64:101451, June 2020.

K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, Apr. 2002.

N. Debonne, J. van Vliet, A. Heinimann, and P. Verburg. Representing large-scale land acquisitions in land use change scenarios for the Lao PDR. *Regional Environmental Change*, 18(6):1857–1869, Aug. 2018.

N. Debonne, J. van Vliet, G. Metternicht, and P. Verburg. Agency shifts in agricultural land governance and their implications for land degradation neutrality. *Global environmental change: human and policy dimensions*, 66:102221, Jan. 2021.

J. DeCarolis, H. Daly, P. Dodds, I. Keppo, F. Li, W. McDowall, S. Pye, N. Strachan, E. Trutnevyte, W. Usher, M. Winning, S. Yeh, and M. Zeyringer. Formalizing best practice for energy system optimization modelling. *Applied energy*, 194:184–198, May 2017.

K. Deininger and D. Byerlee. *Rising Global Interest in Farmland: Can It Yield Sustainable and Equitable Benefits?* World Bank Publications, Jan. 2011.

J. Dell'Angelo, P. D'Odorico, and M. C. Rulli. Threats to sustainable development posed by land and water grabbing. *Current Opinion in Environmental Sustainability*, 26-27:120–128, June 2017.

M. S. DeLonge, A. Miles, and L. Carlisle. Investing in the transition to sustainable agriculture. *Environmental science & policy*, 55:266–273, Jan. 2016.

A. B. Demeke, A. Keil, and M. Zeller. Using panel data to estimate the effect of rainfall shocks on smallholders food security and vulnerability in rural Ethiopia. *Climatic change*, 108(1): 185–206, Sept. 2011.

A. Demissie and B. Legesse. Determinants of income diversification among rural households: The case of smallholder farmers in Fedis district, Eastern hararghe zone, Ethiopia. *Journal of Development and Agricultural Economics*, 5(3):120–128, 2013.

S. Dercon. Growth and shocks: evidence from rural Ethiopia. *Journal of development economics*, 74(2):309–329, 2004.

S. Dercon and L. Christiaensen. Consumption risk, technology adoption and poverty traps: Evidence from Ethiopia. *Journal of development economics*, 96(2):159–173, Nov. 2011.

S. Desta and D. L. Coppock. Cattle Population Dynamics in the Southern Ethiopian Rangelands, 1980-97. *Journal of Range Management*, 55(5):439–451, 2002.

S. Devereux and I. Sussex. Food insecurity in Ethiopia. *a DFID Ethiopia Seminar, London*, 2000.

H. J. Di and K. C. Cameron. Nitrate leaching in temperate agroecosystems: sources, factors and mitigating strategies. *Nutrient Cycling in Agroecosystems*, 64(3):237–256, Nov. 2002.

A. Dorward. Integrating Contested Aspirations, Processes and Policy: Development as Hanging In, Stepping Up and Stepping Out. *Development policy review: the journal of the Overseas Development Institute*, 27(2):131–146, Mar. 2009.

A. Dorward, J. Kydd, J. Morrison, and I. Urey. A Policy Agenda for Pro-Poor Agricultural Growth. *World development*, 32(1):73–89, Jan. 2004.

Y. Dou, P. Deadman, M. Berbés-Blázquez, N. Vogt, and O. Almeida. Pathways out of poverty through the lens of development resilience: an agent-based simulation. *Ecology and Society*, 25 (4), Oct. 2020.

G. Dressler, J. Groeneveld, C. M. Buchmann, C. Guo, N. Hase, J. Thober, K. Frank, and B. Müller. Implications of behavioral change for the resilience of pastoral systems—Lessons from an agent-based model. *Ecological Complexity*, 40:100710, Dec. 2019a.

G. Dressler, F. Hoffmann, I. Breuer, D. Kreuer, M. Mahdi, K. Frank, and B. Müller. Polarization in (post)nomadic resource use in Eastern Morocco: insights using a multi-agent simulation model. *Regional Environmental Change*, 19(2):489–500, Feb. 2019b.

L. E. Drinkwater, P. Wagoner, and M. Sarrantonio. Legume-based cropping systems have reduced carbon and nitrogen losses. *Nature*, 396(6708):262–265, Nov. 1998.

H. Duan, G. Zhang, S. Wang, and Y. Fan. Robust climate change research: a review on multi-model analysis. *Environmental Research Letters*, 14(3):033001, Feb. 2019.

A. J. Duff, P. H. Zedler, J. A. Barzen, and D. L. Knuteson. The capacity-building stewardship model: assessment of an agricultural network as a mechanism for improving regional agroecosystem sustainability. *Ecology and Society*, 22(1), 2017.

T. Dávid-Barrett and R. I. M. Dunbar. Social elites can emerge naturally when interaction in networks is restricted. *Behavioral ecology: official journal of the International Society for Behavioral Ecology*, 25(1):58–68, Oct. 2013.

A. Eckerd, Y. Kim, and H. E. Campbell. Community privilege and environmental justice: An agent-based analysis: Community privilege and environmental justice. *The review of policy research*, 34(2):144–167, Mar. 2017.

A. Efstratiadis and D. Koutsoyiannis. One decade of multi-objective calibration approaches in hydrological modelling: a review. *Hydrological Sciences Journal*, 55(1):58–78, Mar. 2010.

M. Eggen. *Climate Change Impacts to Agriculture of the Blue Nile Highlands, Ethiopia*. The University of Wisconsin-Madison, 2018.

L. Egli, H. Weise, V. Radchuk, R. Seppelt, and V. Grimm. Exploring resilience with agent-based models: State of the art, knowledge gaps and recommendations for coping with multidimensionality. *Ecological Complexity*, 40(B), Aug. 2018.

L. Ekblad and J. D. Herman. Toward data-driven generation and evaluation of model structure for integrated representations of human behavior in water resources systems. *Water resources research*, 57(2), Feb. 2021.

E. Elias, S. Morse, and D. G. R. Belshaw. Nitrogen and phosphorus balances of Kindo Koisha farms in southern Ethiopia. *Agriculture, ecosystems & environment*, 71(1):93–113, Dec. 1998.

J. Elith, J. R. Leathwick, and T. Hastie. A working guide to boosted regression trees. *The Journal of animal ecology*, 77(4):802–813, July 2008.

J. Emmerling and M. Tavoni. Representing inequalities in integrated assessment modeling of climate change. *One Earth*, 4(2):177–180, Feb. 2021.

J. M. Epstein and R. Axtell. *Growing Artificial Societies: Social Science from the Bottom Up*. Brookings Institution Press, Oct. 1996.

T. P. Evans, K. Phanvilay, J. Fox, and J. Vogler. An agent-based model of agricultural innovation, land-cover change and household inequality: the transition from swidden cultivation to rubber plantations in Laos PDR. *Journal of land use science*, 6(2-3):151–173, 2011.

S. U. Ewansiha and B. B. Singh. Relative drought tolerance of important herbaceous legumes and cereals in the moist and semi-arid regions of West Africa. *Journal of Food Agriculture and Environment*, 4(2):188–190, 2006.

FAO. Yield Response to Water. FAO Irrigation and Drainage Paper No. 33. Technical Report 33, Food and Agriculture Organization of the United Nations, 1984.

FAO. Crop Evapo-transpiration: Guidelines for computing crop water requirements. FAO Irrigation and Drainage Paper No. 56. Technical Report 56, Food and Agriculture Organization of the United Nations, 1998.

FAO. An Introduction to the Basic Concepts of Food Security. http://www.fao.org/3/al936e/al936e00.pdf, 2008.

FAO. FAO's Work on Agroecology: A pathway to achieving the SDGs. Technical report, Food and Agriculture Organization, 2018.

S. Farhadi, M. R. Nikoo, G. R. Rakhshandehroo, M. Akhbari, and M. R. Alizadeh. An agent-based-nash modeling framework for sustainable groundwater management: A case study. *Agricultural water management*, 177:348–358, Nov. 2016.

Federal Democratic Republic of Ethiopia. Ethiopia's Climate Resilient Green Economy: National Adaptation Plan. Technical report, 2019.

M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. Certifying and Removing Disparate Impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '15, pages 259–268, New York, NY, USA, Aug. 2015. Association for Computing Machinery.

P. J. Ferraro and M. M. Hanauer. Quantifying causal mechanisms to determine how protected areas affect poverty through changes in ecosystem services and infrastructure. *Proceedings of the National Academy of Sciences of the United States of America*, 111(11):4332–4337, Mar. 2014.

I. E. P. Ferreira, S. S. Zocchi, and D. Baron. Reconciling the Mitscherlich's law of diminishing returns with Liebig's law of the minimum. Some results on crop modeling. *Mathematical biosciences*, 293:29–37, Nov. 2017.

T. Filatova, A. Voinov, and A. van der Veen. Land market mechanisms for preservation of space for coastal ecosystems: An agent-based analysis. *Environmental Modelling & Software*, 26(2): 179–190, Feb. 2011.

T. Filatova, J. G. Polhill, and S. van Ewijk. Regime shifts in coupled socio-environmental systems: Review of modelling challenges and approaches. *Environmental Modelling & Software*, 75: 333–347, Jan. 2016.

R. Finger and N. Buchmann. An ecological economic assessment of risk-reducing effects of species diversity in managed grasslands. *Ecological economics: the journal of the International Society for Ecological Economics*, 110:89–97, Feb. 2015.

E. Fisher, J. Hellin, H. Greatrex, and N. Jensen. Index insurance and climate risk management: Addressing social equity. *Development policy review: the journal of the Overseas Development Institute*, 53:932, May 2019.

M. A. Florentin, M. Peñalva, A. Calegari, and R. Derpsch. Green manure/cover crops and crop rotation in conservation agriculture on small farms. Technical Report Vol. 12-2010, FAO, 2011.

C. Folke. Resilience (Republished). *Ecology and Society*, 21(4), 2016.

J. C. Franco and S. M. Borras. Grey areas in green grabbing: subtle and indirect interconnections between climate change politics and land grabs and their implications for research. *Land use policy*, 84:192–199, May 2019.

N. Fraser. From redistribution to recognition? Dilemmas of justice in a "postsocialist" age. *New Left review*, 1995.

N. Fraser. *Scales of justice: Reimagining political space in a globalizing world*. Columbia University Press, New York, 2009.

R. Frelat, S. Lopez-Ridaura, K. E. Giller, M. Herrero, S. Douxchamps, A. Andersson Djurfeldt, O. Erenstein, B. Henderson, M. Kassie, B. K. Paul, C. Rigolot, R. S. Ritzema, D. Rodriguez, P. J. A. van Asten, and M. T. van Wijk. Drivers of household food availability in sub-Saharan Africa based on big data from small farms. *Proceedings of the National Academy of Sciences of the United States of America*, 113(2):458–463, Jan. 2016.

C. C. Funk, P. J. Peterson, M. F. Landsfeld, D. H. Pedreros, J. P. Verdin, J. D. Rowland, B. E. Romero, G. J. Husak, J. C. Michaelsen, and A. P. Verdin. A quasi-global precipitation time series for drought monitoring. *US Geological Survey Data Series*, 832(4), 2014.

J. Gaffney, M. Challender, K. Califf, and K. Harden. Building bridges between agribusiness innovation and smallholder farmers: A review. *Global Food Security*, 20:60–65, Mar. 2019.

F. W. Geels. Technological transitions as evolutionary reconfiguration processes: a multi-level perspective and a case-study. *Research policy*, 31(8):1257–1274, Dec. 2002.

K. E. Giller, G. Cadisch, C. Ehaliotis, E. Adams, W. D. Sakala, and P. L. Mafongoya. Building Soil Nitrogen Capital in Africa. In *Replenishing Soil Fertility in Africa*, SSSA Special Publication, pages 151–192. Soil Science Society of America and American Society of Agronomy, Madison, WI, 1997.

K. E. Giller, P. Tittonell, M. C. Rufino, M. T. van Wijk, S. Zingore, P. Mapfumo, S. Adjei-Nsiah, M. Herrero, R. Chikowo, M. Corbeels, E. C. Rowe, F. Baijukya, A. Mwijage, J. Smith, E. Yeboah, W. J. van der Burg, O. M. Sanogo, M. Misiko, N. de Ridder, S. Karanja, C. Kaizzi, J. K'ungu, M. Mwale, D. Nwaga, C. Pacini, and B. Vanlauwe. Communicating complexity: Integrated assessment of trade-offs concerning soil fertility management within African farming systems to support innovation and development. *Agricultural systems*, 104(2):191–203, Feb. 2011.

D. E. Goldberg and J. Richardson. Genetic algorithms with sharing for multimodal function optimization. In *Genetic algorithms and their applications: Proceedings of the Second International Conference on Genetic Algorithms*, pages 41–49. books.google.com, 1987.

S. M. Goodreau, E. S. Rosenberg, S. M. Jenness, N. Luisi, S. E. Stansfield, G. A. Millett, and P. S. Sullivan. Sources of racial disparities in HIV prevalence in men who have sex with men in Atlanta, GA, USA: a modelling study. *The Lancet HIV*, 4(7):e311–e320, July 2017.

S. Grieshaber. Equity and research design. *Doing early childhood research: International perspectives on theory and practice*, pages 177–191, 2010.

M. Grillot, F. Guerrin, B. Gaudou, D. Masse, and J. Vayssières. Multi-level analysis of nutrient cycling within agro-sylvo-pastoral landscapes in West Africa using an agent-based model. *Environmental Modelling & Software*, 107:267–280, Sept. 2018.

V. Grimm and S. F. Railsback. *Individual-based Modeling and Ecology*. Princeton University Press, 2005.

V. Grimm and S. F. Railsback. Pattern-oriented modelling: a 'multi-scope'for predictive systems ecology. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 367(1586):298–310, 2012.

V. Grimm, E. Revilla, U. Berger, F. Jeltsch, W. M. Mooij, S. F. Railsback, H.-H. Thulke, J. Weiner, T. Wiegand, and D. L. DeAngelis. Pattern-oriented modeling of agent-based complex systems: lessons from ecology. *Science*, 310(5750):987–991, Nov. 2005.

V. Grimm, U. Berger, F. Bastiansen, S. Eliassen, V. Ginot, J. Giske, J. Goss-Custard, T. Grand, S. K. Heinz, G. Huse, A. Huth, J. U. Jepsen, C. Jørgensen, W. M. Mooij, B. Müller, G. Pe'er, C. Piou, S. F. Railsback, A. M. Robbins, M. M. Robbins, E. Rossmanith, N. Rüger, E. Strand, S. Souissi, R. A. Stillman, R. Vabø, U. Visser, and D. L. DeAngelis. A standard protocol for describing individual-based and agent-based models. *Ecological modelling*, 198(1–2):115–126, Sept. 2006.

V. Grimm, J. Augusiak, A. Focks, B. M. Frank, F. Gabsi, A. S. A. Johnston, C. Liu, B. T. Martin, M. Meli, V. Radchuk, P. Thorbek, and S. F. Railsback. Towards better modelling and decision support: Documenting model development, testing, and analysis using TRACE. *Ecological modelling*, 280:129–139, May 2014.

V. Grimm, S. F. Railsback, C. E. Vincenot, U. Berger, C. Gallagher, D. L. Deangelis, B. Edmonds, J. Ge, J. Giske, J. Groeneveld, A. S. A. Johnston, A. Milles, J. Nabe-Nielsen, J. G. Polhill, V. Radchuk, M. S. Rohwäder, R. A. Stillman, J. C. Thiele, and D. Ayllón. The ODD protocol for describing agent-based and other simulation models: A second update to improve clarity, replication, and structural realism. *Journal of Artificial Societies and Social Simulation*, 23(2), Jan. 2020.

J. Groeneveld, B. Müller, C. M. Buchmann, G. Dressler, C. Guo, N. Hase, F. Hoffmann, F. John, C. Klassert, T. Lauf, V. Liebelt, H. Nolzen, N. Pannicke, J. Schulze, H. Weise, and N. Schwarz. Theoretical foundations of human decision-making in agent-based land use models – A review. *Environmental Modelling & Software*, 87:39–48, Jan. 2017.

J. H. A. Guillaume, J. D. Jakeman, S. Marsili-Libelli, M. Asher, P. Brunner, B. Croke, M. C. Hill, A. J. Jakeman, K. J. Keesman, S. Razavi, and J. D. Stigter. Introductory overview of identifiability analysis: A guide to evaluating whether you have the right type of data for your modeling purpose. *Environmental Modelling & Software*, 119:418–432, Sept. 2019.

L. J. Haider, W. J. Boonstra, G. D. Peterson, and M. Schlüter. Traps and Sustainable Development in Rural Areas: A Review. *World development*, 101:311–321, Jan. 2018.

A. Hailegiorgis, A. Crooks, and C. Cioffi-Revilla. An Agent-Based Model of Rural Households' Adaptation to Climate Change. *Journal of Artificial Societies and Social Simulation*, 21(4):1–4, 2018.

A. B. Hailegiorgis and C. Cioffi-Revilla. Agent-Based Modeling of Large-Scale Land Acquisition and Rural Household Dynamics. In J.-C. Thill and S. Dragicevic, editors, *GeoComputational Analysis and Modeling of Regional Systems*, pages 101–119. Springer International Publishing, Cham, 2018.

R. Hajjar, A. N. Ayana, R. Rutt, O. Hinde, C. Liao, S. Keene, S. Bandiaky-Badji, and A. Agrawal. Capital, labor, and gender: the consequences of large-scale land transactions on household labor allocation. *The Journal of peasant studies*, pages 1–23, May 2019.

R. Hall, M. Edelman, S. M. Borras, Jr., I. Scoones, B. White, and W. Wolford. Resistance, acquiescence or incorporation? An introduction to land grabbing and political reactions 'from below'. *The Journal of peasant studies*, 42(3-4):467–488, July 2015.

J. Hansen, J. Hellin, T. Rosenstock, E. Fisher, J. Cairns, C. Stirling, C. Lamanna, J. van Etten, A. Rose, and B. Campbell. Climate risk management and rural poverty reduction. *Agricultural systems*, 172:28–46, June 2019.

J. W. Hansen, S. J. Mason, L. Sun, and A. Tall. Review of seasonal climate forecasting for agriculture in sub-Saharan Africa. *Experimental Agriculture*, 47(2):205–240, Apr. 2011.

K. W. Harrison, R. D. Dumas, E. Solano, M. A. Barlaz, E. D. Brill, Jr, and S. R. Ranjithan. Decision support tool for life-cycle-based solid waste management. *Journal of Computing in Civil Engineering*, 15(1):44–58, 2001.

F. Hartig, J. M. Calabrese, B. Reineking, T. Wiegand, and A. Huth. Statistical inference for stochastic simulation models–theory and application. *Ecology letters*, 14(8):816–827, Aug. 2011.

P. Hazell, J. Anderson, N. Balzer, A. Hastrup Clemmensen, U. Hess, and F. Rispoli. The Potential for Scale and Sustainability in Weather Index Insurance for Agriculture and Rural Livelihoods. Technical report, International Fund for Agricultural Development and World Food Programme., 2010a.

P. Hazell, C. Poulton, S. Wiggins, and A. Dorward. The Future of Small Farms: Trajectories and Policy Priorities. *World development*, 38(10):1349–1361, Oct. 2010b.

D. Headey, A. S. Taffesse, and L. You. Diversification and Development in Pastoralist Ethiopia. *World development*, 56:200–213, Apr. 2014.

D. D. Headey and O. Ecker. Improving the Measurement of Food Security. Technical Report IFPRI Discussion Paper 1225, International Food Policy Research Institute (IFPRI), Nov. 2012.

A. D. Henry and H. I. Brugger. Agent-Based Explorations of Environmental Consumption in Segregated Networks. *Social systems engineering*, page 199, 2018.

R. T. Herrmann. Large-Scale Agricultural Investments and Smallholder Welfare: A Comparison of Wage Labor and Outgrower Channels in Tanzania. *World development*, 90:294–310, Feb. 2017.

HLPE. Agroecological and other innovative approaches for sustainable agriculture and food systems that enhance food security and nutrition. A report by the High Level Panel of Experts on Food Security and Nutrition of the Committee on World Food Security, Rome. Technical report, 2019.

S. T. Holden, B. Shiferaw, and M. Wik. Poverty, market imperfections and time preferences: of relevance for environmental policy? *Environment and Development Economics*, 3(1):105–130, Feb. 1998.

C. S. Holling. Resilience and Stability of Ecological Systems. *Annual review of ecology and systematics*, 4(1):1–23, Nov. 1973.

A. G. D. Holmes. Researcher positionality – A consideration of its influence and place in qualitative research – A new researcher guide. *Shanlax International Journal of Education*, 8(4):1–10, Sept. 2020.

R. D. Horan, E. P. Fenichel, K. L. S. Drury, and D. M. Lodge. Managing ecological thresholds in coupled environmental–human systems. *Proceedings of the National Academy of Sciences of the United States of America*, 108(18):7333–7338, May 2011.

G. M. Hornberger and R. C. Spear. Approach to the preliminary analysis of environmental systems. *J. Environ. Mgmt.*, 12(1):7–18, 1981.

R. Huber, M. Bakker, A. Balmann, T. Berger, M. Bithell, C. Brown, A. Grêt-Regamey, H. Xiong, Q. B. Le, G. Mack, P. Meyfroidt, J. Millington, B. Müller, J. G. Polhill, Z. Sun, R. Seidl, C. Troost, and R. Finger. Representation of decision-making in European agricultural agent-based models. *Agricultural systems*, 167:143–160, Nov. 2018.

R. P. Hämäläinen, I. Miliszewska, and A. Voinov. Leadership in participatory modelling–Is there a need for it? *Environmental Modelling & Software*, 133:104834, 2020.

J. Ikeme. Equity, environmental justice and sustainability: incomplete approaches in climate change politics. *Global environmental change: human and policy dimensions*, 13(3):195–206, Oct. 2003.

B. Iooss and P. Lemaître. A Review on Global Sensitivity Analysis Methods. In G. Dellino and C. Meloni, editors, *Uncertainty Management in Simulation-Optimization of Complex Systems: Algorithms and Applications*, pages 101–122. Springer US, Boston, MA, 2015.

J. Jacobi, S. Mukhovi, A. Llanque, H. Augstburger, F. Käser, C. Pozo, M. Ngutu Peter, J. M. F. Delgado, B. P. Kiteme, S. Rist, and C. Ifejika Speranza. Operationalizing food system resilience: An indicator-based assessment in agroindustrial, smallholder farming, and agroecological contexts in Bolivia and Kenya. *Land use policy*, 79:433–446, Dec. 2018.

A. Z. Jacobs and H. Wallach. Measurement and fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 375–385. arxiv.org, 2021.

M. A. Janssen, B. H. Walker, J. Langridge, and N. Abel. An adaptive agent model for analysing co-evolution of management and policies in a complex rangeland system. *Ecological modelling*, 131(2):249–268, July 2000.

M. A. Janssen, J. M. Anderies, and B. H. Walker. Robust strategies for managing rangelands with multiple stable attractors. *Journal of environmental economics and management*, 47(1): 140–162, Jan. 2004.

M. A. Janssen, C. Pritchard, and A. Lee. On code sharing and model documentation of published individual and agent-based models. *Environmental Modelling and Software[R]*, 134:104873, Dec. 2020.

I. Jatoi, H. Becher, and C. R. Leake. Widening disparity in survival between white and African-American patients with breast carcinoma treated in the US Department of Defense Healthcare system. *Cancer: Interdisciplinary International Journal of the American Cancer Society*, 98(5): 894–899, 2003.

F. John, R. Toth, K. Frank, J. Groeneveld, and B. Müller. Ecological Vulnerability Through Insurance? Potential Unintended Consequences of Livestock Drought Insurance. *Ecological economics: the journal of the International Society for Ecological Economics*, 157:357–368, Mar. 2019.

J. W. Jones, G. Hoogenboom, C. H. Porter, K. J. Boote, W. D. Batchelor, L. A. Hunt, P. W. Wilkens, U. Singh, A. J. Gijsman, and J. T. Ritchie. The DSSAT cropping system model. *European journal of agronomy: the journal of the European Society for Agronomy*, 18(3):235–265, Jan. 2003.

H. Jumare, M. Visser, and K. Brick. Risk Preferences and the Poverty Trap: A Look at Technology Uptake amongst Smallholder Farmers in the Matzikama Municipality. *ERSA working paper 725*, 2018.

D. Kahneman and A. Tversky. Prospect Theory: An Analysis of Decision under Risk. *Econometrica: journal of the Econometric Society*, 47(2):263–291, 1979.

D. Kaimowitz and A. Angelsen. *Economic Models of Tropical Deforestation: A Review*. CIFOR, Jan. 1998.

M. K. Kansiime, P. van Asten, and K. Sneyers. Farm diversity and resource use efficiency: Targeting agricultural policy interventions in East Africa farming systems. *NJAS - Wageningen Journal of Life Sciences*, 85:32–41, Apr. 2018.

D. Karlan, R. Osei, I. Osei-Akoto, and C. Udry. Agricultural Decisions after Relaxing Credit and Risk Constraints. *The quarterly journal of economics*, 129(2):597–652, May 2014.

J. R. Kasprzyk, S. Nataraj, P. M. Reed, and R. J. Lempert. Many objective robust decision making for complex environmental systems undergoing change. *Environmental Modelling & Software*, 42:55–71, Apr. 2013.

B. E. Kaufman. A new theory of satisficing. *The Journal of behavioral economics*, 19(1):35–51, Mar. 1990.

J. P. Kaye and M. Quemada. Using cover crops to mitigate and adapt to climate change. A review. *Agronomy for Sustainable Development*, 37(1):4, Jan. 2017.

W. Kaye-Blake, C. Schilling, R. Monaghan, R. Vibart, S. Dennis, and E. Post. Quantification of environmental-economic trade-offs in nutrient management policies. *Agricultural systems*, 173: 458–468, July 2019.

B. A. Keating, P. S. Carberry, G. L. Hammer, M. E. Probert, M. J. Robertson, D. Holzworth, N. I. Huth, J. N. G. Hargreaves, H. Meinke, Z. Hochman, G. McLean, K. Verburg, V. Snow, J. P. Dimes, M. Silburn, E. Wang, S. Brown, K. L. Bristow, S. Asseng, S. Chapman, R. L. McCown, D. M. Freebairn, and C. J. Smith. An overview of APSIM, a model designed for farming systems simulation. *European journal of agronomy: the journal of the European Society for Agronomy*, 18(3):267–288, Jan. 2003.

D. Keftasa. Role of crop residues as livestock feed in the Ethiopian highlands. In *African Forage Plant Genetic Resources, Evaluation of Forage Germplasm and Extensive Livestock Production Systems. Proceedings of a Workshop Held in Arusha*, pages 430–439. cgspace.cgiar.org, 1988.

R. A. L. Kelly, A. J. Jakeman, O. Barreteau, M. E. Borsuk, S. ElSawah, S. H. Hamilton, H. J. Henriksen, S. Kuikka, H. R. Maier, A. E. Rizzoli, H. van Delden, and A. A. Voinov. Selecting among five common modelling approaches for integrated environmental assessment and management. *Environmental Modelling & Software*, 47:159–181, Sept. 2013.

S. Khatami, M. C. Peel, T. J. Peterson, and A. W. Western. Equifinality and Flux Mapping: A new approach to model evaluation and process representation under uncertainty. *Water resources research*, 55:2, 2019.

E. Kiesling, M. Günther, C. Stummer, and L. M. Wakolbinger. Agent-based simulation of innovation diffusion: a review. *Central European Journal of Operations Research*, 20(2):183–230, June 2012.

A. Klabunde and F. Willekens. Decision-Making in Agent-Based Models of Migration: State of the Art and Challenges. *European journal of population = Revue europeenne de demographie*, 32:73–97, Feb. 2016.

L. Kleemann and R. Thiele. Rural welfare implications of large-scale land acquisitions in Africa: A theoretical framework. *Economic modelling*, 51:269–279, Dec. 2015.

E. S. Klein, M. R. Barbier, and J. R. Watson. The dual impact of ecology and management on social incentives in marine common-pool resource systems. *Royal Society open science*, 4(8): 170740, Aug. 2017.

J. Kleinberg, S. Mullainathan, and M. Raghavan. Inherent Trade-Offs in the Fair Determination of Risk Scores. In *Proceedings of The 8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, Sept. 2016.

R. Komuro, E. D. Ford, and J. H. Reynolds. The use of multi-criteria assessment in developing a process model. *Ecological modelling*, 197(3):320–330, Aug. 2006.

J. Konow. Which Is the Fairest One of All? A Positive Analysis of Justice Theories. *Journal of economic literature*, 41(4):1188–1239, Dec. 2003.

B. Kramer and F. Ceballos. Enhancing adaptive capacity through climate-smart insurance: Theory and evidence from India. Technical report, 2018.

B. Kramer, J. Hellin, J. Hansen, A. Rose, and M. Braun. Building resilience through climate risk insurance: Insights from agricultural research for development. Technical Report 287, CGIAR, 2019.

D. Kremmydas, I. N. Athanasiadis, and S. Rozakis. A review of Agent Based Modeling for agricultural policy evaluation. *Agricultural systems*, 164:95–106, 2018.

S. V. Kumar, C. D. Peters-Lidard, Y. Tian, P. R. Houser, J. Geiger, S. Olden, L. Lighty, J. L. Eastman, B. Doty, P. Dirmeyer, J. Adams, K. Mitchell, E. F. Wood, and J. Sheffield. Land information system: An interoperable framework for high resolution land surface modeling. *Environmental Modelling & Software*, 21(10):1402–1415, Oct. 2006.

A. M. Kusow. Beyond indigenous authenticity: Reflections on the insider/outsider debate in immigration research. *Symbolic Interaction*, 26(4):591–599, Nov. 2003.

A. Kustov. How ethnic structure affects civil conflict: A model of endogenous grievance. *Conflict Management and Peace Science*, 34(6):660–679, Nov. 2017.

S. J. Lade, A. Tavoni, S. A. Levin, and M. Schlüter. Regime shifts in a social-ecological system. *Theoretical Ecology*, 6(3):359–372, 2013.

C. Lamine and S. Bellon. Conversion to organic farming: a multidimensional research object at the crossroads of agricultural and social sciences. A review. *Agronomy for Sustainable Development*, 29(1):97–112, Mar. 2009.

G. Latombe, L. Parrott, and D. Fortin. Levels of emergence in individual based models: Coping with scarcity of data and pattern redundancy. *Ecological modelling*, 222(9):1557–1568, May 2011.

A. M. Law. *Simulation modeling and analysis, Fourth Edition*, volume 4. McGraw-Hill New York, 2008.

P. R. Lawrence, J. T. Dijkman, and G. P. Jansen. The introduction of animal traction into inland valley regions. 1. Manual labour and animal traction in the cultivation of rice and maize: a comparison. *The Journal of agricultural science*, 129(1):65–70, Aug. 1997.

M. Leach, I. Scoones, and A. Stirling. *Dynamic Sustainabilities: Technology, Environment, Social Justice*. Routledge, 2010.

M. Leach, B. Reyers, X. Bai, E. S. Brondizio, C. Cook, S. Díaz, G. Espindola, M. Scobie, M. Stafford-Smith, and S. M. Subramanian. Equity and sustainability in the Anthropocene: a social–ecological systems perspective on their intertwined futures. *Global Sustainability*, 1, 2018.

D. Leclère, P. Havlík, S. Fuss, E. Schmid, A. Mosnier, B. Walsh, H. Valin, M. Herrero, N. Khabarov, and M. Obersteiner. Climate change induced transformations of agricultural systems: insights from a global model. *Environmental Research Letters*, 9(12):124018, Dec. 2014.

J.-S. Lee, T. Filatova, A. Ligmann-Zielinska, B. Hassani-Mahmooei, F. Stonedahl, I. Lorscheid, A. Voinov, J. G. Polhill, Z. Sun, and D. C. Parker. The complexities of agent-based modeling output analysis. *Journal of Artificial Societies and Social Simulation*, 18(4):4, 2015.

M. C. Lemos, J. C. Arnott, N. M. Ardoin, K. Baja, A. T. Bednarek, A. Dewulf, C. Fieseler, K. A. Goodrich, K. Jagannathan, N. Klenk, K. J. Mach, A. M. Meadow, R. Meyer, R. Moss, L. Nichols, K. D. Sjostrom, M. Stults, E. Turnhout, C. Vaughan, G. Wong-Parodi, and C. Wyborn. To co-produce or not to co-produce. *Nature sustainability*, 1(12):722–724, Dec. 2018.

R. J. Lempert. A new decision sciences for complex systems. *Proceedings of the National Academy of Sciences of the United States of America*, 99 Suppl 3:7309–7313, May 2002.

R. J. Lempert. *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*. Rand Corporation, July 2003.

J. L. Leroy, M. Ruel, E. A. Frongillo, J. Harris, and T. J. Ballard. Measuring the Food Access Dimension of Food Security: A Critical Review and Mapping of Indicators. *Food and nutrition bulletin*, 36(2):167–195, June 2015.

J. Levine, K. M. A. Chan, and T. Satterfield. From rational actor to efficient complexity manager: Exorcising the ghost of Homo economicus with a unified synthesis of cognition research. *Ecological economics: the journal of the International Society for Ecological Economics*, 114 (Supplement C):22–32, June 2015.

A. Ligmann-Zielinska, D. B. Kramer, K. Spence Cheruvelil, and P. A. Soranno. Using uncertainty and sensitivity analyses in socioecological agent-based models to improve their analytical performance and policy relevance. *PloS one*, 9(10):e109779, Oct. 2014.

A. Ligmann-Zielinska, R. L. Church, and P. Jankowski. Spatial optimization as a generative technique for sustainable multiobjective land-use allocation. *International journal of geographical information science: IJGIS*, 22(6):601–622, June 2008.

B. B. Lin. Resilience in Agriculture through Crop Diversification: Adaptive Management for Environmental Change. *Bioscience*, 61(3):183–193, Mar. 2011.

Y. S. Lincoln, S. A. Lynham, and E. G. Guba. Paradigmatic controversies, contradictions, and emerging confluences, revisited. *The Sage handbook of qualitative research*, 4:97–128, 2011.

E. Lindkvist, N. Wijermans, T. M. Daw, B. Gonzalez-Mon, A. Giron-Nava, A. F. Johnson, I. van Putten, X. Basurto, and M. Schlüter. Navigating Complexities: Agent-Based Modeling to Support Research, Governance, and Management in Small-Scale Fisheries. *Frontiers in Marine Science*, 6:733, 2020.

G. E. Liston and K. Elder. A Meteorological Distribution System for High-Resolution Terrestrial Modeling (MicroMet). *Journal of Hydrometeorology*, 7(2):217–234, Apr. 2006.

P. D. Little and M. Watts. *Living Under Contract: Contract Farming and Agrarian Transformation in Sub-Saharan Africa*. Univ of Wisconsin Press, 1994.

P. D. Little, M. P. Stone, T. Mogues, A. P. Castro, and W. Negatu. 'Moving in place': Drought and poverty dynamics in South Wollo, Ethiopia. *The journal of development studies*, 42(2):200–225, 2006.

J. Liu, T. Dietz, S. R. Carpenter, M. Alberti, C. Folke, E. Moran, A. N. Pell, P. Deadman, T. Kratz, J. Lubchenco, E. Ostrom, Z. Ouyang, W. Provencher, C. L. Redman, S. H. Schneider, and W. W. Taylor. Complexity of coupled human and natural systems. *Science*, 317(5844):1513–1516, Sept. 2007.

T. M. Logan, M. J. Anderson, T. G. Williams, and L. Conrow. Measuring inequalities in urban systems: An approach for evaluating the distribution of amenities and burdens. *Computers, environment and urban systems*, 86:101590, Mar. 2021.

S. M. C. Loureiro, J. Romero, and R. G. Bilro. Stakeholder engagement in co-creation processes for innovation: A systematic literature review and case study. *Journal of business research*, 119: 388–409, Oct. 2020.

N. Z. Lupwayi, M. Girma, and I. Haque. Plant nutrient contents of cattle manures from small-scale farms and experimental stations in the Ethiopian highlands. *Agriculture, ecosystems & environment*, 78(1):57–63, Mar. 2000.

W. K. Luseno, J. G. McPeak, C. B. Barrett, P. D. Little, and G. Gebru. Assessing the Value of Climate Forecast Information for Pastoralists: Evidence from Southern Ethiopia and Northern Kenya. *World development*, 31(9):1477–1494, Sept. 2003.

T. J. Lybbert and M. R. Carter. Chapter 22 - Bundling Drought Tolerance and Index Insurance to Reduce Rural Household Vulnerability to Drought. In A. M. Balisacan, U. Chakravorty, and M.-L. V. Ravago, editors, *Sustainable Economic Development*, pages 401–414. Academic Press, San Diego, Jan. 2015.

N. R. Magliocca and E. C. Ellis. Evolving human landscapes: a virtual laboratory approach. *Journal of land use science*, 11(6):642–671, Nov. 2016.

N. R. Magliocca, D. G. Brown, and E. C. Ellis. Exploring Agricultural Livelihood Transitions with an Agent-Based Virtual Laboratory: Global Forces to Local Decision-Making. *PloS one*, 8(9): e73241, Sept. 2013.

N. R. Magliocca, J. van Vliet, C. Brown, T. P. Evans, T. Houet, P. Messerli, J. P. Messina, K. A. Nicholas, C. Ornetsmüller, J. Sagebiel, V. Schweizer, P. H. Verburg, and Q. Yu. From meta-studies to modeling: Using synthesis knowledge to build broadly applicable process-based land change models. *Environmental Modelling & Software*, 72:10–20, Oct. 2015.

N. R. Magliocca, E. C. Ellis, G. R. H. Allington, A. de Bremond, J. Dell'Angelo, O. Mertz, P. Messerli, P. Meyfroidt, R. Seppelt, and P. H. Verburg. Closing global knowledge gaps: Producing generalized knowledge from case studies of social-ecological systems. *Global environmental change: human and policy dimensions*, 50:1–14, May 2018.

B. Mahault, A. Saxena, and C. Nisoli. Emergent inequality and self-organized social classes in a network of power and frustration. *PloS one*, 12(2):e0171832, Feb. 2017.

A. Malawska and C. J. Topping. Evaluating the role of behavioral factors and practical constraints in the performance of an agent-based model of farmer decision making. *Agricultural systems*, 143:136–146, Mar. 2016.

A. Malik, A. Crooks, H. Root, and M. Swartz. Exploring creativity and urban development with agent-based modeling. *Journal of artificial societies and social simulation: JASSS*, 18(2), Mar. 2015.

B. Manyena, G. O'Brien, P. O'Keefe, and J. Rose. Disaster resilience: a bounce back or bounce forward ability? *Local Environment: The International Journal of Justice and Sustainability*, 16(5):417–424, 2011.

S. Manzoni and A. Porporato. Soil carbon and nitrogen mineralization: Theory and models across scales. *Soil biology & biochemistry*, 41(7):1355–1379, July 2009.

N. A. Marshall, I. J. Gordon, and A. J. Ash. The reluctance of resource-users to adopt seasonal climate forecasts to enhance resilience to climate variability on the rangelands. *Climatic change*, 107(3):511–529, Aug. 2011.

E. A. Martini, J. S. Buyer, D. C. Bryant, T. K. Hartz, and R. F. Denison. Yield increases during the organic transition: improving soil quality or increasing experience? *Field crops research*, 86(2): 255–266, Mar. 2004.

G. Martiniello. Bitter sugarification: sugar frontier and contract farming in Uganda. *Globalizations*, pages 1–17, July 2020.

P. A. Matson, W. J. Parton, A. G. Power, and M. J. Swift. Agricultural intensification and ecosystem properties. *Science*, 277(5325):504–509, July 1997.

R. B. Matthews, N. G. Gilbert, A. Roach, J. Gary Polhill, and N. M. Gotts. Agent-based land-use models: a review of applications. *Landscape ecology*, 22(10):1447–1459, Dec. 2007.

M. McDermott, S. Mahanty, and K. Schreckenberg. Examining equity: A multidimensional framework for assessing equity in payments for ecosystem services. *Environmental science & policy*, 33:416–427, Nov. 2013.

P. McIntosh. White privilege and male privilege: A personal account of coming to see correspondences through work. *Privilege and prejudice: Twenty years with the invisible knapsack*, page 7, 2020.

E.-M. Meemken and M. F. Bellemare. Smallholder farmers and contract farming in developing countries. *Proceedings of the National Academy of Sciences of the United States of America*, 117(1):259–264, Jan. 2020.

N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan. A Survey on Bias and Fairness in Machine Learning. Aug. 2019.

P. Messerli, M. Giger, M. B. Dwyer, T. Breu, and S. Eckert. The geography of large-scale land acquisitions: Analysing socio-ecological patterns of target contexts in the global South. *Applied geography*, 53:449–459, Sept. 2014.

K. A. Metherell. CENTURY Soil Organic Matter Model Environment. *Technical Documentation Agroecosystem*, 1993.

P. Meyfroidt, R. Roy Chowdhury, A. de Bremond, E. C. Ellis, K.-H. Erb, T. Filatova, R. D. Garrett, J. M. Grove, A. Heinimann, T. Kuemmerle, C. A. Kull, E. F. Lambin, Y. Landon, Y. le Polain de Waroux, P. Messerli, D. Müller, J. Nielsen, G. D. Peterson, V. Rodriguez García, M. Schlüter,

B. L. Turner, and P. H. Verburg. Middle-range theories of land system change. *Global environmental change: human and policy dimensions*, 53:52–67, Nov. 2018.

H. C. Michelson. Influence of neighbor experience and exit on small farmer market participation. *American journal of agricultural economics*, 99(4):952–970, July 2017.

B. L. Miller and M. J. Shaw. Genetic algorithms with dynamic niche sharing for multimodal function optimization. In *Proceedings of IEEE International Conference on Evolutionary Computation*, pages 786–791. IEEE, 1996.

F. Miller, H. Osbahr, E. Boyd, F. Thomalla, S. Bharwani, G. Ziervogel, B. Walker, J. Birkmann, S. van der Leeuw, J. Rockström, J. Hinkel, T. Downing, C. Folke, and D. Nelson. Resilience and Vulnerability: Complementary or Conflicting Concepts? *Ecology and Society*, 15(3), Sept. 2010.

H. R. Milner. Race, Culture, and Researcher Positionality: Working Through Dangers Seen, Unseen, and Unforeseen. *Educational researcher*, 36(7):388–400, Oct. 2007.

A. M. Mobarak and M. R. Rosenzweig. Informal Risk Sharing, Index Insurance, and Risk Taking in Developing Countries. *The American economic review*, 103(3):375–380, May 2013.

T. Mogues. Shocks, Livestock Asset Dymanics and Social Capital in Ethiopia: DSGD Discussion Paper No. 38. Technical report, IFPRI, 2006.

M.-L. Moore, O. Tjornbo, E. Enfors, C. Knapp, J. Hodbod, J. A. Baggio, A. Norström, P. Olsson, and D. Biggs. Studying the complexity of change: toward an analytical framework for understanding deliberate social-ecological transformations. *Ecology and Society*, 19(4), 2014.

J. Morduch. Income Smoothing and Consumption Smoothing. *The journal of economic perspectives: a journal of the American Economic Association*, 9(3):103–114, Sept. 1995.

T. Moreda. Listening to their silence? The political reaction of affected communities to large-scale land acquisitions: insights from Ethiopia. *The Journal of peasant studies*, 42(3-4):517–539, July 2015.

T. Moreda. The right to food in the context of large-scale land investment in Ethiopia. *Third world quarterly*, 39(7):1326–1347, July 2018.

A. Motchoulski. Adjudicating distributive disagreement. *Synthese*, Oct. 2019.

I. Moya, M. Chica, and Cordón. A multicriteria integral framework for agent-based model calibration using evolutionary multiobjective optimization and network-based visualization. *Decision support systems*, page 113111, July 2019.

S. Moyo and F. J. C. Swanepoel. Multifunctionality of livestock in developing communities. *The role of livestock in developing communities: Enhancing multifunctionality*, 3:69, 2010.

B. Müller and D. Kreuer. Ecologists Should Care about Insurance, too. *Trends in ecology & evolution*, 31(1):1–2, Jan. 2016.

B. Müller, M. F. Quaas, K. Frank, and S. Baumgärtner. Pitfalls and potential of institutional change: Rain-index insurance and the sustainability of rangeland management. *Ecological economics: the journal of the International Society for Ecological Economics*, 70(11):2137–2144, Sept. 2011.

B. Müller, F. Bohn, G. Dressler, J. Groeneveld, C. Klassert, R. Martin, M. Schlüter, J. Schulze, H. Weise, and N. Schwarz. Describing human decisions in agent-based models – ODD + D, an extension of the ODD protocol. *Environmental Modelling & Software*, 48:37–48, Oct. 2013.

B. Müller, L. Johnson, and D. Kreuer. Maladaptive outcomes of climate insurance in agriculture. *Global environmental change: human and policy dimensions*, 46:23–33, Sept. 2017.

B. Müller, F. Hoffmann, T. Heckelei, C. Müller, T. W. Hertel, J. G. Polhill, M. van Wijk, T. Achterbosch, P. Alexander, C. Brown, D. Kreuer, F. Ewert, J. Ge, J. D. A. Millington, R. Seppelt, P. H. Verburg, and H. Webber. Modelling food security: Bridging the gap between the micro and the macro scale. *Global environmental change: human and policy dimensions*, 63:102085, July 2020.

M. F. Müller, G. Penny, M. T. Niles, V. Ricciardi, D. D. Chiarelli, K. F. Davis, J. Dell'Angelo, P. D'Odorico, L. Rosa, M. C. Rulli, and N. D. Mueller. Impact of transnational land acquisitions on local food security and dietary diversity. *Proceedings of the National Academy of Sciences of the United States of America*, 118(4), Jan. 2021.

V. Nanhthavong, C. Oberlack, C. Hett, P. Messerli, and M. Epprecht. Pathways to human well-being in the context of land acquisitions in Lao PDR. *Global environmental change: human and policy dimensions*, 68:102252, May 2021.

M. Narisi, K. F. Birungi, and M. Paul. Smallholder farmers' access to and use of scientific climatic forecast information in Mt. Elgon Area, Eastern Uganda. *International Journal of Agricultural Science, Research and Technology in Extension and Education Systems*, 8(1):29–42, 2018.

National Research Council. *Assessing the Reliability of Complex Models: Mathematical and Statistical Foundations of Verification, Validation, and Uncertainty Quantification*. National Academies Press, July 2012.

K. Newcombe. An economic justification for rural afforestation: The case of Ethiopia. *The Annals of regional science*, 21(3):80–99, Nov. 1987.

T. L. Ng, J. W. Eheart, X. Cai, and J. B. Braden. An agent-based model of farmer decision-making and water quality impacts at the watershed scale under markets for carbon allowances and a second-generation biofuel crop. *Water resources research*, 47(9):W09519, Sept. 2011.

H. K. Nguyen, R. Chiong, M. Chica, R. H. Middleton, and D. K. Pham. Contract Farming in the Mekong Delta's Rice Supply Chain: Insights from an Agent-Based Modeling Study. *Journal of Artificial Societies & Social Simulation*, 22(3), 2019.

C. F. Nicholson, E. C. Stephens, A. D. Jones, B. Kopainsky, D. Parsons, and J. Garrett. Setting priorities to address the research gaps between agricultural systems analysis and food security outcomes in low-and middle-income countries. Technical Report CGIAR Working Paper No. 255, 2019.

M. Nussbaum and A. Sen. *The Quality of Life*. Clarendon Press, Mar. 1993.

C. Oberlack, L. Tejada, P. Messerli, S. Rist, and M. Giger. Sustainable livelihoods in the global land rush? Archetypes of livelihood vulnerability and sustainability potentials. *Global environmental change: human and policy dimensions*, 41:153–171, Nov. 2016.

Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453, Oct. 2019.

C. O'Connor. Soil matters: How the Federal Crop Insurance Program should be reformed to encourage low-risk farming methods with high-reward environmental outcomes. In *Agricultural & Applied Economics Association's 2013 Crop Insurance and the Farm Bill Symposium*. Natural Resources Defense Council, 2013.

P. O'Hare, I. White, and A. Connelly. Insurance as maladaptation: Resilience and the 'business as usual'paradox. *Environment and planning. C, Government & policy*, 34(6):1175–1193, 2016.

F. Olalotiti-Lawal and A. Datta-Gupta. A multi-objective Markov chain Monte Carlo approach for history matching and uncertainty quantification. In *SPE Annual Technical Conference and Exhibition*. onepetro.org, 2015.

G. d. L. T. Oliveira, B. M. McKay, and J. Liu. Beyond land grabs: new insights on land struggles and global agrarian change. *Globalizations*, pages 1–18, Nov. 2020.

C. O'Neil. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown, 2016.

K. Orach, A. Duit, and M. Schlüter. Sustainable natural resource governance under interest group competition in policy-making. *Nature human behaviour*, 4(9):898–909, Sept. 2020.

N. Oreskes, K. Shrader-Frechette, and K. Belitz. Verification, validation, and confirmation of numerical models in the Earth sciences. *Science*, 263(5147):641–646, Feb. 1994.

M. G. Orr, S. Galea, M. Riddle, and G. A. Kaplan. Reducing racial disparities in obesity: simulating the effects of improved education and social network influence on diet behavior. *Annals of epidemiology*, 24(8):563–569, Aug. 2014.

E. Ostrom. A general framework for analyzing sustainability of social-ecological systems. *Science*, 325(5939):419–422, July 2009.

D. O'Sullivan, T. Evans, S. Manson, S. Metcalf, A. Ligmann-Zielinska, and C. Bone. Strategic directions for agent-based modeling: avoiding the YAAWN syndrome. *Journal of land use science*, 11(2):177–187, 2016.

C. Oya. Contract farming in sub-Saharan Africa: A survey of approaches, debates and issues. *Journal of Agrarian Change*, 12(1):1–33, 2012.

S. E. Page. On incentives and updating in agent based models. *Computational Economics*, 10(1): 67–87, 1997.

S. E. Page. *Diversity and Complexity*. Princeton University Press, Nov. 2010.

E. H. Pangaribowo, N. Gerber, and M. Torero. Food and Nutrition Security Indicators: A Review. Technical Report ZEF Working Paper 108, Feb. 2013.

H.-Y. Park, A. Datta-Gupta, and M. J. King. Handling conflicting multiple objectives using Pareto-based evolutionary algorithm for history matching of reservoir performance. In *SPE reservoir simulation symposium*. onepetro.org, 2013.

D. C. Parker and V. Meretsky. Measuring pattern outcomes in an agent-based model of edge-effect externalities using spatial metrics. *Agriculture, ecosystems & environment*, 101(2):233–250, Feb. 2004.

D. C. Parker, S. M. Manson, M. A. Janssen, M. J. Hoffmann, and P. Deadman. Multi-agent systems for the simulation of land-use and land-cover change: a review. *Annals of the Association of American Geographers. Association of American Geographers*, 93(2):314–337, 2003.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *The Journal of Machine Learning Research*, 12:2825–2830, 2011.

R. H. M. Pereira, T. Schwanen, and D. Banister. Distributive justice and equity in transportation. *Transport Reviews*, 37(2):170–191, Mar. 2017.

C. A. Peterson, V. T. Eviner, and A. C. M. Gaudin. Ways forward for resilience research in agroecosystems. *Agricultural systems*, 162:19–27, May 2018.

J. R. Platt. Strong Inference. *Science*, 146(3642):347–353, 1964.

S. Polasky, S. R. Carpenter, C. Folke, and B. Keeler. Decision-making under great uncertainty: environmental management in an era of global change. *Trends in ecology & evolution*, 26(8): 398–404, Aug. 2011.

C. Ponsiglione, V. Roma, F. Zampella, and G. Zollo. The Fairness/Efficiency Issue Explored Through El Farol Bar Model. In *Scientific Methods for the Treatment of Uncertainty in Social Sciences*, pages 309–327. Springer International Publishing, 2015.

J. M. Powell, R. A. Pearson, and P. H. Hiernaux. Crop–Livestock Interactions in the West African Drylands. *Agronomy journal*, 96:469–483, 2004.

C. E. Pray, W. A. Masters, and S. Ayoub. Impacts of agricultural research on poverty, malnutrition and resilience. Technical report, USAID/Bureau for Food Security, Office of Agriculture, Research and Policy, 2017.

B. L. Preston, A. W. King, K. M. Ernst, S. M. Absar, S. S. Nair, and E. S. Parish. Scale and the representation of human agency in the modeling of agroecosystems. *Current Opinion in Environmental Sustainability*, 14:239–249, June 2015.

J. Pretty, T. G. Benton, Z. P. Bharucha, L. V. Dicks, C. B. Flora, H. C. J. Godfray, D. Goulson, S. Hartley, N. Lampkin, C. Morris, G. Pierzynski, P. V. V. Prasad, J. Reganold, J. Rockström, P. Smith, P. Thorne, and S. Wratten. Global assessment of agricultural system redesign for sustainable intensification. *Nature Sustainability*, 1(8):441–446, Aug. 2018.

J. N. Pretty. Participatory learning for sustainable agriculture. *World development*, 23(8):1247–1263, Aug. 1995.

M. Prowse. Contract farming in developing countries: a review. In *A Savoir 12*. Paris: Agence Française de Développement., 2012.

M. Quaas, S. Baumgärtner, and M. De Lara. Insurance value of natural capital. *Ecological economics: the journal of the International Society for Ecological Economics*, 165:106388, Nov. 2019.

S. Radosavljevic, L. J. Haider, S. J. Lade, and M. Schlüter. Implications of poverty traps across levels. *World development*, 144:105437, Aug. 2021.

C. Ragasa, I. Lambrecht, and D. S. Kufoalor. Limitations of Contract Farming as a Pro-poor Strategy: The Case of Maize Outgrower Schemes in Upper West Ghana. *World development*, 102:30–56, Feb. 2018.

A. Rajkomar, M. Hardt, M. D. Howell, G. Corrado, and M. H. Chin. Ensuring Fairness in Machine Learning to Advance Health Equity. *Annals of internal medicine*, 169(12):866–872, Dec. 2018.

J. Rawls. *Justice as Fairness: A Restatement*. Harvard University Press, May 2001.

J. Rawls. *A theory of justice*. Harvard University Press, 2009.

T. Redda. Small-scale milk marketing and processing in Ethiopia. In *Smallholder dairy production and marketing: Opportunities and constraints*, pages 352–367, 2002.

M. S. Reed. Stakeholder participation for environmental management: A literature review. *Biological conservation*, 141(10):2417–2431, Oct. 2008.

P. M. Reed, D. Hadka, J. D. Herman, J. R. Kasprzyk, and J. B. Kollat. Evolutionary multiobjective optimization in water resources: The past, present, and future. *Advances in water resources*, 51: 438–456, Jan. 2013.

D. W. Reeves. The role of soil organic matter in maintaining soil quality in continuous cropping systems. *Soil and Tillage Research*, 43(1):131–167, Nov. 1997.

A. C. Reilly, S. D. Guikema, L. Zhu, and T. Igusa. Evolution of vulnerability of communities facing repeated hazards. *PloS one*, 12(9):e0182719, Sept. 2017.

A. C. Reilly, R. L. Dillon, and S. D. Guikema. Agent-based models as an integrating boundary object for interdisciplinary research. *Risk analysis: an official publication of the Society for Risk Analysis*, June 2018.

V. Ricciardi, Z. Mehrabi, H. Wittman, D. James, and N. Ramankutty. Higher yields and more biodiversity on smaller farms. *Nature Sustainability*, pages 1–7, Mar. 2021.

J. Risbey, M. Kandlikar, H. Dowlatabadi, and D. Graetz. Scale, context, and decision making in agricultural adaptation to climate variability and change. *Mitigation and Adaptation Strategies for Global Change*, 4(2):137–165, June 1999.

G. P. Robertson and P. M. Vitousek. Nitrogen in Agriculture: Balancing the Cost of an Essential Resource. *Annual review of environment and resources*, 34(1):97–125, Nov. 2009.

S. Robinson, H. van Meijl, D. Willenbockel, H. Valin, S. Fujimori, T. Masui, R. Sands, M. Wise, K. Calvin, P. Havlik, D. Mason d'Croz, A. Tabeau, A. Kavallari, C. Schmitz, J. P. Dietrich, and M. von Lampe. Comparing supply-side specifications in models of global agriculture and the food system. *Agricultural economics*, 45(1):21–35, Jan. 2014.

J. P. Rodríguez, T. D. Beard, E. M. Bennett, G. S. Cumming, S. J. Cork, J. Agard, A. P. Dobson, and G. D. Peterson. Trade-offs across Space, Time, and Ecosystem Services. *Ecology and Society*, 11(1), 2006.

J. Rosa-Schleich, J. Loos, O. Mußhoff, and T. Tscharntke. Ecological-economic trade-offs of Diversified Farming Systems – A review. *Ecological economics: the journal of the International Society for Ecological Economics*, 160:251–263, June 2019.

M. C. Rulli and P. D'Odorico. Food appropriation through large scale land acquisitions. *Environmental Research Letters*, 9(6):064030, May 2014.

E. Sadoulet and A. De Janvry. *Quantitative development policy analysis*, volume 5. Johns Hopkins University Press Baltimore, 1995.

P. A. Sanchez, K. D. Shepherd, M. J. Soule, F. M. Place, R. J. Buresh, A.-M. N. Izac, A. U. Mokwunye, F. R. Kwesiga, C. G. Ndiritu, and P. L. Woomer. Soil Fertility Replenishment in Africa: An Investment in Natural Resource Capital. In *Replenishing Soil Fertility in Africa*, SSSA Special Publication, pages 1–46. Soil Science Society of America and American Society of Agronomy, Madison, WI, 1997.

M. Sandker, B. M. Campbell, M. Ruiz-Pérez, J. A. Sayer, R. Cowling, H. Kassa, and A. T. Knight. The Role of Participatory Modeling in Landscape Approaches to Reconcile Conservation and Development. *Ecology and Society*, 15(2), 2010.

J. C. Schank, P. E. Smaldino, and M. L. Miller. Evolution of fairness in the dictator game by multilevel selection. *Journal of theoretical biology*, 382:64–73, Oct. 2015.

S. Schaub, N. Buchmann, A. Lüscher, and R. Finger. Economic benefits from plant species diversity in intensively managed grasslands. *Ecological economics: the journal of the International Society for Ecological Economics*, 168:106488, Feb. 2020.

T. C. Schelling. Dynamic models of segregation. *The Journal of mathematical sociology*, 1(2): 143–186, July 1971.

J. Schindler. Rethinking the tragedy of the commons: The integration of Socio-psychological dispositions. *Journal of artificial societies and social simulation: JASSS*, 15(1), Jan. 2012.

M. Schlüter, A. Baeza, G. Dressler, K. Frank, J. Groeneveld, W. Jager, M. A. Janssen, R. R. J. McAllister, B. Müller, K. Orach, N. Schwarz, and N. Wijermans. A framework for mapping and comparing behavioural theories in models of social-ecological systems. *Ecological economics: the journal of the International Society for Ecological Economics*, 131(Supplement C):21–35, Jan. 2017.

M. Schlüter, L. J. Haider, S. J. Lade, E. Lindkvist, R. Martin, K. Orach, N. Wijermans, and C. Folke. Capturing emergent phenomena in social-ecological systems: an analytical framework. *Ecology and Society*, 24(3), 2019a.

M. Schlüter, B. Müller, and K. Frank. The potential of models and modeling for social-ecological systems research: the reference frame ModSES. *Ecology and Society*, 24(1), Mar. 2019b.

M. Schlüter, K. Orach, E. Lindkvist, R. Martin, N. Wijermans, Bodin, and W. J. Boonstra. Toward a methodology for explaining and theorizing about social-ecological phenomena. *Current Opinion in Environmental Sustainability*, 39:44–53, Aug. 2019c.

M. Schlüter, E. Lindkvist, and X. Basurto. The interplay between top-down interventions and bottom-up self-organization shapes opportunities for transforming self-governance in small-scale fisheries. *Marine Policy*, page 104485, Mar. 2021.

M. W. I. Schmidt, M. S. Torn, S. Abiven, T. Dittmar, G. Guggenberger, I. A. Janssens, M. Kleber, I. Kögel-Knabner, J. Lehmann, D. A. C. Manning, P. Nannipieri, D. P. Rasse, S. Weiner, and S. E. Trumbore. Persistence of soil organic matter as an ecosystem property. *Nature*, 478(7367): 49–56, Oct. 2011.

P. J. Schmidt, M. B. Emelko, and M. E. Thompson. Recognizing Structural Nonidentifiability: When Experiments Do Not Provide Information About Important Parameters and Misleading Models Can Still Have Great Fit. *Risk analysis: an official publication of the Society for Risk Analysis*, 40(2):352–369, Feb. 2020.

A. Schmolke, P. Thorbek, D. L. DeAngelis, and V. Grimm. Ecological models supporting environmental decision making: a strategy for the future. *Trends in ecology & evolution*, 25(8): 479–486, Aug. 2010.

G. C. Schoneveld. The anatomy of large-scale farmland acquisitions in sub-Saharan Africa. *CIFOR Working Paper*, (85), 2011.

G. C. Schoneveld. Host country governance and the African land rush: 7 reasons why large-scale farmland investments fail to contribute to sustainable development. *Geoforum; journal of physical, human, and regional geosciences*, 83:119–132, July 2017.

P. Schreinemachers, T. Berger, and J. B. Aune. Simulating soil fertility and poverty dynamics in Uganda: A bio-economic multi-agent systems approach. *Ecological economics: the journal of the International Society for Ecological Economics*, 64(2):387–401, Dec. 2007.

F. Schuenemann, J. Thurlow, and M. Zeller. Leveling the field for biofuels: comparing the economic and environmental impacts of biofuel and other export crops in Malawi. *Agricultural economics*, 48(3):301–315, May 2017.

J. Schulze, B. Müller, J. Groeneveld, and V. Grimm. Agent-Based Modelling of Social-Ecological Systems: Achievements, Challenges, and a Way Forward. *Journal of Artificial Societies and Social Simulation, The*, 20(2):1–8, Mar. 2017.

P. Schwartz. *The Art of the Long View: Planning for the Future in an Uncertain World*. Crown, Feb. 2012.

I. Scoones, A. Stirling, D. Abrol, J. Atela, L. Charli-Joseph, H. Eakin, A. Ely, P. Olsson, L. Pereira, R. Priya, P. van Zwanenberg, and L. Yang. Transformations to sustainability: combining structural, systemic and enabling approaches. *Current Opinion in Environmental Sustainability*, 42: 65–75, Feb. 2020.

R. Serraj, T. R. Sinclair, and L. C. Purcell. Symbiotic N2 fixation response to drought. *Journal of experimental botany*, 50(331):143–155, 1999.

K. A. Shankarnarayan, L. N. Harsh, and S. Kathju. Agroforestry in the arid zones of India. *Agroforestry Systems*, 5(1):69–88, 1987.

C. Shennan. Biotic interactions, ecological knowledge and agriculture. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 363(1492):717–739, Feb. 2008.

M. Shete and M. Rutten. Impacts of large-scale farming on local communities' food security and income levels–Empirical evidence from Oromia Region, Ethiopia. *Land use policy*, 2015.

T. Sikor, A. Martin, J. Fisher, and J. He. Toward an Empirical Analysis of Justice in Ecosystem Governance: Justice in ecosystem governance. *Conservation letters*, 7(6):524–532, Nov. 2014.

B. Simane, B. F. Zaitchik, and M. Ozdogan. Agroecosystem Analysis of the Choke Mountain Watersheds, Ethiopia. *Sustainability: Science Practice and Policy*, 5(2):592–616, Feb. 2013.

H. A. Simon. A Behavioral Model of Rational Choice. *The quarterly journal of economics*, 69(1): 99–118, Feb. 1955.

H. A. Simon. Rational Decision Making in Business Organizations. *The American economic review*, 69(4):493–513, 1979.

G. Singh and K. Deb, Dr. Comparison of Multi-modal Optimization Algorithms Based on Evolutionary Algorithms. In *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, GECCO '06, pages 1305–1312, New York, NY, USA, 2006. ACM.

R. Singh, M. Worku, S. Bogale, A. Cullis, A. Adem, B. Irwin, S. Lim, L. Bosi, and C. C. Venton. Reality of resilience: Perspectives of the 2015-16 drought in Ethiopia. *BRACED Resilience Intel*, (6):1–28, 2016.

S. Singh. Contracting Out Solutions: Political Economy of Contract Farming in the Indian Punjab. *World development*, 30(9):1621–1638, Sept. 2002.

H. Smart. Operationalizing a Conceptual Model of Colorism in Local Policing. *Social justice research*, 32(1):72–115, Mar. 2019.

L. T. Smith. *Decolonizing Methodologies: Research and Indigenous Peoples*. Zed Books Ltd., Oct. 2013.

P. Smith, K. J. Beven, and J. A. Tawn. Informal likelihood measures in model assessment: Theoretic development and investigation. *Advances in water resources*, 31(8):1087–1100, Aug. 2008.

S. S. Snapp, S. M. Swinton, R. Labarta, D. Mutch, J. R. Black, R. Leep, J. Nyiraneza, and K. O'Neil. Evaluating Cover Crops for Benefits, Costs and Performance within Cropping System Niches. *Agronomy journal*, 97:322–332, 2005.

P. Sobkowicz. Agent based model of effects of task allocation strategies in flat organizations. *Physica A: Statistical Mechanics and its Applications*, 458:17–30, Sept. 2016.

N. Speybroeck, C. Van Malderen, S. Harper, B. Müller, and B. Devleesschauwer. Simulation models for socioeconomic inequalities in health: a systematic review. *International journal of environmental research and public health*, 10(11):5750–5780, Nov. 2013.

C. Steger, S. Hirsch, C. Cosgrove, S. Inman, E. Nost, X. Shinbrot, J. P. R. Thorn, D. G. Brown, A. Grêt-Regamey, B. Müller, R. S. Reid, C. Tucker, B. Weibel, and J. A. Klein. Linking model design and application for transdisciplinary approaches in social-ecological systems. *Global environmental change: human and policy dimensions*, 66(102201):102201, Jan. 2021a.

C. Steger, J. A. Klein, R. S. Reid, S. Lavorel, C. Tucker, K. A. Hopping, R. Marchant, T. Teel, A. Cuni-Sanchez, T. Dorji, G. Greenwood, R. Huber, K.-A. Kassam, D. Kreuer, A. Nolin, A. Russell, J. L. Sharp, M. Šmid Hribar, J. P. R. Thorn, G. Grant, M. Mahdi, M. Moreno, and D. Waiswa. Science with society: Evidence-based guidance for best practices in environmental transdisciplinary work. *Global environmental change: human and policy dimensions*, 68:102240, May 2021b.

E. C. Stephens, C. F. Nicholson, D. R. Brown, D. Parsons, C. B. Barrett, J. Lehmann, D. Mbugua, S. Ngoze, A. N. Pell, and S. J. Riha. Modeling the impact of natural resource-based poverty traps on food security in Kenya: The Crops, Livestock and Soils in Smallholder Economic Systems (CLASSES) model. *Food Security*, 4(3):423–439, 2012.

F. Stonedahl and U. Wilensky. Evolutionary Robustness Checking in the Artificial Anasazi Model. In *AAAI Fall Symposium: Complex Adaptive Systems*, pages 120–129. aaai.org, 2010.

A. E. Stratton, L. Kuhl, and J. Blesh. Ecological and nutritional functions of agroecosystems as indicators of smallholder resilience. *Frontiers in Sustainable Food Systems*, 4:173, 2020.

Z. Sun, I. Lorscheid, J. D. Millington, S. Lauf, N. R. Magliocca, J. Groeneveld, S. Balbi, H. Nolzen, B. Müller, J. Schulze, and C. M. Buchmann. Simple or complicated agent-based models? A complicated issue. *Environmental Modelling & Software*, 86:56–67, Dec. 2016.

C. Swinerd and K. R. McNaught. Simulating the diffusion of technological innovation with an integrated hybrid agent-based system dynamics model. *Journal of Simulation*, 8(3):231–240, Aug. 2014.

SwissRe. Partnering for food security in emerging markets. *SIGMA. Symmetry, Integrability and Geometry. Methods and Applications*, 1:1–44, 2013.

A. Sánchez and J. A. Cuesta. Altruism may arise from individual selection. *Journal of theoretical biology*, 235(2):233–240, July 2005.

T. Tabucchi, R. Davidson, and S. Brink. Simulation of post-earthquake water supply system restoration. *Civil Engineering and Environmental Systems*, 27(4):263–279, Dec. 2010.

Y. Tamura. Contexts behind differentiated responses to contract farming and large-scale land acquisitions in Central Mozambique: Post-war experiences, social relations, and power balance of local authorities. *Land use policy*, 106:105439, July 2021.

S. Tarawali, M. Herrero, K. Descheemaeker, E. Grings, and M. Blümmel. Pathways for sustainable development of mixed crop livestock systems: Taking a livestock and pro-poor approach. *Livestock science*, 139(1):11–21, July 2011.

H. Teklewold, M. Kassie, B. Shiferaw, and G. Köhlin. Cropping system diversification, conservation tillage and modern seed adoption in Ethiopia: Impacts on household income, agrochemical use and demand for labor. *Ecological economics: the journal of the International Society for Ecological Economics*, 93:85–93, 2013.

G. A. ten Broeke, G. A. K. van Voorn, A. Ligtenberg, and J. Molenaar. Cooperation can improve the resilience of common-pool resource systems against over-harvesting. *Ecological Complexity*, 40:100742, Dec. 2019.

J. C. Thiele, W. Kurth, and V. Grimm. Facilitating parameter estimation and sensitivity analysis of agent-based models: A cookbook using NetLogo and R. *Journal of Artificial Societies and Social Simulation*, 17(3):11, 2014.

P. K. Thornton and M. Herrero. Adapting to climate change in the mixed crop and livestock farming systems in sub-Saharan Africa. *Nature climate change*, 5(9):830–836, Sept. 2015.

Q. Tian, J. H. Holland, and D. G. Brown. Social and economic impacts of subsidy policies on rural development in the Poyang Lake Region, China: Insights from an agent-based model. *Agricultural systems*, 148:12–27, Oct. 2016.

P. Tittonell. Livelihood strategies, resilience and transformability in African agroecosystems. *Agricultural systems*, 126:3–14, Apr. 2014.

P. Tittonell and K. E. Giller. When yield gaps are poverty traps: The paradigm of ecological intensification in African smallholder agriculture. *Field crops research*, 143:76–90, Mar. 2013.

P. Tittonell, P. A. Leffelaar, B. Vanlauwe, M. T. van Wijk, and K. E. Giller. Exploring diversity of crop and soil management within smallholder African farms: A dynamic model for simulation of N balances and use efficiencies at field scale. *Agricultural systems*, 91(1):71–101, Nov. 2006.

R. S. J. Tol, T. E. Downing, O. J. Kuik, and J. B. Smith. Distributional aspects of climate change impacts. *Global environmental change: human and policy dimensions*, 14(3):259–272, Oct. 2004.

T. P. Tomich, P. Lidder, M. Coley, D. Gollin, R. Meinzen-Dick, P. Webb, and P. Carberry. Food and agricultural innovation pathways for prosperity. *Agricultural systems*, 172:1–15, June 2019a.

T. P. Tomich, P. Lidder, J. Dijkman, M. Coley, P. Webb, and M. Gill. Agri-food systems in international research for development: Ten theses regarding impact pathways, partnerships, program design, and priority-setting for rural prosperity. *Agricultural systems*, 172:101–109, June 2019b.

G. Ton, W. Vellema, S. Desiere, S. Weituschat, and M. D'Haese. Contract farming for improving smallholder incomes: What can we learn from effectiveness studies? *World development*, 104: 46–64, Apr. 2018.

G. L. Tonn and S. D. Guikema. An Agent-Based Model of Evolving Community Flood Risk. *Risk analysis: an official publication of the Society for Risk Analysis*, 38(6):1258–1278, June 2018.

S. Touhidul Mustafa, J. Nossent, G. Ghysels, and M. Huysmans. Integrated Bayesian Multi-model approach to quantify input, parameter and conceptual model structure uncertainty in groundwater modeling. *Environmental Modelling & Software*, page 104654, Feb. 2020.

M. C. Turley and E. D. Ford. Definition and calculation of uncertainty in ecological process models. *Ecological modelling*, 220(17):1968–1983, Sept. 2009.

T. R. Tyler. Social justice: Outcome and procedure. *International journal of psychology: Journal international de psychologie*, 35(2):117–125, Apr. 2000.

UN General Assembly. Transforming our world : the 2030 Agenda for Sustainable Development. Technical report, United Nations, 2015.

P. W. Unger and M. F. Vigil. Cover crop effects on soil water relationships. *Journal of Soil and Water Conservation*, 53(3):200–207, July 1998.

N. Urruty, D. Tailliez-Lefebvre, and C. Huyghe. Stability, robustness, vulnerability and resilience of agricultural systems. A review. *Agronomy for Sustainable Development*, 36(1):15, Mar. 2016.

USDA. *CENTURY Model*, 1993.

D. Valbuena, O. Erenstein, S. Homann-Kee Tui, T. Abdoulaye, L. Claessens, A. J. Duncan, B. Gérard, M. C. Rufino, N. Teufel, A. van Rooyen, and M. T. van Wijk. Conservation Agriculture in mixed crop–livestock systems: Scoping crop residue trade-offs in Sub-Saharan Africa and South Asia. *Field crops research*, 132:175–184, June 2012.

D. Valente, P. P. Miglietta, D. Porrini, M. R. Pasimeni, G. Zurlini, and I. Petrosillo. A first analysis on the need to integrate ecological aspects into financial insurance. *Ecological modelling*, 392: 117–127, Jan. 2019.

G. Van den Broeck, J. Swinnen, and M. Maertens. Global value chains, large-scale farming, and poverty: Long-term effects in Senegal. *Food policy*, 66:97–107, Jan. 2017.

J. van Vliet, A. K. Bregt, D. G. Brown, H. van Delden, S. Heckbert, and P. H. Verburg. A review of current calibration and validation practices in land-change modeling. *Environmental Modelling & Software*, 82:174–182, Aug. 2016.

M. T. van Wijk. From global economic modelling to household level analyses of food security and sustainability: How big is the gap and can we bridge it? *Food policy*, 49:378–388, Dec. 2014.

T. J. VanderWeele and W. R. Robinson. On the causal interpretation of race in regressions adjusting for confounding and mediating variables. *Epidemiology*, 25(4):473–484, July 2014.

P. H. Verburg, J. A. Dearing, J. G. Dyke, S. v. d. Leeuw, S. Seitzinger, W. Steffen, and J. Syvitski. Methods and approaches to modelling the Anthropocene. *Global environmental change: human and policy dimensions*, 39:328–340, July 2016.

P. H. Verburg, P. Alexander, T. Evans, N. R. Magliocca, Z. Malek, M. D. A. Rounsevell, and J. van Vliet. Beyond land cover change: towards a new generation of land use models. *Current Opinion in Environmental Sustainability*, 38:77–85, June 2019.

S. J. Vermeulen, P. K. Aggarwal, A. Ainslie, C. Angelone, B. M. Campbell, A. J. Challinor, J. W. Hansen, J. S. I. Ingram, A. Jarvis, P. Kristjanson, C. Lau, G. C. Nelson, P. K. Thornton, and E. Wollenberg. Options for support to agriculture and food security under climate change. *Environmental science & policy*, 15(1):136–144, Jan. 2012.

G. B. Villamor, M. van Noordwijk, U. Djanibekov, M. E. Chiong-Javier, and D. Catacutan. Gender differences in land-use decisions: shaping multifunctional landscapes? *Current Opinion in Environmental Sustainability*, 6:128–133, Feb. 2014.

A. Voinov and F. Bousquet. Modelling with stakeholders. *Environmental Modelling & Software*, 25(11):1268–1281, Nov. 2010.

A. Voinov, R. Seppelt, S. Reis, J. E. M. S. Nabel, and S. Shokravi. Values in socio-environmental modelling: Persuasion for action or excuse for inaction. *Environmental Modelling & Software*, 53:207–212, Mar. 2014.

A. Voinov, N. Kolagani, M. K. McCall, P. D. Glynn, M. E. Kragt, F. O. Ostermann, S. A. Pierce, and P. Ramu. Modelling with stakeholders – Next generation. *Environmental Modelling & Software*, 77:196–220, Mar. 2016.

A. Voinov, K. Jenni, S. Gray, N. Kolagani, P. D. Glynn, P. Bommel, C. Prell, M. Zellner, M. Paolisso, R. Jordan, E. Sterling, L. Schmitt Olabisi, P. J. Giabbanelli, Z. Sun, C. Le Page, S. Elsawah, T. K. BenDor, K. Hubacek, B. K. Laursen, A. Jetter, L. Basco-Carrera, A. Singer, L. Young, J. Brunacini, and A. Smajgl. Tools and methods in participatory modeling: Selecting the right tool for the job. *Environmental Modelling & Software*, 109:232–255, Nov. 2018.

J. von Braun and R. Meinzen-Dick. "Land grabbing" by foreign investors in developing countries: Risks and opportunities. Technical Report 13, International Food Policy Research Institute (IFPRI), 2009.

T. Vongvisouk, R. B. Broegaard, O. Mertz, and S. Thongmanivong. Rush for cash crops and forest protection: Neither land sparing nor land sharing. *Land use policy*, 55:182–192, Sept. 2016.

J. A. Vrugt. Markov chain Monte Carlo simulation using the DREAM software package: Theory, concepts, and MATLAB implementation. *Environmental Modelling & Software*, 75:273–316, Jan. 2016.

J. A. Vrugt and K. J. Beven. Embracing equifinality with efficiency: Limits of Acceptability sampling using the DREAM(LOA) algorithm. *Journal of Hydrology*, 559:954–971, Apr. 2018.

J. A. Vrugt, C. J. F. Ter Braak, M. P. Clark, J. M. Hyman, and B. A. Robinson. Treatment of input uncertainty in hydrologic modeling: Doing hydrology backward with Markov chain Monte Carlo simulation. *Water resources research*, 44(12), 2008.

J. A. Vrugt, C. J. F. ter Braak, H. V. Gupta, and B. A. Robinson. Equifinality of formal (DREAM) and informal (GLUE) Bayesian approaches in hydrologic modeling? *Stochastic environmental research and risk assessment: research journal*, 23(7):1011–1026, Oct. 2009.

B. Walker. Resilience: what it is and is not. *Ecology and Society*, 25(2), 2020.

E. Walter. *Identifiability of Parametric Models*. Elsevier, May 2014.

M. Wang, N. White, V. Grimm, H. Hofman, D. Doley, G. Thorp, B. Cribb, E. Wherritt, L. Han, J. Wilkie, and J. Hanan. Pattern-oriented modelling as a novel way to verify and validate functional-structural plant models: a demonstration with the annual growth module of avocado. *Annals of botany*, 121(5):941–959, Apr. 2018.

P. Webb and S. Block. Support for agriculture during economic transformation: impacts on poverty and undernutrition. *Proceedings of the National Academy of Sciences of the United States of America*, 109(31):12309–12314, July 2012.

H. Weise, H. Auge, C. Baessler, I. Bärlund, E. M. Bennett, U. Berger, F. Bohn, A. Bonn, D. Borchardt, F. Brand, A. Chatzinotas, R. Corstanje, F. De Laender, P. Dietrich, S. Dunker, W. Durka, I. Fazey, J. Groeneveld, C. S. E. Guilbaud, H. Harms, S. Harpole, J. Harris, K. Jax, F. Jeltsch, K. Johst, J. Joshi, S. Klotz, I. Kühn, C. Kuhlicke, B. Müller, V. Radchuk, H. Reuter, K. Rinke, M. Schmitt-Jansen, R. Seppelt, A. Singer, R. J. Standish, H. Thulke, B. Tietjen, M. Weitere, C. Wirth, C. Wolf, and V. Grimm. Resilience trinity: safeguarding ecosystem functioning and services across three different time horizons and decision contexts. *Oikos*, 40:97, Jan. 2020.

L. Welch and S. Ekwaro-Osire. Fairness in Agent Based Simulation Frameworks. *Journal of Computing and Information Science in Engineering*, 10(1), Mar. 2010.

M. A. Wendimu, A. Henningsen, and P. Gibbon. Sugarcane Outgrowers in Ethiopia: "Forced" to Remain Poor? *World development*, 83(Supplement C):84–97, July 2016.

M. A. Wendimu, A. Henningsen, and T. G. Czekaj. Incentives and moral hazard: plot level productivity of factory-operated and outgrower-operated sugarcane production in Ethiopia. *Agricultural economics*, 48(5):549–560, 2017.

T. Wiegand, F. Jeltsch, I. Hanski, and V. Grimm. Using pattern-oriented modeling for revealing hidden information: a key for reconciling ecological theory and application. *Oikos*, 100(2): 209–222, 2003.

S. Wiggins, J. Kirsten, and L. Llambí. The Future of Small Farms. *World development*, 38(10): 1341–1348, Oct. 2010.

M. Will, G. Dressler, D. Kreuer, H. Thulke, A. Grêt-Regamey, and B. Müller. How to make socio-environmental modelling more useful to support policy and management? *People and Nature*, (pan3.10207), Mar. 2021.

T. G. Williams, S. Trush, J. Sullivan, C. Liao, N. Chesterman, A. Agrawal, S. D. Guikema, and D. G. Brown. Direct and indirect land-use changes associated with large-scale land acquisitions in Ethiopia. *Under Review in Ecology and Society*.

T. G. Williams, S. D. Guikema, D. G. Brown, and A. Agrawal. Resilience and equity: Quantifying the distributional effects of resilience-enhancing strategies in a smallholder agricultural system. *Agricultural systems*, 182:102832, June 2020a.

T. G. Williams, S. D. Guikema, D. G. Brown, and A. Agrawal. Assessing model equifinality for robust policy analysis in complex socio-environmental systems. *Environmental Modelling & Software*, 134:104831, Dec. 2020b.

T. G. Williams, G. Dreßler, A. E. Stratton, and B. Müller. Ecological and financial strategies provide complementary benefits for smallholder climate resilience: insights from a simulation model. *Ecology and Society*, 2021.

S. Wilson. What is an indigenous research methodology? *Canadian journal of native education*, 25(2):175–179, 2001.

P. Windrum, G. Fagiolo, and A. Moneta. Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation*, 10(2):8, 2007.

R. A. Wittwer, B. Dorn, W. Jossi, and M. G. A. van der Heijden. Cover crops support ecological intensification of arable cropping systems. *Scientific reports*, 7:41911, Feb. 2017.

H. L. Wong, X. Wei, H. B. Kahsay, Z. Gebreegziabher, C. Gardebroek, D. E. Osgood, and R. Diro. Effects of input vouchers and rainfall insurance on agricultural production and household welfare: Experimental evidence from northern Ethiopia. *World development*, 135:105074, Nov. 2020.

I. H. Worku, M. Dereje, B. Minten, and K. Hirvonen. Diet transformation in Africa: the case of Ethiopia. *Agricultural economics*, 48(S1):73–86, Nov. 2017.

World Bank, FAO, IFAD, and UNCTAD. Principles for Responsible Agricultural Investment that Respect Rights, Livelihood and Resources. Technical report, 2010.

T. Wossen, T. Berger, N. Swamikannu, and T. Ramilan. Climate variability, consumption risk and poverty in semi-arid Northern Ghana: Adaptation options for poor farm households. *Environmental Development*, 12(Supplement C):2–15, Oct. 2014.

T. Wossen, T. Berger, and S. Di Falco. Social capital, risk preference and adoption of improved farm land management practices in Ethiopia. *Agricultural economics*, 46(1):81–97, Jan. 2015.

T. Wossen, T. Berger, M. G. Haile, and C. Troost. Impacts of climate variability and food price volatility on household income and food security of farm households in East and West Africa. *Agricultural systems*, Feb. 2017.

J. Xiong, , P , Thenkabail,, J. C. Tilton, , M , Gumma,, P. Teluguntla, , R , Congalton,, K. Yadav, J. Dungan, , A , Oliphant,, J. Poehnelt, C. Smith, and M. R. NASA Making Earth System Data Records for Use in Research Environments (MEaSUREs) Global Food Security-support Analysis Data (GFSAD) Cropland Extent 2015 Africa 30 m V001, 2017.

H. Yen, X. Wang, D. G. Fontane, R. D. Harmel, and M. Arabi. A framework for propagation of uncertainty contributed by parameterization, input data, model structure, and calibration/validation data in watershed modeling. *Environmental Modelling & Software*, 54:211–221, Apr. 2014.

C. Zagaria, C. J. E. Schulp, M. Zavalloni, D. Viaggi, and P. H. Verburg. Modelling transformational adaptation to climate change among crop farming systems in Romagna, Italy. *Agricultural systems*, 188:103024, Mar. 2021.

E. M. Zechman and S. R. Ranjithan. An evolutionary algorithm to generate alternatives (EAGA) for engineering optimization problems. *Engineering Optimization*, 36(5):539–553, Oct. 2004.

J. Zhang and D. T. Robinson. Replication of an agent-based model using the Replication Standard. *Environmental Modelling & Software*, 139:105016, May 2021.

G. Ziervogel, M. Bithell, R. Washington, and T. Downing. Agent-based social simulation: a method for assessing the impact of seasonal climate forecast applications among smallholder farmers. *Agricultural systems*, 83(1):1–26, 2005.

Z. Çağnan, R. A. Davidson, and S. D. Guikema. Post-earthquake restoration planning for Los Angeles electric power. *Earthquake Spectra*, 22(3):589–608, 2006.