

**Systemic Inequality, Technological Innovation, and
the Limits of Human Understanding**

by

Christopher L. Quarles

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Information)
in the University of Michigan
2023

Doctoral Committee:

Professor Scott E. Page, Co-Chair
Associate Professor Daniel Romero, Co-Chair
Associate Professor Ceren Budak
Associate Professor Fabian T. Pfeffer

Christopher L. Quarles

cquarles@umich.edu

ORCID iD: [0000-0002-3953-7564](https://orcid.org/0000-0002-3953-7564)

© Christopher L. Quarles 2023

Dedication

For all those throughout history who gave a part of their life, without reward or fame,
to make the world a better place.

Acknowledgements

The acknowledgements for my doctoral dissertation need to start with Mickey Davis, just like my career as a social scientist did. He helped me go from being a math teacher with no experience in social science research to a published author. Mickey taught me that anyone can be a mentor - a principle I've carried with me through the years. I wouldn't be traveling down this path if he hadn't helped me find the trailhead many years ago. Thank you, Mickey.

I'm also grateful for other mentors who helped me get to UMSI including Alec Campbell, Bonnie January, Darryl Dieter, and Al Friedman.

I want to thank my committee: Daniel Romero, Scott Page, Ceren Budak & Fabian Pfeffer. Daniel & Scott took me on and supported me after some big bumps in my graduate career. They've been a source of wise advice, solid administrative support, and clever research ideas, and I will be eternally grateful.

The number of people that I want to thank here is a sign of how rich my experience over the last few years has been. So many people have contributed significantly to my success, whether through friendship, advice, or being a fabulous colleague. Thanks to (in no particular order): Rob Hladky, Koissi Savi, Bhartendu Pandey, Christine Aidala, Abraham Mhaidli, Yixin Zou, Hao Peng, Shubham Atreja, Ashwin Rajadesingan, Tamy Guberek, Chanda Phelan, Ed Platt, Minje Choi, Aparna Ananthasubramaniam, David Gamba, Stacey Xiang, Yulin Yu, Zhuofeng Wu, Allison Tyler, Devon Keen, Tom Boyce, Melissa Bora, Brinda Gokul, Jared Wadley, Jane Im, Sangmi Kim, Minje Choi, Vaishnav Kameswaran, Linfeng Li, Robert Markum, Jordan Hemingway, Lu Xian, Ihudiya Finda Williams, Howard Manly, Joe LaBriola, Julia Mendelsohn, Alec Kirkley, Siqi Wu, Melissa Chalmers, Laura Schram, Paula Di Rita Wishart, Abigail Jacobs, David Jurgens, Misha Teplitskiy, Tawanna Dillahunt, Kelly Kowatch, Joanna Kroll, Tanya Rosenblat, Colleen van Lent, Chris Teplovs, Jessica Sendra, Charles Yun, and Todd Stuart, as well as others that I have surely forgotten.

I've particularly valued being part of communities like Daniel Romero's lab group, who have been fabulous, and the group that Fabian Pfeffer grew from an "Inequality Lab" to the much fancier "Stone Center for Inequality Dynamics". That growth is due largely to Fabian's superpower: making new people feel welcome and appreciated.

Many thanks to the UM School of Information, Poverty Solutions, and the Santa Fe Institute for research backing. In particular, Beth Yakel, Veronica Falandino, Erin Krupka, and Amy Eaton gave me different types of support that kept me going through the pandemic and beyond. The managers who capably handle all the little things are the ones who create a lot of well-being and are also often underappreciated. I appreciate them.

I also want to acknowledge the loss of Charlie Doering, who gave a heck of a chalk talk. If he was still around, I would surely be thanking him for his contribution to this dissertation.

Super-special thanks go out to Mohamed Abbadi, Pablo Mitnik, Leo & Epiphani Raabe, Cindy Bank, Elizabeth Bruch, Lia Bozarth, Danaja & Nirmala Maldeniya, Song Mi Lee, and Alex Liu. All of them know what they've done, so no explanation is necessary.

It may sound silly, but the individual that I've spent the most time with was my "Republic of Gamers" laptop computer. My computer was with me when I came to Ann Arbor, and it's where I wrote my thesis. We both got considerably beaten up along the way, but we made it through together. I suppose you could say this dissertation was written by a cyborg.

I'll end this section with my gratefulness to my family, who will be with me long after I've left UMSI. My parents & step-parents put in years of work, and deserve more appreciation than I show them (as most parents do). I also want to thank my siblings. My son, Raphael, has been great, putting up with the nights where Dad was working late and all the interesting(?!) conversations about social science. And, of course, words can't express how much I appreciate my wife Ellen. Being successful takes a huge number of skills, traits, and resources, and we don't always have everything we need. Having a strong, capable partner to complement your skills, grow together, and support you through the shocks of life is the most powerful single-factor intervention I can think of. I've been so very lucky.

Table of Contents

Dedication.....	ii
Acknowledgements.....	iii
List of Tables	x
List of Figures.....	xi
Abstract.....	xv
Chapter 1 Introduction	1
1.1 Chapter Two – Patterns of Student Success are Universal and Suggest a Latent Limited Resource	2
1.2 Chapter Three – Conflicts in Understanding of Language on Social Media Can Be Counterproductive and Increase Polarization	2
1.3 Chapter Four – Success Arises from Many, Interacting Factors.....	3
1.4 Chapter Five – Systemic Inequality	4
Chapter 2 The Shape of Educational Inequality	6
2.1 Introduction	6
2.1.1 Student Capital	6
2.1.2 Operationalizing Student Capital	8
2.1.3 The Shape of Inequality.....	9
2.1.4 Models for Student Capital Distribution	9
2.2 Results	12
2.2.1 Testing the Models	13
2.2.2 Student Capital as a Finite Resource	17

2.3 Discussion	19
2.4 Materials and Methods	21
2.4.1 Data.....	21
2.4.2 Statistical Analysis	22
Chapter 3 How the Term <i>White Privilege</i> Affects Participation, Polarization, and Content in Online Communication.....	24
3.1 Introduction	24
3.1.1 How people respond to white privilege.....	25
3.1.2 Online conversations	26
3.2 Study design	29
3.2.1 Respondents.....	30
3.2.2 Instrument.....	30
3.2.3 Coding for stance and frames	32
3.2.4 Frames	33
3.2.5 Comparing frames	35
3.3 Experiment A results:.....	36
3.3.1 Avoidance.....	36
3.3.2 Stance	36
3.3.3 Response Quality.....	37
3.3.4 Frames	37
3.3.5 Composition of responses.....	39
3.3.6 Avoidance differences between whites	40
3.4 Experiment B results	41
3.4.1 Avoidance.....	42
3.4.2 Stance	43
3.4.3 Response quality.....	43

3.4.4 Frames	43
3.4.5 Composition of responses.....	43
3.5 Summary of Results	44
3.6 Discussion	44
3.6.1 Limitations.....	45
3.6.2 Implications	46
3.7 Conclusion to Chapter Three.....	48
Chapter 4 Modeling the Accumulation of Success.....	49
4.1 Introduction	49
4.1.1 Literature Review	49
4.2 Model	51
4.2.1 A Thought Experiment	51
4.2.2 Model Limitations	53
4.2.3 The Accumulation Model.....	55
4.2.4 Basic Properties of the Accumulation Model.....	57
4.2.5 Stories All the Way Down.....	58
4.2.6 Inequality in an Accumulation Model.....	58
4.3 Reduction to One Dimension	60
4.4 One-Dimensional Accumulated Distributions	63
4.4.1 General Features of 1-Dimensional Accumulated Distributions.....	64
4.4.2 Characterizing 1-D Accumulated Distributions	66
4.4.3 Dictionary of Distributions.....	71
4.4.4 Dictionary of Distributions for $\mathbf{a} \sim \mathbf{1}$	72
4.5 Examples	74
4.5.1 Numerical Simulation: Accumulated Bernoulli	75

4.5.2 Numerical Simulation: Accumulated Normal Distribution.....	79
4.5.3 Numerical Simulation: Accumulated Pareto	80
4.5.4 Numerical Simulation: Accumulated Exponential	81
4.5.5 Explicit Distribution Functions of Accumulated Exponential	82
4.6 Conclusion.....	85
4.6.1 Future Work.....	86
4.7 Proofs of Major Theorems	87
Chapter 5 Systemic Inequality	106
5.1 Introduction	106
5.1.1 Understanding Systemic Inequality.....	106
5.1.2 Hidden Variables	107
5.1.3 Reinforcing Networks	108
5.1.4 Progression of Inequality over Time	109
5.1.5 Inequality Between Groups	110
5.1.6 Models of Inequality.....	111
5.2 Systemic Inequality Arising Over Time.....	113
5.2.1 Statistics of Systemic Inequality Change	114
5.2.2 The Causes of Success Need Not be the Indicators of It.....	117
5.2.3 Increasing relationship between causal factors increases inequality.....	120
5.2.4 Adding a new factor that contributes to success	123
5.3 Interventions	127
5.3.1 Model.....	127
5.3.2 A Note about the Difference between Gaps and Inequality	129
5.3.3 Targeted Interventions.....	130
5.3.4 Single vs Multifactor Interventions	137

5.4 Conclusion.....	145
5.5 Appendix to Chapter 5	146
Chapter 6 General Discussion.....	154
6.1 Implications for Policymakers and Interventionists.....	155
6.2 Implications for Researchers.....	156
6.3 Conclusion.....	159
References.....	160

List of Tables

Table 1: Demographics of respondents	30
Table 2: Experiment A - Likelihood of responding, stance, and response quality by treatment group and race	36
Table 3: Experiment B - Probability of responding, stance, and response quality by treatment group and race	42
Table 4: Cumulants of Accumulated Distributions	72
Table 5: Approximate Cumulants of Accumulated Distribution for $a \sim 1$	73
Table 6: Initial values of knowledge, wealth, and the hidden factor	132

List of Figures

Figure 1: (top) The probability distribution function for each model. **(bottom)** The hazard rate of dropping out specified by each model. Specifically, the vertical axis gives the probability that a student who has k units of student capital will stop their education before earning $k + 1$ units. In both cases, the trends suggested are qualitative, designed to show the shape of the distribution rather than any specific numbers. 10

Figure 2: Distribution of Credits Earned Each graph corresponds to the distribution of students in one college in the dataset within 5 years of enrolling. White bars represent students who dropped out. 13

Figure 3: QQ plots for the three parametric models. Each plot compares a real data set to a simulated population generated using a fitted model. The columns correspond to discrete normal, geometric, and zeta distributions, respectively. The first three rows each correspond to a single college-year cohort. The bottom row infers the distributions for the complete set of 156,712 students from all 28 colleges. Points close to the red line indicate that the quantiles of the simulated data are very close to the quantiles of the actual data, signifying that the model fits the data better. 15

Figure 4: Further analysis of the finite resource model. (A) Histogram of average student capital. Each data value is the average student capital of a single college-year cohort. (B) Comparison of actual dropout rates and the dropout rates estimated by the finite resource model ($n = 140$). Each point corresponds to one college-year cohort. The red line corresponds to both values being equal. The blue line is the line of best fit. 18

Figure 5: Percentage of responses in Experiment A that used each frame. Squares give the proportion of responses that used a given frame, among all responses that supported renaming buildings. Diamonds represent frame use among all responses that opposed renaming buildings. Starred frames were categorized as low-quality. 38

Figure 6: Composition of posts in a hypothetical online conversation among 100 responders who are representative of our sample. For Experiment A, the figure represents likely responders. For Experiment B, the figure represents those who responded to the renaming-buildings question. Shape corresponds to the race of each responder. Points are colored based on support for renaming buildings. The Other category includes responses that were neutral, unclear, or said that it should depend on the situation. 40

Figure 7: Average self-reported likelihood of responding in Experiment A. Respondents rated their likelihood of responding on a scale from 2 = very likely to respond to -2 = very unlikely to respond. Error bars represent standard errors. 41

Figure 8: Example matrix of interactions between social skills, professional network strength, and knowledge. Numbers are for demonstration only.	52
Figure 9: Exogenous variables (blue) may be uninfluenceable (or minimally influenceable) by other variables in the model.	55
Figure 10: The relationship between a and its effect on inequality.	60
Figure 11: Income distribution in the United States for individuals between ages 25 and 65. Data from 2019 American Community Survey.....	75
Figure 12: Accumulated Bernoulli YT for various values of T , with $a = 1.01$ and $p = 0.1$	76
Figure 13: Accumulated Bernoulli YT for various values of T , with $a = 1.05$ and $p = 0.1$	77
Figure 14: Accumulated Bernoulli YT for various values of T , with $a = 1.2$ and $p = 0.1$	78
Figure 15: Accumulated Bernoulli Y^∞ for various values of a , with $p = 0.1$. Plotted with a simulated probability distribution function.....	79
Figure 16: Accumulated normal distribution YT for various values of T with $a = 1.1$, $\mu = 1$, $\sigma = 0.2$	80
Figure 17: Accumulated Pareto distribution YT for various values of T with exponent 2 and $a = 1.1$. Plotted with a simulated probability distribution function.	81
Figure 18: Accumulated exponential distribution Y^∞ for various values of a with constant $s = 1$. Plotted with a simulated probability distribution function.	82
Figure 19: Graphical representation of systemic inequality	109
Figure 20: Distribution of Income (A) Earnings distribution for adult Americans. (B) Earnings distributions and medians for white and black Americans. Data taken from 2019 American Community Survey, and includes all people age 25-65 with positive earnings.	111
Figure 21: Summary statistics of accumulation models with different growth rates over time. (A) Median values as a varies. (B) Standard deviation of outcomes. (C) Relative social mobility, as measured by the probability that an individual below the median at time t would have a value above the median some time before $t + 50$. (D) Inequality, as measured by the coefficient of variation.	115
Figure 22: Causal relationship between social skills, professional network, and knowledge in a toy model.	118
Figure 23: Distribution of correlations between causal weights u and relative outcome size u of randomly generated exp-exp matrices. 10,000 matrices were drawn from <i>ExpExp</i> .04, 40. For each matrix, the correlation between the dominant left and right eigenvector was calculated.	120

Figure 24: Relationship between the average effect of each factor on another factor and inequality. The blue line gives the mean, and the gray bar gives the standard deviation. Both mean and standard deviation were calculated using a sliding window using a total of 30,000 iterations..... 121

Figure 25: Relationship between the average effect of each factor on another factor and growth. The blue line gives the mean, and the gray bar gives the standard deviation. Both mean and standard deviation were calculated using a sliding window using a total of 30,000 iterations. 122

Figure 26: Graphical representation of the accumulation network with social skills, professional network strength, knowledge, and an added new factor (red) which we can call computer skills. 123

Figure 27: Effects on Inequality of a New Factor (A) Effect of new factor on overall inequality at time $t = 50$ as a function of the factor’s effect as a proportion of the effects of the other factors. (B) Systemic growth multiplier a . (C) Inequality in short-term investments $cvu \cdot X$. (D) Percent change of overall inequality broken up in terms of the two factors $a - 1$ and $a + 1$ and $cvu \cdot X$. The y-values on the purple curve are the product of the y-values from the red and blue curves. 125

Figure 28: Effects of a uniform intervention on two individuals...... 131

Figure 29: Effects of the health intervention on Emmy, Gregor, and Johann...... 133

Figure 30: Effects of interventions on the average well-being gap between Group High and Group Low. Each gray distribution represents 1000 simulations. Positive y-values correspond to interventions that increase inequality. The two graphs on the top correspond to the case where interventions are given to all individuals, while the graphs on the bottom represent the case where interventions are targeted only at the disadvantaged Group Low. The left graphs correspond to where the variables that cause inequality are hidden, so interventions focus on a single known variable where both groups are similar. In the right graphs, interventions focus on one of the dimensions where Group High is better off. Inequality becomes more larger and more systemic as the number of inequality-generating variables increases. 135

Figure 31: Optimal interventions when success is dependent on two factors. The blue line is the equilibrium value. Arrows represent optimal intervention types for points in different regions. 138

Figure 32: Relationship between α_j and the factor $1\alpha_j - 1\alpha$ in the proportional loss equation. This graph assumes that $\alpha = 1$. A similar shape holds for other values of α 141

Figure 33: Two heterogeneous groups of people. Group A has knowledge, but less wealth. Group B is more balanced in their relative amounts of wealth and knowledge. 142

Figure 34: Comparative return to single-factor vs multi-factor interventions as a group moves closer to equilibrium. The x-axis describes the distance between the centroid of the data

and the equilibrium line. The y-axis represents the increase in well-being divided by the size of the investment. 143

Abstract

The reasons for success and well-being, and therefore of inequality, are numerous. A large number of systemic factors influence who succeeds or fails, who is healthy or sick. These factors, many of which are hard to understand or measure, interact in important and surprising ways. In contrast, humans tend to choose simple explanations of the world. Early humans evolved heuristics to help them effectively navigate a highly social, low-data world. Researchers choose intuitive, low-dimensional models of complex problems. While these approaches have provided humanity with tremendous benefits, they can sometimes fall short when faced with the systemic problems of the modern day. Technology plays a central role here, since tech can drive economic growth & improvements in health, as well as online polarization. This dissertation takes an interdisciplinary approach to examine the multifaceted nature of inequality in four studies. The first study presents a conceptual framework for the total effect of these factors on education, develops a novel statistical method, and uses administrative data to show universal patterns in community college students' ability to be successful. The second study is an experiment involving racially charged language that shows how individuals' tendencies toward simple explanations for inequality can combine with technology to create polarization and decrease open democratic discussion. The third study uses mathematics and statistics to develop a causal network model of multifactor cumulative (dis)advantage. The fourth study develops a conceptual framework for systemic inequality, and then uses models to explore the relationship between the causes of inequality and principles for designing effective interventions.

Chapter 1

Introduction

Researchers have found many interacting factors influencing different forms of success (as in the case of education) or well-being (as in the case of health). Since inequality is the unequal distribution of success or well-being, the causes of inequality are also typically manifold and complex. These causal factors interact, since resources that can be used to increase one dimension of well-being are also likely to be able to influence other dimensions, either directly or indirectly. In addition, due to their idiosyncratic nature, many of the causal factors and interactions are hard for both researchers and laypeople to understand and study. However, humans benefit from being able to make sense of the world. So we create low-dimensional representations of a high-dimensional, and relatively data-scarce, world. Evolution and individuals have come up with a variety of such approaches including heuristics (Gigerenzer & Gaissmaier, 2011), sense-making (Chater & Loewenstein, 2016), coarse-graining (Saunders & Voth, 2013), and theory development (Shalley, 2012). These strategies have led to significant benefit and improvements in well-being. However, they can also sometimes be maladaptive in the modern, technology-laden world where inequalities are often entrenched.

The studies in this dissertation use different methods to approach this idea. Chapter Two uses low-dimensional administrative data and a novel conceptual framework to show that universal forces govern community college student success. Chapter Three is experimental, showing how our evolutionary heuristics can be counterproductive and polarizing when combined with the systemic affordances of the internet. Chapter Four generates a mathematical and statistical model for understanding how success (or the lack thereof) can grow through a multifactor cumulative (dis)advantage process. Chapter Five uses models to draw general conclusions about systemic inequality and the interventions that aim to address it.

1.1 Chapter Two – Patterns of Student Success are Universal and Suggest a Latent Limited Resource

Hundreds of thousands of students drop out of school each year, despite billions of dollars of funding and myriad educational reforms (Shapiro et al., 2018; U.S. Department of Education., 2017). Existing research tends to look at the effect of easily measurable student characteristics. However, a vast number of harder-to-measure student traits, skills, and resources affect educational success (J. Johnson & Rochkind, 2009; Pascarella & Terenzini, 2005; Porchea et al., 2010). Furthermore, community colleges often refer to easily measurable variables when measuring student success, such as term-to-term retention or completion of college-level math & English. This can lead to a tail-wagging-the-dog approach to decision-making, where interventions focus on improving measurable metrics rather than improving students' ability to be successful (Wood et al., 2019).

Chapter Two addresses these ideas by presenting a conceptual framework for the cumulative effect of all factors influencing student success. We call this quantity *student capital*. The framework assumes that degree-seeking students apply their resources to achieving their goal in their own way and on their own time, rather than using measurable quantities chosen by academics or administrators. We develop a method for estimating student capital in groups of students and find that student capital is distributed exponentially in each of 140 cohorts of community college students. Students' ability to be successful does not behave like standard tests of intelligence. Instead, it acts like a limited resource, distributed unequally. The results suggest that rather than removing barriers related to easily measured characteristics, interventions should focus on building up the skills and resources needed to be successful in school.

1.2 Chapter Three – Conflicts in Understanding of Language on Social Media Can Be Counterproductive and Increase Polarization

The language used in online discussions affects who participates in them and how they respond (Zhu et al., 2017), which can influence perceptions of public opinion (McGregor, 2019; Neubaum & Krämer, 2017). However, language itself is a dimensional reduction technique, facilitating the understanding and communication of complex topics. Furthermore, language use

varies between geographic regions, demographic groups, and political factions. Even definitions of words can vary between groups (Banton, 2015). The internet brings these diverse groups people together, creating potential misunderstanding and conflict. The term *white privilege* is a useful instrument for understanding these effects, since its definition encompasses multiple factors (McIntosh, 1990) and because it targets a group identity which is personal for many people.

Chapter Three examines how the term *white privilege* affects communication on social media. In two lab experiments, US residents were given a chance to respond to a post asking their opinions about renaming college buildings. Using the term *white privilege* in the question decreased the percentage of whites who supported renaming. In addition, those whites who remained supportive when *white privilege* was mentioned were less likely to create an online post, while opposing whites and non-whites showed no significant difference. The term also led to more low-quality posts among both whites and non-whites. We find evidence that the effects of the term *white privilege* on the content of people's responses is primarily affective – based in feelings. Overall, use of the term *white privilege* seems to create internet discussions that are less constructive, more polarized, and less supportive of racially progressive policies. The findings suggest that individuals who hope to gain broad support, promote meaningful conversations between diverse parties, and reduce polarization should try to use language that creates a shared understanding across groups

1.3 Chapter Four – Success Arises from Many, Interacting Factors

Even though success and well-being arise from many factors, these factors are often hard to measure. There is solid research examining the effects of various idiosyncratic factors on well-being, such as mindsets (Yeager & Walton, 2011), dominance (Hawley, 2002), or conscientiousness (Bogg & Roberts, 2004). However, these studies are costly, and typically use methods designed to isolate the effects of the studied factor. In reality, the causes of success and well-being are manifold, interrelated, and hard to study together (Blau & Duncan, 1967; Hovmand, 2014; Levy et al., 2020).

Chapter Four builds on the idea of causal networks to develop a model of multi-factor well-being as an accumulation process. It assumes that, when individuals can use a resource to improve their

other resources, some will do so. This creates a situation where positively construed traits, skills, and resources are often mutually causally related in a complex web. I build a network of these relationships, and examine how growth in each of the factors happens over time. The contributions of this chapter are conceptual, mathematical, and methodological. I present a novel way of understanding how well-being changes through a high-dimensional cumulative advantage (or disadvantage) process. I prove theorems about the model, showing that this process leads to stable one-dimensional distributions. I create a method to connect the model to large, one-dimensional datasets – a process which can lead to deeper understanding about the process that generated the data.

1.4 Chapter Five – Systemic Inequality

Inequality between and within groups is often hard to change. Despite simple explanations of this persistence, many forms of success or well-being arise from a large number of related skills, traits, and resources, including one's ability to participate in social & economic institutions (Braveman et al., 2022). Being disadvantaged in one way can make it harder to succeed in other ways, as that disadvantage makes everything else just a little bit harder (Dannefer, 2003; Torche, 2018). This idea has been studied in a variety of specific contexts (Dannefer, 2003; Kraus & Park, 2017; Lorenc et al., 2013; McLanahan & Jacobsen, 2015). However, studying any particular context is costly and academics tend to be siloed. So there is value in attempting to provide more generalizable results about systemic inequality.

Chapter Five presents a general conceptual framework for systemic inequality and uses two models to draw general principles about how systemic inequality grows and how best to tailor interventions. I define systemic inequality as arising from many interacting causes which are often hidden from analysts. I provide examples of systemic and non-systemic inequality, and give historical examples of how non-systemic inequality evolved into systemic inequality over time. I then use the accumulation model from Chapter Four to (a) show the causes and indicators of success need not be the same, (b) demonstrate that making some factors more useful can increase systemic inequality, and (c) explore how adding a new potential cause of success might either increase *or* decrease inequality. To explore the effects of interventions on high-dimensional inequality, I use a Cobb-Douglas utility model as well as more general mathematics.

I demonstrate, using both simple examples and simulations, how the best interventions for addressing systemic inequality are targeted towards specific groups and involve many factors at once.

Chapter 2

The Shape of Educational Inequality

2.1 Introduction¹

Over 500,000 high school students and over 600,000 college students drop out of school every year (Shapiro et al., 2018; U.S. Department of Education., 2017). Practitioners, researchers, and pundits have proposed a variety of explanations for why so many students are unable to achieve their goals. However, despite a variety of different education policies and billions of dollars spent ensuring no children are left behind, millions of children and adults are unable to achieve their academic goals. Unfortunately, there is no simple explanation that can point to simple interventions. The process of becoming successful in school can be complicated and difficult, requiring the right combination of social, personal, academic, and financial traits and skills. Researchers and policy makers have not yet found the secret to consistently cultivating success in students. So it is no surprise that so many students are unable to successfully navigate the educational system.

2.1.1 Student Capital

In this paper, we present a conceptual framework for studying students' capacity to be successful in school, which we call student capital. We also demonstrate an analytical method for measuring this quantity in community college students. Broadly speaking, we define *student capital* as the cumulative amount of resources a student can bring to bear to be successful in a particular school context. These resources might come in many forms, such as economic resources, social, cognitive, non-cognitive and academic skills. Like other forms of capital, these resources both help drive students toward their goals and insulate them against the random shocks that affect all of us. For example, consider the unlucky situation of a commuter college

¹ A version of this chapter was published as:
Quarles, C.L., Budak, C. & Resnick, P. (2020) The shape of educational inequality. *Science Advances*. 6(29). doi: 10.1126/sciadv.aaz5954

student whose car has broken down, and therefore may have to miss class. If the student has a supportive social network, then she might be able to catch a ride with a friend. Strong academic skills might ensure that missing a class does not affect her learning. Cultural capital and self-confidence might give her the ability to communicate clearly with her instructor to minimize any effects on grades. Economic resources might have allowed her to live in a campus dormitory, possibly at a more elite college, thus avoiding the situation in the first place. And of course the more resources she has, the better off she will probably be.

There is a rich body of literature examining the factors that affect student persistence in college. An incomplete list of personal factors related to college success includes academic preparation (Long et al., 2012; Scott-Clayton et al., 2014), students' self-discipline, self-confidence, commitment to college, amount of social activity, race, age, full-time/part-time enrollment, degree expectations, distance between home and school, number of hours worked at a job, parents' income, parents' education level (Porchea et al., 2010), perceptions of faculty, peer groups, campus engagement (Pascarella & Terenzini, 2005), familial responsibilities, interest in school, lack of money (J. Johnson & Rochkind, 2009), unreliable housing, and food insecurity (Goldrick-Rab, 2018). In addition, the skills required to be successful in a classroom can be confusing and vary from class to class (Boaler, 1998; Cox, 2015), with some researchers calling classrooms a "black box" (Cuban, 2013; Grubb, 2001). So college success may be partly attributable to a student's ability to learn new classroom expectations. Many researchers have attempted to disentangle this web of causal relationships (Mayhew et al., 2016). They face significant challenges from selection bias and hidden variables. Our goal is not to wade into that discussion, but to consider the cumulative effects of all factors as a single variable.

We call this quantity student capital to fit with prior literature on other forms of capital, which can be used for both a source of investment income and a reserve of resources to insulate against shocks. Social scientists have examined a variety of forms of capital, including economic, social, and human capital. Social capital is some form of social relationship that can be used to benefit individuals or groups, such as membership in neighborhood groups (Putnam, 2001), networks of parents (Coleman, 1988), or professional connections (Granovetter, 1983). Human capital is the set of skills embodied in a given workforce, typically measured in terms of economic benefit. For example, researchers have examined the wage benefits of specific skills, on-the-job training, and a variety of college degrees (Becker, 1994; Castex & Dechter, 2014;

Deming, 2017). Different forms of capital can manifest in multiple ways, and there have been many debates about how to measure them (Becker, 1994; Coleman, 1988; Killewald et al., 2017; Mulder et al., 2009; Putnam, 2001). Economic and social capital are most often defined in terms of what they consist of: money and social relationships, respectively. In contrast, human capital is usually defined in terms of economic benefit, and can consist of many types of embodied skills and traits. Student capital is more like human capital, since it is defined by the educational success it can provide.

2.1.2 Operationalizing Student Capital

We consider student capital as the amount of success that a student is *able* to achieve. In our study, this is the number of credits they *could* earn if that many credits were required for their goals. To use a metaphor with more financial forms of capital: In an economic system without a standardized currency, an individual's wealth can only be determined by what it can be traded for. Depending on the circumstances, a bag of gold might be worth more or less than a bag of rice. Similarly, a supportive family or good social skills may have a varying effect on a student's outcome, depending on a variety of situational factors. Furthermore, those factors might interact in a way that helps or hinders the student. We operationalize student capital as the total amount of educational success (credits completed) that can be "bought" by a student, in their particular context, using their skills, traits, and resources. This runs the risk of conflating student capital with the *returns* on student capital. However, given the striking universalities in our results, we think this metric measures something meaningful.

It is important to distinguish student capital from student outcomes, which might include whether a student graduated or transferred to a four-year school. Outcomes are measurable representations of whether a student reached a certain goal, rather than giving a measure of how well they *could have* done. Student capital is harder to measure in individuals. The student capital of students who have dropped out of school can be directly observed as the number of credits they earned. However, students who graduated or transferred may not have run out of student capital. We only know their capacity to earn credits is greater than or equal to the number they earned. This is good for those students, but makes data analysis more challenging. It makes it impossible to measure the student capital of every student. Instead, we can estimate the

distribution of student capital in a *group* using right-censored maximum likelihood estimation. This is still useful, since statistics like this are commonly used to describe groups of students.

2.1.3 The Shape of Inequality

Despite humankind's best efforts, inequality has always been with us (Milanovic et al., 2011). However, the amount of inequality has varied, depending on the era and location (Killewald et al., 2017; Milanovic et al., 2011). This suggests that, beyond the idiosyncratic forces unique to specific groups, there are systemic forces that keep resources allocated unequally. To better understand those forces, we look at the *shape* of educational inequality. In a more economic context studying the shape of inequality might involve examining income or wealth distributions (Nirei & Souma, 2007; Tao et al., 2019). In the context of education, we look at the shape of student capital distributions. If we find that these distributions have the same universal shape across colleges, then this will give us insight into the underlying macro-scale processes that create educational inequality.

To analyze the shape of educational inequality, we used data from 156,712 students from 28 Washington community colleges. We grouped students into 140 cohorts, all of whom started at the same college during the same academic year. We focus on degree-seeking community college students who aim to transfer to a four-year college. This group has the benefit of being fairly diverse (Goldrick-Rab et al., 2017), while sharing the same goals and educational context. This allows us to measure their student capital on the same credit-based scale. For each student, we calculated the total number of community college credits they had earned within five years of enrolling, and whether the student dropped out of school without earning a degree or transferring to a four-year college.

2.1.4 Models for Student Capital Distribution

To explore the shape of educational inequality, we consider a number of different models for how student capital might be distributed. Each model represents a universe with plausible educational behavior that leads to a particular distribution of student capital. Graphical comparison of the models is shown in Figure 1. In the next section, we test whether these models fit the data.

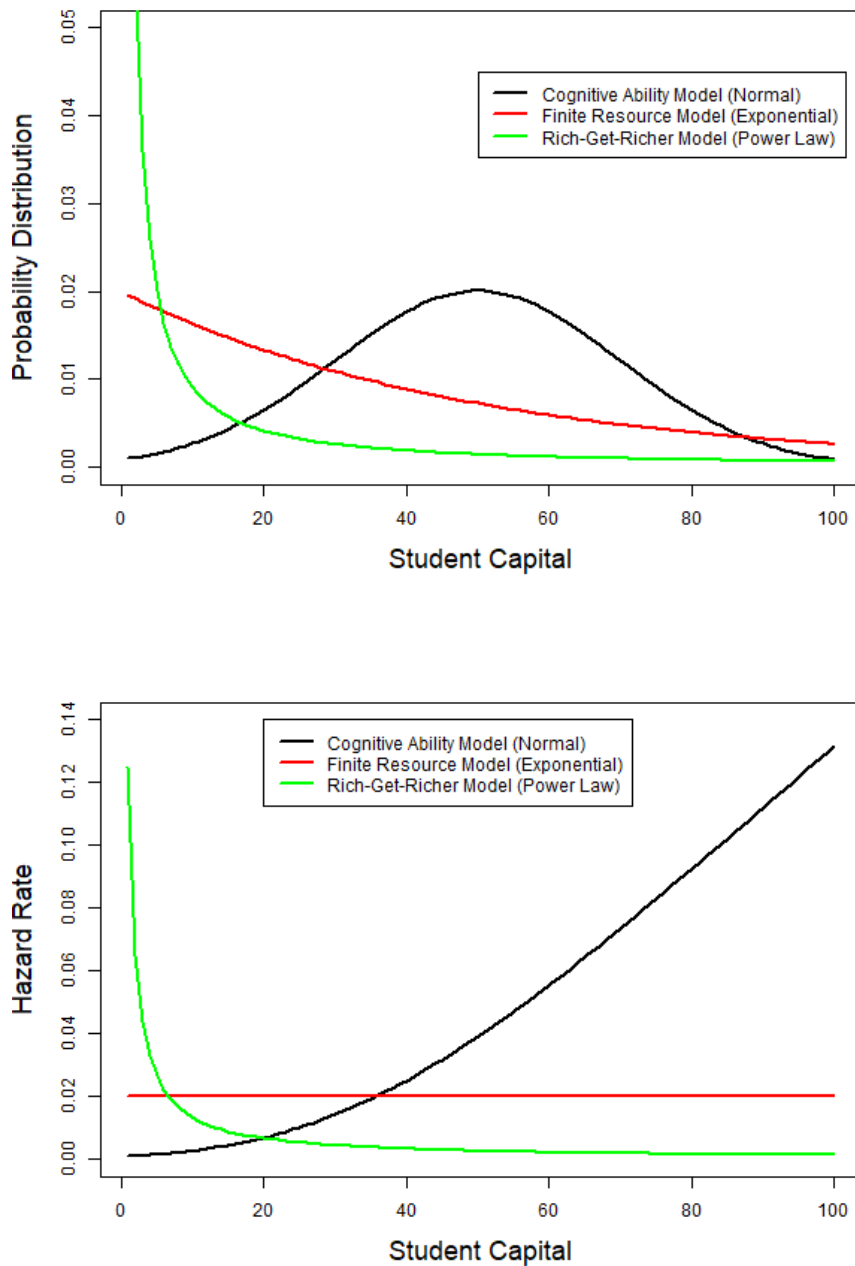


Figure 1: (top) The probability distribution function for each model. **(bottom)** The hazard rate of dropping out specified by each model. Specifically, the vertical axis gives the probability that a student who has k units of student capital will stop their education before earning $k + 1$ units. In both cases, the trends suggested are qualitative, designed to show the shape of the distribution rather than any specific numbers.

The *cognitive ability model* comes from the claim that educational outcomes are largely determined by, or equivalent to, cognitive ability as measured by achievement tests such as IQ

(Herrnstein & Murray, 1994). This model is consistent with the common practice in the education literature of using standardized tests such as the SAT as measures of ability or achievement (Dillon & Smith, 2017; Sirin, 2005). These tests measure specific cognitive abilities or knowledge at the time the student takes the test. The cognitive ability model assumes that student capital, which is nominally a student's ability to navigate successfully through the complex social, personal, and academic demands of a school system, is mostly dependent on IQ or cognitive abilities. Since these cognitive abilities, as measured, tend to be normally distributed, this model suggests that student capital might also be shaped like a bell curve.

It could also be possible that, like forms of financial capital, a given community has a limited amount of student capital that can be generated in their college-going population. Collectively, parents, family, and friends may have a finite amount of experience, noncognitive skills, social stability, and financial resources to share with children and college-bound adults. Some communities, particularly wealthier and more educated ones, have more of this resource than others. This is consistent with the well-established fact that children of wealthier and more educated parents tend to have more of the skills necessary for academic success (Lareau, 2011; Porchea et al., 2010; Sirin, 2005). The *finite resource model* assumes that the only thing constraining students' capacity to complete college is the limited nature of this resource in a population. For this model, we assume that resources are at least partially substitutable, so that we can treat these resources as coming from a single pool. For instance, a student whose home is too unstable to study at may be able to spend money to work at a coffee shop. If society distributes this finite student capital in the least informative way, we would expect to see an exponential distribution of student capital for a given community. Put another way: If, in a given population, the only major limitation is that student capital is finite, then there are many ways that it could be distributed to individuals. However, in this case the vastly most probable distribution is exponential – or something very close. More details and examples of the principle of maximum entropy, which underlies this model, can be found in (De Martino & De Martino, 2018; Harte, 2011; Montroll, 1981; A. G. Wilson, 1970). A similar model was proposed in (Drăgulescu & Yakovenko, 2000) to explain why income distributions between the 10th and 90th percentiles are distributed exponentially (Nirei & Souma, 2007).

The *rich-get-richer model* assumes that the student capital gained from a new resource is roughly proportional to the student capital they already have. For instance, a student with good

study skills might be able to benefit more from increased wealth, because they might be able to use the time not working at a job to study more efficiently. The rich-get-richer phenomenon has been well-studied in a variety of other areas (Easley & Kleinberg, 2010; Muchnik et al., 2013; Yakovenko & Rosser, 2009) and leads to a heavy-tailed distribution such as a power law or log-normal distribution. Since these distributions often have similar behavior and can be difficult to distinguish from each other, we focus on whether the distribution of student capital fits a power law.

Of course another mental model, often implicitly assumed among those who do educational interventions, is that (a) interventions and college policies can have a significant effect on student progress at various points in the college process, and (b) the policies and supports in different colleges vary enough to see this effect. If this *context-specific model* were true, we would expect distributions of student capital to have varying, idiosyncratic shapes depending on the school itself and perhaps even the year. For example, a college with a strong student onboarding program might have a mode at 15 or 30 credits, while other colleges with regular enrollment cycles might have periodic distributions of student capital. In this case, institutional structures would be more important for student success than the resources, skills and traits student brought with them. Student capital would be a relatively unimportant consideration in educational success. This model is not shown in Figure 1, because the context-specific model would imply that each cohort of students has its own distinctive curve.

2.2 Results

Unfortunately, we cannot directly measure the number of credits that every student *could have* earned. Instead, we only have data for the number of credits students *actually* earned, and whether they dropped out, graduated, or transferred. Figure 2 shows the distribution of credits, graduation, and transfer for two colleges. White bars represent students who dropped out, so that their observed number of credits is equal to their student capital. Blue, green, and yellow students represent censored data points. These individuals' student capital is greater than or equal to the observed number of credits shown on the graph. The number of successful students peaks around 90-100 credits, because associate's degrees in Washington require at least 90 credits. Note that Figure 2 does not show student capital, just the observed number of earned credits. Most graduating/transferred students will have student capital values larger than the number of

credits they earned. So we can imagine what the student capital of successful students might look like by flattening the colored bars to the right. The distribution of student capital might look like Figure 2, but with a smaller bump. Or it could be continually decreasing, so that the number of students who have k credits of student capital decreases as k increases.




Figure 2: Distribution of Credits Earned Each graph corresponds to the distribution of students in one college in the dataset within 5 years of enrolling. White bars represent students who dropped out.

2.2.1 Testing the Models

The cognitive ability, finite resource, and rich-get-richer models each assumed that the distribution of student capital follows a given parametric model: normal, exponential, or power law. So we explored them all using the same approach. We assumed that student capital is distributed according to the specified model, with a censoring process corresponding to graduation/transfer which is estimated individually at each credit level. We used right-censored maximum likelihood estimation to estimate the parameters for each model. We then examined goodness of fit for each model using both the Akaike information criterion (AIC) and quantile-quantile (QQ) plots.

AIC is a standard information-theoretic method for comparing distribution fit. If a model has K parameters and log-likelihood \mathcal{L} , then $AIC = 2K - 2\mathcal{L}$. We used AIC to compare the fit of the three parametric models on each of the 140 cohorts. The finite resource/exponential model gave the best fit to the data on every cohort. So the inferred distribution of student capital fits an exponential distribution better than a normal or power law distribution.

While the AIC analysis shows which of the chosen distributions is better, it does not show if that fit is *good*. To qualitatively examine goodness-of-fit, we used QQ plots. To generate a QQ plot for a given cohort of students, we fit the parameters for each of the three models and then used those parameters to generate a set of simulated students. We then compared the distribution of simulated students' credits earned to the actual distribution of credits earned. Figure 3 shows QQ plots using this process for three representative cohorts and also for the combined set of all students.




Figure 3: QQ plots for the three parametric models. Each plot compares a real data set to a simulated population generated using a fitted model. The columns correspond to discrete normal, geometric, and zeta distributions, respectively. The first three rows each correspond to a single college-year cohort. The bottom row infers the distributions for the complete set of 156,712 students from all 28 colleges. Points close to the red line indicate that the quantiles of the simulated data are very close to the quantiles of the actual data, signifying that the model fits the data better.

The finite resource model fit the data remarkably well. Using both QQ plots and AIC, this model fit best across colleges and across years. After accounting for the censoring effect of students graduating and transferring, student capital seems to follow an exponential distribution.

The rich-get-richer model fit the data very poorly. The power law seems to expect more students dropping out early and more students with very high student capital than is found in the real data. The heavy-tailed behavior found in the power law distribution is inconsistent with our data, suggesting that other heavy-tailed distributions such as log-normal would be poor fits as well.

Because students in our data set only earned positive integer numbers of credits, we used a truncated discrete normal distribution for the cognitive ability model. At first glance, the cognitive ability model seems to fit the data almost as well as the finite resource model. The QQ plots for the normal distribution are reasonably close to the diagonal. However, the results were not consistent with what we would consider a normally distributed population. In such a population, the mean will be between the minimum value and the maximum value. However, for every cohort in our data set, the inferred mean was $\hat{\mu} = 1$. This was the minimum possible number of credits earned, and also the minimum allowed $\hat{\mu}$ using our algorithm. Inferred standard deviations $\hat{\sigma}$ were distributed between 90.5 and 169.0 credits ($\text{mean}(\hat{\sigma})=120.1$, $\text{sd}(\hat{\sigma})=12.7$). This inferred standard deviation is larger than the number of credits earned by most degree-receiving students. These pathological results are consistent with a continually decreasing probability distribution of student capital. The best way to fit a normal distribution to a decreasing distribution is to just fit the right tail. The resulting simulated data is missing the characteristic bell curve shape of the normal distribution. Therefore, we can not say that the cognitive ability model is supported by our results.

The context-specific model assumes that the shape of the distribution of student capital is highly dependent on the college. Evidence for this model would involve very different distributions of student capital, with some colleges having high dropout rates for students with low numbers of credits, and others having high dropout rates at higher credit levels. However, our previous analysis shows that distributions of student capital, across years and across colleges, all fit an exponential model very well. It seems that colleges do not have a significant impact on the *shape* of educational inequality. At every college, there are more low-resourced students than high-resourced students.

2.2.2 Student Capital as a Finite Resource

We now explore the finite resource/exponential model in more depth. Colleges often want to compare the experiences of different groups of students. Because the exponential distribution can be uniquely characterized by a single parameter, we can use our model to assign a number to any group of students. This number can be used like any other statistic, such as graduation rate. One possible such parameter is the per-credit retention rate q , which is one minus the traditional exponential decay rate. For example, one college might find that 95% of their students, at any credit level, will take one more 5-credit class. This corresponds to $q^5 = .95$, or $q = .9898$. Another such parameter is the mean of the distribution $\mu_S = \frac{1}{1-q}$. This has the units of credits, and is reasonably easy to interpret as the average student capital in the student population. Equivalently, $n\mu_S$ is the total amount of student capital collectively possessed by a group of students. Both q and μ_S can be easily inferred with the algorithm we used. Figure 4(A) shows the distribution of average student capital μ_S for the 140 cohorts in the dataset. Student populations in most of the colleges we studied have an average student capital between 90 and 130 credits, with a peak around 110 credits. It may seem surprising that most students drop out of school, given that the average student capital in most cohorts is larger than the typical 90-100 credits required for an associate's degree in Washington. The high dropout rate comes from the fact that the exponential distribution is right-skewed. Some students would be able to achieve very high levels of education, which pulls up the average student capital but only increases the number of graduates by one. Note that the values in Figure 4(A) are specific to Washington state community colleges. Schools that measure credits differently, such as those on a semester system, will not be able to compare their average student capital with Washington's quarter system.




Figure 4: Further analysis of the finite resource model. (A) Histogram of average student capital. Each data value is the average student capital of a single college-year cohort. (B) Comparison of actual dropout rates and the dropout rates estimated by the finite resource model ($n = 140$). Each point corresponds to one college-year cohort. The red line corresponds to both values being equal. The blue line is the line of best fit.

Typically, education regression models include R^2 values to show how well the model explains variation in a set of data. So we calculated the amount of variance in college dropout rates explained by our model. Again, our process is: (a) Select a cohort of students. (b) Fit the finite resource model, which involves inferring the decay rate for the exponential model, and inferring the full distribution of success points. (c) Generate 10,000 simulated students using this new fitted model. (d) Compare the percentage of simulated students who dropped out with the percentage of actual students who dropped out. Figure 4(B) is a plot of these percentages, with one point for each college-year cohort. The figure shows that the estimated dropout rate is close to the actual dropout rate. However, the estimates are systematically biased so that the model estimates are systematically higher than the true values. Notably, the relationship is very strongly linear ($R^2 = .982$, $F(1, 138) = 7731$, $p < .001$), which means that the actual dropout rate could be reconstructed with high accuracy from this biased estimate. This reconstructability means that the combination of the exponential parameter and the distribution of success points contain effectively all of the information contained in college dropout rates.

This approach assumes the full distribution of success points, which involves estimating the percentage of graduating/transferring students at every credit level. In pursuit of simplicity, we repeated the process without such a strong assumption. For each cohort, we took the mean of

the success point distribution, effectively assuming that all students would graduate or transfer at the same credit value. This simplified two parameter model still explained 92.4% of the variation in dropout rates by cohort ($F(1, 138) = 1680, p < .001$).

2.3 Discussion

This paper has presented a conceptualization of student capital as a many-faceted resource, operationalized it, and shown that there is a universal shape to educational inequality. This shape suggests that, in a given population, the amount of student success is finite. The results have ramifications for how colleges think about student success and interventions. In addition, the informationally-equivalent parameters mentioned here: the per-credit retention rate q and the average student capital μ_S , might be used to compare groups of students. For instance, they could be used to compare demographic groups.

We defined student capital using an input-based approach: as the resources that students can marshal toward achieving their academic goals. In contrast, the more common practice of measuring student outcomes is an output-based approach. The cohorts in this study started in different years, and came from colleges with different policies, geographies, and populations. However, in all cases, the shape of educational inequality was the same. Student capital distributions across colleges and years were surprisingly all exponential distributions. This model explained 98% of the variation in graduation rates of cohorts in our data set.

Our explanation for this systemic inequality is that student capital is a finite resource in a given population. Society has a limited amount of student capital to distribute to the community college-going population, and distributes that capital in the least informative way possible. Student capital as a finite resource makes sense if ability to be successful in school is truly a form of capital that one gathers from parents, mentors, and friends. Throughout their life, people gain things like social skills (Lareau, 2011), academic skills (S. Reardon, 2011), emotional regulation (Morris et al., 2007), and economic resources (Pfeffer & Killewald, 2018) from their environment. Geographic areas that are less educated and poorer have fewer resources like this. So they have less ability to share that student capital with their college-bound population.

We have also discarded a number of hypotheses that are common in scholarly and popular conceptions of academic achievement. Many of the colleges were running interventions

focused on student success (Jenkins et al., 2009; Moore et al., 2013), which we might expect would change the shape of the student capital distribution. Despite their attempts, the general shape of the student capital distribution was remarkably similar across cohorts. It seems that small-scale interventions don't have a significant effect without affecting students throughout the college-going process.

Students' ability to earn college credits has a fundamentally different distribution than that of intelligence and academic achievement tests. This is consistent with previous research showing that tests of knowledge have limited relationship with more comprehensive measures of ability to be successful in school, like GPA (Quarles & Davis, 2017; Scott-Clayton et al., 2014). Even students who are academically knowledgeable are subject to different types of knowledge tests and to instructors with wildly varying grading practices (Brimi, 2011; Starch & Elliott, 1912, 1913) and pedagogical practices (Boaler, 1998; Cox, 2015). Successfully navigating school at least partially amounts to learning and adapting to the particular expectations of teachers and school bureaucracy. The results also caution researchers against cavalier use of the word *ability* to describe test scores. An individual's ability to do well on standardized tests, which might more aptly be called *cognitive ability*, is clearly not the same as student capital, the ability to complete schooling.

Nor does student capital follow the power law behavior of a rich-get-richer model. In some sense, this is unsurprising. Many of the examples we have of power law behavior, such as social media follower networks (Myers et al., 2014) and academic citation (Price, 1976) require a negligible cost for each additional unit of capital. Given that each additional college credit has, at minimum, a financial cost, we would not expect to see power law behavior in this regime. However, wealth and income distributions do have heavy tails at the high end (Drăgulescu & Yakovenko, 2001; Nirei & Souma, 2007). So it would not be surprising if there was an unobserved tail of students who had nearly unlimited ability to be successful in school.

We think that these results will be useful for the design of student success interventions. These interventions often focus on finding and reducing barriers in the college-going process. However, students face a great many barriers, most of which are outside of the college's influence (J. Johnson & Rochkind, 2009). This paper suggests that successful educational interventions should be focused on building up resources and skills in students, rather than minimizing barriers. Interventions that focus on resource-building are also likely to improve life

outcomes in the broader sense. Results from comprehensive, resource-building interventions show significant returns (Kolenovic et al., 2013). Even though these interventions are more costly, the benefit to society is lower than the cost (Levin & Garcia, 2013).

A common concept in community colleges is student momentum. Our results suggest that we might instead think of student capital as a form of energy. The exponential distribution of student capital is very similar to the Boltzmann-Gibbs distribution in physics, which has been used to study economic capital (Drăgulescu & Yakovenko, 2001). In this formulation, the average student capital μ_S is a state variable corresponding to the average energy of the students in the system. Colleges might conceptualize interventions that focus on increasing the energy of their student body.

A few notes of caution are warranted to readers trying to generalize or extend our work. For our analytical technique to work, there needs to be a sufficient number of uncensored data points to infer the distribution. These are dropouts which, sadly, community colleges have in plenty. High schools and more selective colleges likely have too few dropouts to accurately make an inference.

It is also worth emphasizing that randomness can play a significant role in a student's ability to be successful. An inspirational teacher or an unexpected financial challenge may have a huge effect on a student's outcomes. This randomness creates error in the use of credits to measure *individuals'* student capital. When looking at groups of students, this error should average out. Some people will have the inspirational teachers and some won't.

The institutional context also plays a role in student persistence and completion. For example, the skills necessary to thrive in a low-income high school may be very different from those required in an elite university. So a student who has a lot of student capital in one school may have less in another. Most differences in student persistence by college are associated with the differences between two and four-year institutions and college selectivity. After controlling for the student populations, other factors seem to have a relatively small effect (Clotfelter et al., 2013; Mayhew et al., 2016; Pascarella & Terenzini, 2005).

2.4 Materials and Methods

2.4.1 Data

We use deidentified data provided by the Washington State Board for Community and Technical Colleges (SBCTC), which included all students who started at 30 of the 31 community/technical colleges in Washington within the five year period between Summer 2006 and Spring 2011. One college declined to participate. The original dataset contained 303,390 students. To create a group of people with nominally similar goals, we only included degree-seeking students who self-identified as academic transfer students during their first quarter. We excluded reverse transfer students, dual-enrollment high school students enrolled through Washington’s Running Start program, and anyone who enrolled but earned zero credits. We also excluded two colleges that had less than 100 transfer students. The remaining colleges each had over 1000 transfer students. This reduced the dataset to 156,712 students, split into 140 college-year cohorts. Data exploration was initially performed on four of the 28 colleges. These four were chosen to have different general shapes and to have a sufficient sample size. Once the statistical methods were designed and written, we then examined the remaining colleges.

There were two main observable variables of interest. The first was x_i , the number of community college credits each student earned within five years of enrolling in the Washington community college system. We did not differentiate between credits based on when they were earned. We assume that students are putting resources into being successful in college at the rate that is optimal for them. Our other observable variable is \check{y}_i , a binary variable which describes whether a student dropped out. We say a student graduated if they earned an associate’s or bachelor’s degree in the SBCTC system within five years of initial enrollment. All degrees required at least 90 credits, though some students brought credits into the SBCTC system and graduated with fewer than 90 credits in our data. We say a student transferred if they enrolled at a four year college within five years of initial community college enrollment. Transfer data was obtained by SBCTC from the National Student Clearinghouse. A student dropped out if they did not transfer or graduate. Analysis was performed using R version 3.5.1 (R Core Team, 2018) using the VGAM package (Yee, 2018).

2.4.2 Statistical Analysis

The parametric models assume that each student has two independent latent variables: their student capital y_i , which is the number of credits they can earn before they have to dropout, and their success point g_i , the credit level where they achieve their academic goals by

transferring or graduating with an associate's degree. The observed number of credits is then $x_i = \min(y_i, g_i)$. Students who have dropped out ($\hat{y}_i = 1$) correspond to the case where $y_i < g_i$. Otherwise, $\hat{y}_i = 0$.

To test the parametric models, we assume that y_i and g_i are drawn from theoretical probability distributions, infer the parameters of those distributions, and then compare the inferred distributions with the real data. Let Y_k be the probability that a randomly drawn student will have a student capital of exactly k credits. Let G_k be the probability that a randomly drawn student has a success point of exactly k credits. This gives the likelihood function:

$$\mathcal{L} = \sum_i \left[Y_{x_i} \sum_{k=x_i+1}^{\infty} G_k \right]^{\check{y}_i} \left[G_{x_i} \sum_{k=x_i}^{\infty} Y_k \right]^{1-\check{y}_i}$$

Taking logs and simplifying gives the log-likelihood function

$$\log \mathcal{L} = \left[\sum_i \check{y}_i \log Y_{x_i} + (1 - \check{y}_i) \log \left(\sum_{k=x_i}^{\infty} Y_k \right) \right] + \left[\sum_i (1 - \check{y}_i) \log G_{x_i} + \check{y}_i \log \left(\sum_{k=x_i+1}^{\infty} G_k \right) \right]$$

Notice that the only distribution in the left sum is Y_k , while the right sum only includes G_k . So we can maximize the log-likelihood by maximizing each term separately. The distribution $\{Y_k\}$ that best fits student capital does so regardless of distribution of success point $\{G_k\}$.

We tested the three different parametric models for $\{Y_k\}$: the discrete normal distribution $Y_k(\mu, \sigma) = \frac{1}{A(\mu, \sigma)} e^{-\frac{(k-\mu)^2}{2\sigma^2}}$ where $A(\mu, \sigma)$ is a normalizing constant calculated numerically, the geometric distribution $Y_k(q) = (1 - q)q^{k-1}$, and the zeta distribution $Y_k(\alpha) = \frac{1}{\zeta(\alpha)} k^{-\alpha}$ where $\zeta(\alpha)$ is the Riemann zeta function. To validate the model, we had to create simulated students, which meant inferring $\{G_k\}$ as well. Unlike $\{Y_k\}$, where we were trying to find a simple parametric form with few parameters, we were only interested in $\{G_k\}$ as a validation tool. Since $\{G_k\}$ definitely is dependent on college policies and transfer options, we did not expect a parametric form for it. So for each value of k we inferred G_k as its own parameter.

Chapter 3

How the Term *White Privilege* Affects Participation, Polarization, and Content in Online Communication

3.1 Introduction²

Billions of people use the internet and social media as a window to the world. Rather than being made of glass, this window is manufactured and shaped by the collective choices and language of billions of people. Online behavior is shaped by a community's language (Danescu-Niculescu-Mizil et al., 2013), norms (Rajadesingan et al., 2020), moderation policies (Gillespie, 2018), initial posts (Salganik et al., 2006), and the perceived demographic and social status of the participants (Munger, 2017).

This study aims to understand how the content that is posted online is affected by one particular piece of controversial language: the term *white privilege*. While the term *white privilege* existed in academic writings as early as the 1980s (McIntosh, 1990), the general public has become increasingly aware of it amid the heightened racial tension of the past decade (Saad, 2020). At the same time, social media has increased the availability of extreme, and often vitriolic, views online. A search for “white privilege” on any major social media platform will show a range of posts representing strong feelings from multiple ideological angles.

Social media has given people more options than ever for how to spend their time. Individuals today can scroll through a near-infinite stream of cat videos or talk about their favorite video game instead of engaging in uncomfortable discussions of race. Small changes in initial language have the potential to create large effects in both the content that gets posted and the traits of those engaged. To understand the effects of the term *white privilege* on social media discussions, we

² A version of this chapter was published as:

Quarles, C.L., Bozarth, L. (2022) How the term ‘white privilege’ affects participation, polarization, and content in online communication. PLOS ONE 17(5): e0267048. doi: 10.1371/journal.pone.0267048

ran two experiments in a simulated online environment. Respondents were asked, “Should colleges rename buildings that were named after people who actively supported X?” where X is either *racial inequality* or *white privilege*. We studied *how* people responded by looking at stance (pro/con), the frames (arguments, topics, and ideas) used in the response, and response quality. We also examined *who* would respond to the post by looking at both stated and actual likelihood of response. In addition, we use the posts to simulate the composition of responses in a real online forum.

3.1.1 How people respond to white privilege

Privilege is “unearned advantage derived from one’s group membership” (Phillips & Lowery, 2018). In the present study, white privilege refers to racial privilege in the American context. The concept of white privilege is central in areas such as contemporary diversity training (Case & Rios, 2017) and whiteness studies scholarship (Doane, 2003). However, in public discussion, the term is more controversial. Popular media has variously talked about white privilege as a topic to be taught to children (M. Brown, 2020), a racist term (Adams, 2020), and a distraction from the root causes of racial inequality (Malik, 2020). To be clear, this study does not directly examine the concept of white privilege itself, or whether whites think they have advantages due to their race. Instead, our goal here is to look at behavior: How individuals respond to the *term* in the context of an online forum. We expect that whites will respond differently to the term *white privilege* than other groups for two reasons.

Social identity theory suggests that we often define ourselves, and others, in terms of the groups that we are members of (Tajfel & Turner, 1979). A person’s behavior or perception of their social status might change based on which group membership is most salient at the time (Tajfel & Turner, 1979). The term *white privilege* evokes images of whites as a coherent group with representative traits. So we expect that the term will lead to increased salience of racial identity among whites, which will affect their responses.

In addition, whites have different views, on average, than members of other races about the advantages that whites have. In a recent Pew study, 47% of whites said that whites benefit either a great deal, or a fair amount, from advantages that Blacks don’t have (Pew Research Center, 2019a). In contrast, 89% of Blacks and 74% of Hispanics said that whites benefited from these advantages. While this difference in perception may come from motivated reasoning (Lowery et

al., 2007) or from genuinely different life experiences (McIntosh, 1990), by itself it is likely to affect how whites respond to the term *white privilege*.

Some individuals identify more strongly with their race than others. The strength of this pre-existing identification can give a differential effect on responses to racial priming, which has been shown in a variety of contexts with a variety of identities (R. P. Brown et al., 2008; Doosje et al., 1998; Lowery & Wout, 2010). American whites have repeatedly shown less identification with their race, on average, than other groups (Pew Research Center, 2019b), likely because being in the minority reinforces category differences and increases the salience of racial identity (Steck et al., 2003; Yang et al., 2008). However, whites vary in the strength of their racial identity, and this affects their thoughts, feelings, and behavior (Branscombe et al., 2007). While the current study does not include a measure of strength of racial identification, it is reasonable to expect that different groups of whites may respond differently to the term *white privilege*. Responses to the term *white privilege* do not come purely from a place of reasoned disagreement. One meta-study found that emotions were twice as important as beliefs in predicting discrimination (Talaska et al., 2008). Just like we can define ourselves using group stereotypes (Abrams & Hogg, 2010), the theory of intergroup emotions describes how group membership can cause us to feel emotions (Mackie et al., 2000). Anger has been shown to mediate the effects of perceived injustice on retributive action (Seip et al., 2014). And guilt has been shown to mediate framing effects on support for Dutch-Indonesian reparations (Doosje et al., 1998) and on perceptions of American racial inequality (Powell et al., 2005) among members of the dominant group. Those emotions do not stop when people go on social media (Duncombe, 2019). Since discussions of white privilege create uncomfortable feelings among some people, these heightened race group-based emotions may cause individuals to avoid engaging in online discussions.

3.1.2 Online conversations

Online information plays a significant role in shaping twenty-first century society. From the 24-hour clickbait-based news cycle, to discussion forums with infinite scrollers, to group-based conversations with friends on messaging apps, online media affects how we think about current events (Diehl et al., 2016), who our friends are (John & Dvir-Gvirsman, 2015), and how we feel about ourselves (Woods & Scott, 2016). However, our perceptions built using the online world

don't always represent reality (Bunker & Varnum, 2021; Lerman et al., 2016). The artificial reality we see online is sensitive to affordances and moderation policies of individual platforms (Budak et al., 2017; Gillespie, 2018) and is highly dependent on initial conditions (Salganik et al., 2006). In addition, media consumers interpret what they read based on pre-existing beliefs and biases (Taber et al., 2009). Ultimately, online media enables different groups of people to have very different perceptions of truth. Race is especially problematic in this respect, since differences in offline lived experiences have the potential to create barriers to a shared reality. We look at that online reality by examining four individual-level dimensions: *avoidance*, *conversation quality*, *stance* (support or opposition towards a topic under discussion), and the *frames* that are used in responses. To understand the system-level impressions of public opinion on a real discussion forum, we also examine the overall *composition* of posts.

Avoidance: Individuals' decisions about *whether* to participate in discussions play a central role in the social media landscape. Individuals avoid posting for a variety of reasons, including lack of time or interest, concern about offending someone or giving a bad representation of themselves (Sleeper et al., 2013). Individuals are also less likely to share negative and emotion-laden content (Bazarova et al., 2015), and are less likely to post in general if they are female, afraid of isolation, didn't feel strongly, or felt like their opinion didn't match the way the country was moving (Fox & Holt, 2018). While avoidance has the potential to be protective of social relationships, it can also lead to adverse personal effects from stifling expression (Butler et al., 2003). More systemically, avoidance is a key component of the "spiral of silence" (Noelle-Neumann, 1974), which leads to perceived minority opinions being underrepresented on social media (Lee & Kim, 2014). Of course, the vast majority of social media consumers are lurkers – people who consume content without contributing (Sun et al., 2014). And even regular posters read more than they post. In the context of race, people have been shown to distance themselves from sources of identity threat (Goff et al., 2008). So we expect that whites will be more likely to avoid responding to the *white privilege* question, particularly those whites who might feel like their ideas are in the minority or who experience identity threat.

Conversation Quality: Incivility and toxicity are important metrics for online spaces, and race-related topics are more likely to draw uncivil comments (Salminen et al., 2020). Even if posts can be categorized as civil, they may be confusing or add little to the conversation. So we operationalized a *low-quality* response as one that attacked people, challenged the question itself,

contained little content, or was hard to understand. Given the toxic nature of some online conversations around race (Mittos et al., 2020) and the discomfort many whites have with the concept of white privilege (Lowery et al., 2007; Pew Research Center, 2019a), we expect that the term will lead to lower average conversation quality among whites.

Stance & Frames: We measure the content of a post in two ways. **Stance** describes whether an individual supports or opposes the proposed topic. We also look at the topics, or arguments, mentioned in each response. These could be described as the *ideas* that the writers have about the topic. Alternatively, if we think of social media consumption, those same ideas become a way of **framing** the conversation. In this paper, we will use the term *frames* to describe this concept.

In the current context, we know that many whites do not believe they have race-based advantages (Pew Research Center, 2019a). The idea of *white privilege* is not consistent with their understanding of the world. Consequently, we hypothesize that fewer whites will be supportive of renaming building when *white privilege* is brought up.

Note that *stance* and *frames* are separate, but highly related. Supporters of a proposition typically find certain frames more salient than opponents do. For instance, abortion opponents often frame the procedure as ending a life, which puts the fetus at the center of attention. While pro-choice advocates tend to frame the issue around the needs and rights of the mother. Speakers and writers will influence support for a topic by framing the issue in different terms (Jacoby, 2000). In our experiments, we expect treatment condition to influence both stance and frames. Previous work suggests that that *white privilege* will have a primarily affective effect on individuals (Lieberman et al., 2005; Talaska et al., 2008). We expect this blunt mechanism to influence stance, instead of the frames used in complex reasoning. In this case, frame use would arise from motivated reasoning, as individuals tried to explain the stance that they had already chosen. So we hypothesize that there will be no significant difference in frames after controlling for stance.

Composition of Posts: Social media is used by individuals (Neubaum & Krämer, 2017), researchers (Prichard et al., 2015), journalists (McGregor, 2019) and policy makers (McGregor, 2020) to understand public opinion. However, responses on social media are not usually representative of the population as a whole (Hargittai, 2018). Online behavior depends on the community members, the affordances of the forum, and framing. To understand how the term *white privilege* affects this perception, we summarize the **composition** of responses in each treatment condition. By this we mean the set of responses, taken as a whole, as a reader might

perceive them. Unlike the other four dimensions, which focus on individual behavior, this variable describes the system's behavior. For instance, does an online community seem supportive of renaming buildings? Or does the community seem to oppose it? This composition can also create higher-order effects on the community, as individuals make decisions about what to post (Matthes et al., 2018; Sleeper et al., 2013). Given the relatively strong responses to the term *white privilege* online, and the lack of debate about whether racial equality is an important social value in the U.S., we expect that *white privilege* and *racial inequality* will create simulated communities with different compositions.

In summary, the literature suggests the following hypotheses:

Hypothesis 1 (Avoidance): Whites will be less likely to respond when asked about *white privilege*.

Hypothesis 2 (Stance): Whites will, on average, be less supportive of renaming buildings when asked about *white privilege*.

Hypothesis 3 (Conversation Quality): Whites will, on average, have lower quality responses when asked about *white privilege*.

Hypothesis 4 (Frames): Supporters and opponents of renaming buildings will bring up different sets of frames. And, after controlling for support, asking about *white privilege* will not affect the frames used.

While not a formal hypothesis, prior work suggests non-whites will either show no mean difference between treatment conditions in these first four dimensions, or show a trend in the opposite direction from whites. Overall, the first four hypotheses should lead to:

Hypothesis 5 (Composition of responses): In an online conversation, the use of the terms *racial inequality* and *white privilege* will result in a different composition of posts.

3.2 Study design

We explored these hypotheses through two experiments. Experiment A enabled us to gather responses from both individuals who would have posted online and those who would have self-censored. Because Experiment A asked people to self-rate their likelihood of responding, Experiment B examined revealed preferences by giving respondents a choice of questions to answer. A lab experiment was chosen to isolate the effects of language, avoid higher-order

network effects on peoples’ responses, and ensure that we could gather data about people who would otherwise avoid responding.

3.2.1 Respondents

Participants were US residents, drawn from Amazon Mechanical Turk (MTurk), who had completed 1000 tasks with 98% or higher acceptance rate. Both experiments were listed as the same task in the MTurk system. US resident MTurkers have been shown to be generally representative of the national population (Coppock, 2019). Participants were randomly assigned to experiment (A or B) and to treatment condition (*racial inequality* or *white privilege*). After excluding respondents who did not respond to the prompt, we were left with 478 people in Experiment A and 446 in Experiment B. Descriptive statistics about the sample are in Table 1.

Table 1: Demographics of respondents

	Experiment A		Experiment B	
	Racial Inequality	White Privilege	Racial Inequality	White Privilege
Number of Respondents	250	228	233	213
Male	51%	53%	56%	50%
Female	48%	46%	43%	49%
White	82%	78%	81%	84%
Black	11%	8%	6%	8%
Asian	6%	13%	9%	6%
Hispanic/Latino	6%	6%	5%	5%
Other	2%	2%	3%	3%
Multiracial	7%	7%	6%	7%
Bachelor’s Degree	59%	57%	67%	65%
Politics				
Mean	-0.42	-0.35	-0.37	-0.44
Standard Deviation	1.2	1.2	1.2	1.2

Politics was rated on a scale from -2 = *strongly liberal* to 2 = *strongly conservative*. Race percentages add to more than 100% because some people identified as multiracial.

We expected that people who identified *only* as white (74%) would tend to respond differently to the term *white privilege* than those who identified, at least in part, as a member of another race. To describe this latter group, we use the term *non-white* to signify that we don’t expect them to have the same white identity as those who identify as only white. Four respondents did not provide a race. They are included in any analyses which don’t involve race.

3.2.2 Instrument

Respondents in both studies received an online survey broken into two parts. After giving informed consent, respondents were sent to the Part 1 that corresponded to their experiment. In Part 1, each respondent was randomly assigned one of the two questions: “Should colleges rename buildings that were named after people who actively supported racial inequality?” or “Should colleges rename buildings that were named after people who actively supported white privilege?” The question language was chosen based on conversations with colleagues and vetting interviews during the study design phase. We purposely tried to use general language that might evoke a broad, identity-based response. *Racial inequality* was chosen as a counterpoint to *white privilege* because it seemed less likely to increase the salience of racial identity. Equality is an American ideal that we thought most respondents would support. And the topic of renaming college buildings seemed to give enough opinion diversity to see meaningful differences in the data.

In Part 1 of Experiment A, each respondent was randomly shown either the *racial inequality* or *white privilege* question. They were then asked: (a) “How likely would you be to respond to this question if you saw it in an online community?” and (b) “If you did reply to this question, what would you post in the online forum? Write the reply exactly as you might post it online.”

Responses to (a) were on a 5-point Likert scale from *very likely* (2) to *very unlikely* (-2).

Responses to (b) were free-written into a text box. After submitting Part 1, respondents were sent to Part 2.

Each participant in Experiment B was also randomly assigned to either the *racial inequality* or the *white privilege* condition. However in this case, for Part 1 participants were given the choice of two questions in a randomly chosen order. They were told that they could respond to either question, but only one. The questions were the renaming-buildings question (which depended on their treatment condition): “Should colleges rename buildings that were named after people who actively supported *racial inequality/white privilege*?” and the college-loans question: “Should college tuition loans be forgiven for people who choose to go into public service, such as social workers and teachers?” The college-loans question was chosen to avoid race and provoke a similarly diverse range of opinions. Text responses to the college-loans questions were not coded or used. After responding to their chosen question in a text box, respondents were sent to the same Part 2 as in Experiment A.

The benefit of the design of Experiment B is that it elicits behavior in a way that better approximates a real social media site. Attention is a precious commodity online. Ads and posts vie for time on consumers' screens. The option of an alternative question simulates that environment. Unlike in Experiment A, however, we do not get the censored responses from individuals who chose not to respond to the renaming-buildings question. These data are sensitive to the attractiveness of the other question. If the college-loans question is something that many or few of the sample would reply to, this will affect the effect size. The results are also sensitive to the college-loans question being *differentially* attractive to special groups, which has the potential to bias the sample in a way unrelated to our hypotheses.

Part 2 was a survey which asked primarily multiple-choice demographic questions. These included gender, age, race/ethnicity, preferred political party, and highest level of education. Part 2 was the same for both experiments.

3.2.3 Coding for stance and frames

The survey gave text responses for the renaming-buildings question from participants in Experiment A and from those who chose this question in Experiment B. We manually coded text responses to the renaming-buildings question for both *stance* and for the *frames* used in the response. Based on its written content, every text response was assigned to one of five stance categories: pro (supported renaming buildings), con (opposed renaming buildings), neutral, conditional (it depends on the person/situation), and unclear (when we could not discern support). For the purposes of analysis, we focused mainly on the pro and con categories. To create the framing codebook, each member of the research team initially independently coded 100 responses according to labels from Moral Foundations Theory (Graham et al., 2009), the Media Frames codebook (Boydston et al., 2014), and with frames generated by the responses themselves. We then collectively tried to synthesize our frames into a set of consistent, reasonable codes. Ultimately, neither Moral Foundations nor the Media Frames Codebook aligned with our sample's responses on renaming college buildings. So we developed and used our own set of codes through an iterative process: We coded a new set of responses using the previously created labels and with frames found in the new data. We then met and synthesized the codebook. This process repeated until the set of codes stabilized. Our codebook was

informed by the other two sets of frames, but definitions are different. For instance, our definition of *harm* does not exactly match the one used in Moral Foundations.

Once the codebook was created, each author independently coded every response in sets of about 100 responses. After each set, we met to discuss our codes until a consensus was reached on every response. Coders were blinded, so we did not know the treatment condition or respondents' demographics. Many responses had multiple frame codes. In the rare cases where there were more than three frames used in a response, we chose the three frame codes that were repeated the most often. In the case of ties, we chose the frames that were used earlier in the response. To calculate test-retest reliability, we performed this process again on a randomly chosen subset of 100 responses. This led to a test-retest reliability, using fuzzy kappa (Kirilenko & Stepchenkova, 2016), of $\kappa = .817$.

3.2.4 Frames

Here is the list of frame codes and the criteria used:

Erasing history – Any reference to erasing history or rewriting the past.

History as lesson – Mentions how we can learn from history and/or historical building names.

College's role – Refers to the college's image, relationship between the college and the community, or the values of the college. Must explicitly mention the college.

Cost – Mentions a scarcity of resources, or the amount of work required to take an action.

Progress – Reference to moving on from a problematic past, making progress on social issues, or solving problems today that we had in the past. Includes metaphors of motion or growth from a past state.

History is past – History is in the past, and is therefore not important or less important than contemporary issues.

Fairness – Equal treatment or preferential treatment. Interpreted narrowly. For example, a reference to equality doesn't automatically fall into this category.

Same people, different times – People are the same as they always have been. Or different times have different standards.

Individuals' contributions – The specific contributions of the individuals who the buildings were named after should be considered. Includes references to relative contributions of different

people, looking up to them as role-models, not honoring people who have done bad things, and references to worthiness due to monetary contributions.

Unintended consequences – There will be an unintended or surprising effect if buildings are renamed (or not renamed).

Inconsistency – There are inconsistencies in the present/future that would be created by renaming/not renaming. Typically referred to hypocrisy arising from some things being renamed when others aren't.

Different action – Suggests a different action, besides renaming buildings.

Harm – Someone will be harmed in the present or future. Includes people taking offense, disrespect, damage to social well-being, and supporting students. Both increasing harm and reducing harm fall in this category.

Authority – Any reference to the individuals who have the right to make the decision.

Doesn't matter – The decision to rename buildings will not have a practical impact. Or the discussion about renaming doesn't matter.

Ad hominem* – Attacks the parties involved in the debate, rather than focusing on the merits of renaming. Includes criticizing their character, calling names, suggesting they are hypocrites, or implying they have the wrong mentality.

Challenges question* – Attacks the language used in the question or challenges the question itself.

Other* – Response unrelated to the question, using a frame not listed above, or no clear frame. Includes simple answers like “yes”. Originally coded as three categories: *off topic*, *other frame*, and *no frame*. However, it was hard to separate these categories, since these responses were often not clearly written.

*Any response that included either the *ad hominem*, *challenges question*, or *other frame* was coded as a *low-quality* response.

To test for differences in proportions, we used Boschloo's test (Boschloo, 1970) using the Exact library (Calhoun, 2019) in R (R Core Team, 2018). The Fisher exact test is inappropriate to analyze contingency tables if column sums are not fixed by design. Boschloo's test adapts Fisher's approach by comparing p-values across different column sums. It is uniformly more

powerful than Fisher’s design. All Boschloo’s tests were one-tailed. The Plotrix library (Lemon, 2006) was also used for visualization.

3.2.5 Comparing frames

We were interested in inferring whether two groups C and D, such as whites and non-whites, were likely to use a different set of frames in their responses. This statistical analysis is challenging, since each response may have used 0, 1, 2, or 3 frames. In addition, there is no obvious statistical model which might explain how the groups use different frames.

So we used a random assignment Monte Carlo approach to infer whether two groups had similar frame use. We assumed as a null hypothesis that membership in Group C and Group D was independent of the probability of using each frame. We created a sampling distribution under the null by first tossing out the original group labels. We then randomly assigned every response to either Group C or Group D, ensuring that simulated groups had the same size as the actual groups. We calculated the test statistic under this simulated division. This process was repeated until we had 10,000 simulated test statistics. Our p-value is the percentage of these simulated test statistics which are larger than the test statistic for the actual sample.

For a test statistic, we used a variant of the Kullback-Liebler (KL) divergence (Kullback & Liebler, 1951). Let p_f^C be the observed proportion of responses from Group C that use frame f . Set p_f^D in a similar fashion. For the null hypothesis, let q_f be the proportion of responses in the complete sample $C \cup D$ that used frame f . Then, the test statistic is:

$$\sum_f p_f^C \log \left(\frac{p_f^C}{q_f} \right) + \sum_f p_f^D \log \left(\frac{p_f^D}{q_f} \right)$$

Note that this is not a true KL divergence, which is typically defined on a probability space where probabilities sum to one. In our case, each response can have multiple frames, so $\sum_f q_f > 1$. However, like KL divergence, this test statistic does measure how different the observed group probabilities p_f^C, p_f^D are from the reference distribution q_f corresponding to the null hypothesis.

All respondents gave informed consent through a digital interface. The University of Michigan institutional review board approved this study.

3.3 Experiment A results:

Experiment A was designed to understand both the responses of people who would respond in an online forum, as well as responses from people who would avoid posting online. So we asked everyone in the sample to respond to the prompt, and then self-rate how likely they would be to respond to it in an online community.

For the purposes of this analysis, we defined someone as a *likely responder* if they said they would be *somewhat likely* or *very likely* to respond to the question. We used this group to understand what might actually be posted online.

Table 2 gives some results from Experiment A.

Table 2: Experiment A - Likelihood of responding, stance, and response quality by treatment group and race

	Whites			Non-Whites			Likely Responders	
	Racial Inequality	White Privilege		Racial Inequality	White Privilege		Racial Inequality	White Privilege
Count	189	161		59	66		133	97
Average self-reported likelihood of responding	0.169 (.11)	-0.255 (.11) **		0.203 (.19)	0.288 (.18)			
% Supported renaming	48	24 ***		42	42		64	38 ***
% Opposed renaming	29	41 **		27	30		18	38 ***
Low quality response	22	37 **		24	36 +		20	36 **

+ p < .1, * p < .05, ** p < .01, *** p < .001

Respondents rated their likelihood of responding on a scale from 2 = *very likely to respond* to -2 = *very unlikely to respond*. Values in parentheses are standard errors. P-values represent differences between treatment groups. Three individuals did not provide a race.

3.3.1 Avoidance

Based on their self-reported likelihood of responding, whites were less likely to respond to the *white privilege* question than the *racial inequality* question ($t(344) = 2.73, p = .003$). In contrast, non-whites were not significantly more likely to respond to the *white privilege question* ($t(121) = -0.33, p = .372$).

3.3.2 Stance

Because we had coded multiple categories for stance, we separately report the percentages of people who supported (pro) and opposed (con) renaming buildings. The other stance categories did not have enough responders to draw reliable conclusions.

Whites in Experiment A were less likely to support ($p < .001$) and more likely to oppose ($p = .008$) renaming buildings when the question was phrased in terms of *white privilege*. This overall

shift in stance among whites was surprising. When asked about *racial inequality*, whites were 67% more likely to be supportive than opposing. However, when *white privilege* was mentioned, 74% more whites opposed renaming college buildings than supported it.

As with avoidance, the choice of *racial inequality* versus *white privilege* did not affect average support ($p = .505$) or opposition ($p = 0.667$) among non-whites. This reinforces previous work that shows individuals have different responses when primed to think about their own group compared with another group.

Among likely responders, the term *white privilege* significantly decreased support for renaming buildings. In the *white privilege* condition, support dropped by 26 percentage points ($p < .001$), and opposition increased by 20 percentage points ($p < .001$). Unlike the results for whites and non-whites, these differences are caused by differences in *who* would respond in addition to stance changes.

3.3.3 Response Quality

Framing the question in terms of *white privilege* increased the percentage of low-quality responses. This was true among whites ($p = .001$), non-whites ($p = .069$), and likely responders ($p = .003$). The percentages for all groups were similar, so the decreased significance among non-whites is likely due to a smaller sample size.

3.3.4 Frames

As predicted, the biggest difference in frame use was between supporters and opposers of renaming buildings ($p < .001$). The frequency of frame use for supporters and opposers is shown in Figure 5. We did not find a difference between the frames that whites and non-whites used in their responses ($p = 0.768$). This result held when we restricted the analysis to only those who received the *racial inequality* ($p = 0.912$) and *white privilege* ($p = 0.649$) questions.




Figure 5: Percentage of responses in Experiment A that used each frame. Squares give the proportion of responses that used a given frame, among all responses that supported renaming buildings. Diamonds represent frame use among all responses that opposed renaming buildings. Starred frames were categorized as low-quality.

Treatment condition did affect the frames that people used in their responses in both the complete sample ($p = 0.018$) and among likely responders ($p = 0.029$). Was this because the terms *racial inequality* and *white privilege* bring up different ideas in peoples' minds? Or was it due to the fact that there are more supporters in the *racial inequality* condition, and supporting arguments generally use different frames?

To answer this, we performed a mediation analysis. We ran a logistic regression predicting the use of each frame based on treatment condition, controlling for support and opposition:

$$\text{logit}(F_i) = \alpha + \beta(\text{treatment}_i) + \gamma(\text{pro}_i) + \delta(\text{con}_i) + \epsilon_i$$

Here F indicates whether individual i used the chosen frame, treatment tells whether the individual received the *racial inequality* or *white privilege* question, and pro/con are binary variables that describe whether the individual supported or opposed renaming buildings. We ran this regression on every frame except the low-quality frames, which as described above did seem to show a difference between treatment conditions, and the *consistency* frame, which was used so rarely that the regression was not valid.

If the frames that people use in each treatment condition can be explained by their stance, then we would expect the coefficient of *treatment* to be uniformly distributed and mostly statistically insignificant. Though we do expect statistical significance ($\alpha = .05$) to occur by random chance around 5% of the time. This is what we found. Of the 17 regressions only one frame, *erasing history*, had a p-value less than .05 ($p = .014$). The p-values seemed uniformly distributed, with the largest p-value for *authority* ($p = .862$). The effect of the term *white privilege* on framing was explained by individuals' stances.

3.3.5 Composition of responses

How does the question language affect the overall composition of responses that get posted online? We turn to the set of likely responders to analyze this question. Figure 6 gives a snapshot of what an online conversation might look like in each condition. The *racial inequality* question led to a set of likely responses that was overwhelmingly supportive of renaming buildings, with 7 supporters for every 2 opponents. In contrast, the *white privilege* framing led to a more divided set of responses, with roughly equal numbers of supporters and opponents. Different frames were brought up in the two conditions as well. Though, as mentioned, this seemed completely driven by differences in support. The *white privilege* question brought 80% more low-quality responses than the *racial inequality* question.

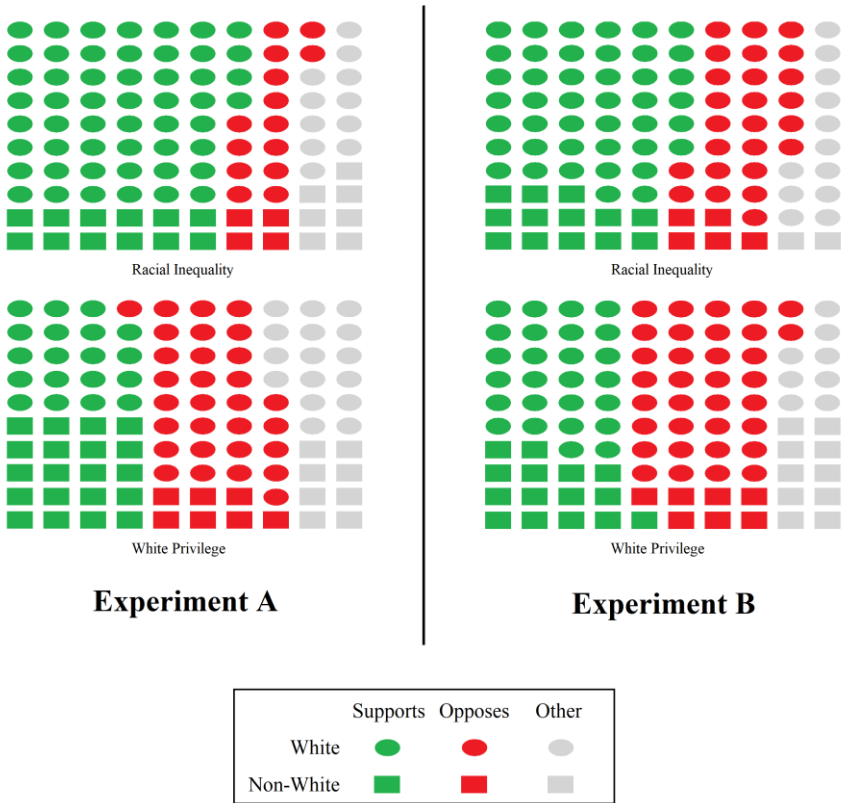


Figure 6: Composition of posts in a hypothetical online conversation among 100 responders who are representative of our sample. For Experiment A, the figure represents likely responders. For Experiment B, the figure represents those who responded to the renaming-buildings question. Shape corresponds to the race of each responder. Points are colored based on support for renaming buildings. The Other category includes responses that were neutral, unclear, or said that it should depend on the situation.

3.3.6 Avoidance differences between whites

The effect of using the term *white privilege* did not affect all whites equally, as shown in Figure 7. Supportive whites were less likely to respond to the *white privilege* question than the *racial inequality* question ($t(62) = 3.03, p = .004$). However, whites who opposed renaming buildings were approximately equally likely to respond in both conditions ($t(114) = -0.48, p = .635$).

Language choice did not affect the likelihood of responding among either supportive or opposing non-whites.

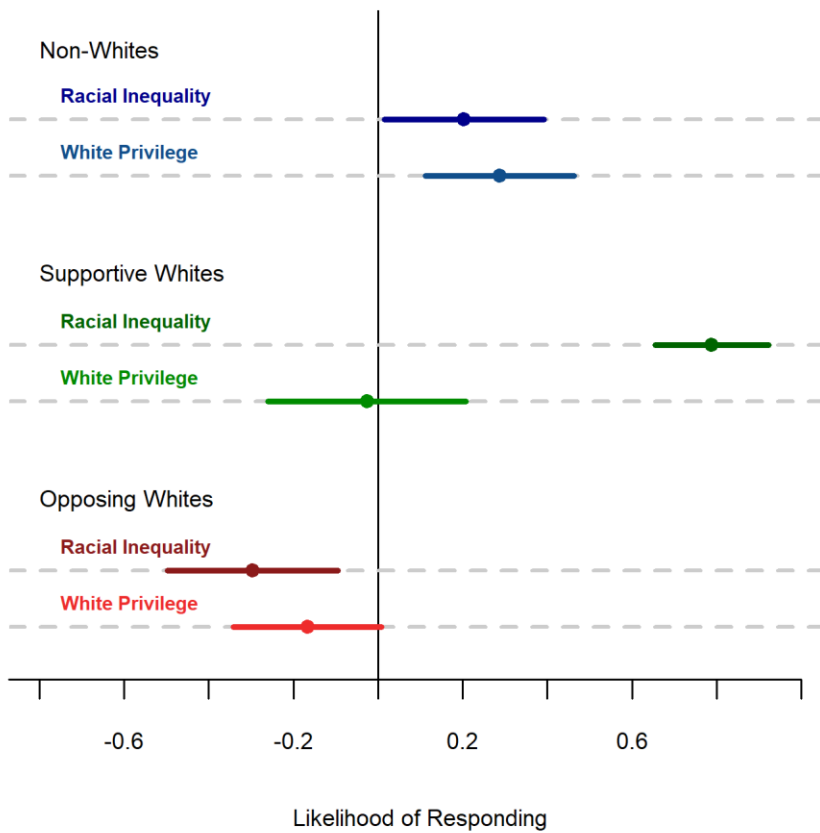


Figure 7: Average self-reported likelihood of responding in Experiment A. Respondents rated their likelihood of responding on a scale from 2 = very likely to respond to -2 = very unlikely to respond. Error bars represent standard errors.

Overall, the results show that the shift from a set of overwhelmingly supportive responses under *racial inequality* to the divided responses under *white privilege* comes from two factors: (a) whites were, on average, less supportive of the *white privilege* question, and (b) supportive whites were less likely to respond to the *white privilege* question.

3.4 Experiment B results

As a counterpoint to Experiment A, where people self-rated their likelihood of responding, Experiment B was designed to examine revealed behavior and see how people might respond in a simulated online environment. Respondents were given (a) the renaming-buildings question that corresponded to their randomly assigned treatment group and (b) the college-loans question. They were told to respond to only one of the questions.

37 respondents filled in the text boxes under both questions. This meant they provided a response for the college loans question, but that it was unclear whether they preferred to answer that question. Since our analysis focused on people who chose to respond to the renaming-buildings question over the college loans question, we excluded those 37 data points from the analysis in this section. For completeness, we performed a robustness check with those individuals included. The results were qualitatively similar to the results below but with smaller effect sizes.

Table 3: Experiment B - Probability of responding, stance, and response quality by treatment group and race

	Whites			Non-Whites			All Combined	
	Racial Inequality	White Privilege		Racial Inequality	White Privilege		Racial Inequality	White Privilege
Count	163	152		49	44		213	196
% Responding to Renaming Buildings Question	37	28 *		33	43		36	31
Among those...								
% Supported Renaming	54	38 +		62	47		56	41 *
% Opposed Renaming	31	50 *		25	21		30	41 +
% Low Quality Response	32	40 *		0	25 **		19	38 **

+ $p < .1$

* $p < .05$

** $p < .01$

*** $p < .001$

P-values represent differences between treatment groups. One individual did not provide a race.

The results in Table 3 tell a story consistent with the results from Experiment A. However, these results have generally weaker statistical significance. In particular, some of the effect sizes for non-whites seem to be similar to whites' effect sizes, but without sufficiently small p-values. This is likely due to a smaller sample size. The alternate question about college loans seems to have been too attractive, with only about 1/3 of respondents answering the renaming-buildings question. This preference for the financial question over the race-related question held regardless of race or treatment condition, and warrants investigation in future studies.

3.4.1 Avoidance

As in Experiment A, whites were less likely to respond to the *white privilege* question by nine percentage points ($p = .035$). Non-whites in the sample were 10 percentage points more likely to respond to the *white privilege* question ($p = .160$), but this did not rise to the level of statistical significance. So the effect for non-whites could be due to sampling variation. These results support Hypothesis 1.

3.4.2 Stance

Whites who responded to the *racial inequality* question were, on average, more positive about renaming college buildings than those who responded to the *white privilege question*. They were 16 percentage points more likely to be supportive ($p = .058$) and 19 percentage points more likely to oppose ($p = .030$). Interestingly, non-white responders also seemed more positive about the *racial inequality* question. Though the sample size was small enough that neither the difference in support ($p = .202$) nor opposition ($p = .427$) were significant. When we consider the set of people who responded to the renaming buildings as a whole, the people who received the *racial inequality* question were more likely to be supportive ($p = .043$) and less likely to oppose ($p = .091$).

3.4.3 Response quality

Responses to the *white privilege* question garnered a higher percentage of low-quality responses among whites ($p = .047$), non-whites ($p = .010$), and all responders ($p = .010$).

3.4.4 Frames

As in Experiment A, there was a large difference in frame use between supporters and opponents of renaming buildings ($p < .001$). There also was a significant difference in the frames between treatment conditions ($p < .001$). To analyze the effect of stance on frame use, we ran a logistic regression for each frame as described in Experiment A. The frames *unintended consequences* and *cost* were omitted from this analysis due to low use. The low-quality frames were also omitted. After controlling for stance, there was no effect of treatment condition on frame use beyond what we would expect by chance. The p-values were distributed fairly uniformly with the smallest p-value corresponding to the *consistency* frame ($p = .040$) and the largest corresponding to *erasing history* ($p = .076$). Again, the effect of question (*racial inequality/white privilege*) on frame use was completely explained by stance. These results support Hypothesis 4.

3.4.5 Composition of responses

Figure 6 shows the overall composition of responses. As before, *racial inequality* led to more supportive responses and fewer low-quality responses than when the question was framed in terms of *white privilege*. As in Experiment A, there were equal numbers of supporters and opponents when asked about *white privilege*, and responders were generally supportive when

asked about *racial inequality*. There were 1.9 supporters for every opposer in the *racial inequality* condition. This was weaker than in Experiment A, where the support/opposition ratio was 3.5. It is unclear whether this weaker support is caused by the attractiveness of the college-loans question, a difference between stated preferences (Experiment A) and revealed preferences (Experiment B), or random chance.

3.5 Summary of Results

These results shed light on our hypotheses. Hypothesis 1 and Hypothesis 2 are both confirmed by the data. Whites who received the *white privilege* questions were less likely to respond and less supportive of renaming buildings. We also found support for Hypothesis 3. Use of the term *white privilege* led to more low-quality responses. This result was not only true among whites, but also among non-whites. The results also support Hypothesis 4, which focused on motivated reasoning. Supporters and opponents of renaming college buildings used different arguments. However, differences in framing between people who received the *white privilege* and *racial inequality* question disappeared after taking into account their stance. These experiments also provided evidence for Hypothesis 5. The term *racial inequality* created a set of responses that supported renaming college buildings. *White privilege* led to a more divided, polarized set of posts. While the effects of the term *white privilege* on whites was unambiguous, the effect on non-whites was less clear due to a combination of smaller sample sizes and seemingly weaker effects. The only reliable result among non-whites was that *white privilege* led to more low-quality responses.

3.6 Discussion

Using two experiments, we studied how individuals respond to the term *white privilege* in an online environment. Mentioning *white privilege* was enough to flip white support for renaming college buildings from primarily supportive to primarily opposing. Furthermore, the term *white privilege* deters some supportive whites from engaging in the conversation. Surprisingly, we did not see this avoidance effect among opposing whites. In addition, the term *white privilege* led to less constructive responses among both whites and non-whites.

If these were posts on a real online discussion board, asking about *racial inequality* would give the impression of general support for renaming college buildings. Asking about *white privilege* would lead to a seemingly less supportive, more divided public opinion with lower-quality online

debate. This decreased support is driven by two factors: (a) whites were, on average, less supportive when *white privilege* was brought up, and (b) supportive whites were more likely to avoid talking about *white privilege*.

Responses to *white privilege* tended to use different arguments from arguments about *racial inequality*. However, that difference was completely explained by differences in stance toward renaming buildings. This lends credence to the claim that the term *white privilege* leads first to a change in stance, followed by motivated reasoning to support that stance. If the causality went the other way, where the choice of language first affects the ideas people have, which leads to them changing their support, then we might expect at least some of the frames to be unexplained by stance.

Prior literature suggests that both emotion (Mackie & Smith, 2015; Powell et al., 2005) and the strength of racial identity (Doosje et al., 1998) play a significant role in our results. We hypothesize that the increased tendency of supportive whites to avoid discussing *white privilege* is mediated by both these factors. It could be that the term made racial identity more salient for all whites, but was more likely to generate guilt and therefore avoidance in supportive whites. Another possibility is that opposing whites tended to identify highly with their race already, so that mentions of *white privilege* had a greater average effect on both racial identity salience and emotion on lower-identifying whites. Future research might test these hypotheses.

In writing about this study, we had to refer to groups, such as “non-whites” and “supportive whites”. There is a lot of variation among the individuals in any group, especially racially-defined groups with millions of members. However, humans have an unfortunate tendency to generalize a statement about a group of people to each individual member (Abrams & Hogg, 2010). This overgeneralization can cause harm, for instance through stereotyping (Zaniboni et al., 2019). Our study, like many research studies, is about averages. So we have been careful to use language that minimizes overgeneralization to individuals. For instance, instead of writing, “Whites were less supportive of the *white privilege* question”, we wrote “Whites were, *on average*, less supportive of the *white privilege* question.” Our results should be interpreted as describing how language affects large-scale social dynamics, not as a way to understand traits or behaviors of individuals.

3.6.1 Limitations

In a real online site, social desirability bias, the design of the forum, and back-and-forth between posters may magnify or dampen the effects we saw here. Another limitation comes from the fact that most social media users post very rarely. Online, the desires for information and entertainment are major drivers of behavior. Indeed, some researchers emphasize the value of active listening (Thill, 2015), which can bring a more diverse set of perspectives. All participants in our study were motivated to respond. It is unclear how the desire to read others' points of view might affect these results. In addition, Experiment A and Experiment B had quantitatively different but qualitatively similar results. So in a true online environment, we might expect a similar effect, but with potentially different effect sizes.

The present study does not capture long-term attitude changes. Further research is required to understand the circumstances under which long-term exposure to the term *white privilege* affects support for racially progressive policies, whether it increases animosity and polarization, and how this effect might differ between demographic groups.

While we chose the language in the study to broadly evoke group-based identity, the terms *racial inequality* and *white privilege* do have different literal meanings. The survey prompt asked individuals to think about buildings named after people who supported these two separate concepts. It's not clear whether that difference in meaning affected their responses. Concerns about building names have cited a variety reasons, from the honoree being a Confederate to supporting eugenics. Perhaps *white privilege* and *racial inequality* suggest different reasons, which led to different responses by treatment group.

3.6.2 Implications

Our study has several practical implications. The first is already known, but often ignored: Opinions on social media do not represent public opinion. Social media posts are highly dependent on how a question is phrased, as well as the norms, community members, and moderation practices of the site. Individual and system-level forces, such as self-categorization (Abrams & Hogg, 2010), the spiral of silence (Matthes et al., 2018), and algorithmic filters (Thorson et al., 2021) affect what shows up on our feeds. In our study, which did not include the moderation found on social media platforms, a two-word change in language was sufficient to shift a community from appearing divided to appearing supportive. This result will not be surprising to survey researchers, who need to be very attentive to choice of language (Fowler &

Cosenza, 2008). However, policy-makers (McGregor, 2020), journalists (McGregor, 2019), and others who use social media to understand the opinions of others may want to turn to more valid sources.

Those who want inclusive online conversations around race and/or support for racially sensitive policies should think carefully about the use of language like *white privilege* that targets the racial identity of specific groups. This language can deter the targeted group from participating. It has the potential to increase affective polarization by creating the image of a politically divided online space. Using slightly different language, such as *racial inequality*, that has more of a shared meaning across cultures can lead to conversations with broader participation and greater shared support.

In discussing this study with academic colleagues, a common response was, “Even if the term *white privilege* makes whites feel uncomfortable, they still need to hear it. It’s part of learning about race.” Indeed, numerous scholars have argued for raising awareness of race-based privilege (Case & Rios, 2017). Spending time thinking about racial advantages and disadvantages can affect individuals’ perceptions of systemic discrimination (Branscombe et al., 2007; Stewart et al., 2012). However, these effects vary significantly depending on the details of the intervention and the individuals involved (Branscombe et al., 2007; Case & Rios, 2017; Stewart et al., 2012). Our results, which focused on a simple change of language in an impersonal context, show that mention of *white privilege* can decrease engagement and lead to opinion shifts opposite to what was intended. It’s reasonable to expect that this identity-based disengagement decreases learning for some whites – an effect which has been documented in other settings (Heikamp et al., 2020; Steele, 2010; Zhao et al., 2019). Humanity has an evolutionarily useful, but usually incorrect, tendency to treat all members of a group as being the same (Abrams & Hogg, 2010; Turchin, 2007). As commonly used, the phrase *white privilege* draws on this tendency to conflate individual traits with group averages, in a way that creates unpleasant emotions. A more effective approach might be to distinguish between individuals’ experiences and group averages through a combination of personal storytelling and large-scale data in a way that is consciously inclusive of whites (Plaut et al., 2011).

3.7 Conclusion to Chapter Three

With online political polarization on the rise (Iyengar et al., 2019) and race in the forefront of today's news, it is important to make cross-cultural online communication effective and inclusive. The present work adds to what we know about communication on racially challenging topics. This study has shown that the term *white privilege* in online conversations tends to decrease support for racially ameliorative policies among whites, cause some supportive whites to avoid participating in discussions, decrease overall online conversation quality, and lead online forums to seem more polarized. Other, more inclusive, ways of speaking about race online, such as the term *racial inequality* are more likely to create a sense of shared purpose. There are very real racial inequities in society today. Choosing language that promotes constructive conversation will not solve those problems. But it is an important step towards collectively understanding their dimensions and working together towards a solution.

Chapter 4

Modeling the Accumulation of Success

4.1 Introduction

Decades of social science and health research have shown that the factors influencing success and well-being are highly interrelated. It is clear that a highly connected web of relationships underlies our social system. In this chapter, I develop a model describing multifactor cumulative advantage leading to a probability distribution of outcomes which can be interpreted in terms of real data. I then calculate the relationship between the inequality in the outcome distribution and the cumulative advantage process that generated it. I show that, under certain conditions, such a model can be approximated by a one-dimensional model which can be used to understand the higher dimensional context. I explore the behavior of these one dimensional models and, to the extent that their moments exist, I build a “dictionary” relating outcome distributions with the processes that generate them. I then explore various examples of generating processes and outcome distributions, including the accumulated Bernoulli, accumulated Pareto, and accumulated exponential. The results provide a pathway for researchers to understand multifactor cumulative advantage and how it might influence outcomes.

4.1.1 Literature Review

It has long been recognized in social science that various resources, traits, and skills reinforce themselves to create what we call professional success or well-being. In *The American Occupational Structure* (1967), Blau & Duncan used an impressive set of data and path analyses to point out the fundamental relationships between parent and child’s socioeconomic status. This approach was adapted by others, for instance by adding new variables to create greater precision (Sewell et al., 1969) or putting the idea of path analysis in a new context (Tyree et al., 1971). This approach to examining the relationships between quantities has propagated widely in the research literature, much as its initial proponents called for. For instance, Sewell, Haller & Portes

called for the inclusion of social psychological variables in the Blau & Duncan model (1969). In the years since, a plethora of research has examined how myriad social & psychological variables influence success (Heckman et al., 2006; Kraus & Park, 2017; Yeager & Walton, 2011), delving into more and more detail. In one fine-grained review, Farrington, et al., (2012) explored the effects of noncognitive factors on academic performance. They highlighted five categories of skills which influenced each other as well as educational success: academic behaviors, academic perseverance, academic mindsets, learning strategies, and social skills. In their classification, each category was made up of multiple, more fine-grained skills.

Path diagram models used in research are often acyclic. That is, they do not include pathways that either allow a factor to influence itself. However, there is a recognition in the research literature that the factors influencing success can be mutually reinforcing (Cortright, 2006; Hovmand, 2014; Levy et al., 2020). And *a priori*, it makes sense that factors influencing success should reinforce themselves in a cycle of cumulative advantage (DiPrete & Eirich, 2006). For instance, good health might mean fewer sick days, leading to a higher income, which in turn could provide money and insurance to visit the doctor and better health.

This mutually reinforcing nature of causal relationships shows up in health as well. The term *syndemic* describes the situation where multiple, reinforcing health and/or social conditions co-occur (Singer et al., 2020). Aging researchers have documented the condition of *frailty*, a term which encompasses a wide variety of physical states related to old age. Frailty is best defined in a systems biology framework, as it is characterized by a large number of interacting factors that occur between the cellular and human scales, as outlined in (Fried et al., 2005). At the genetic level, the omnigenic model of complex physical traits highlights how almost all genes have an important role to play in human health (Boyle et al., 2017).

A number of models have explored how multiple factors can simultaneously influence success and inequality. Critical quantitative researchers have explored intersectionality, the simultaneous effect of being in multiple categories, using latent class analysis (Landale et al., 2017; A. S. P. Wilson & Urick, 2022). This strategy groups individuals based on multivariate categorical data, looking for commonly shared identities which can then be used as variables in further analyses. Bloome (2015) used a transition matrix approach to examine the forces governing black-white income inequality trends. She classified individuals into bins based on family structure, income quintile, and age, and then used a Markov chain approach by building a stochastic transition

matrix describing the probability of moving between bins each year. The Social Genome Model used a series of regressions to predict behavior at five different stages of life (Sawhill & Reeves, 2016). Each regression predicted an individual's outcome using circumstances at birth and the outcomes of previous life stages. This allowed them to simulate experiments using the model to show, for instance, that early and repeated interventions can improve outcomes for disadvantaged children (Sawhill & Karpilow, 2015). In studying frailty, Mitnitski, Bao & Rockwood (2006) developed a model based on the “accumulation of deficits” (clinically measurable health issues). They used a Markov chain approach using a Poisson process, where the average number of deficits gained at the next time step is a linear function of the number of deficits one already has. In essence, their model treated all deficits as practically equivalent, and having an additive effect on future outcomes. Interestingly, their model was highly predictive ($R^2 = .979$) despite being functionally independent of age. Though it is not clear whether they used out-of-sample validation to calculate this R^2 . This suggests that physiological age is a much stronger predictor of future health decay than chronological age.

The work in this chapter builds on these ideas. However, this model varies from most of the earlier work in an important aspect: I do not name or put a limit on the number of variables. Nor do I assume that the variables involved are measurable in practice. This disconnects the model from fine-grained microdata like Blau & Duncan used. Rather than validating or training the model on inputs, we can only examine the outputs to see if they are similar to variables of interest. However, many of the factors that influence our outcomes are *truly* hard to measure. So my hope is that this “fuzzing of the eyes” allows deeper understanding in some ways. In particular, a major goal of this work is to connect distributions of outcome data, such as income or mental health to the accumulative processes that might generate them.

4.2 Model

4.2.1 A Thought Experiment


Consider a situation where everyone possesses three non-negative valued traits, which we'll call social skills, professional network strength, and knowledge, $\{s, p, k\}$. Now these are useful traits, any of which, to the extent possible, people will tend to “reinvest” to gain improve the other traits. For instance, a person can use their social skills to gain a wider professional network. In

turn, they can gain knowledge from a large network of colleagues and friends. Consistent with this framework, evidence shows these variables are correlated (Algan et al., 2022; Wang et al., 2018). An individual’s social skills s_t at time t might be described by a linear equation. (The coefficients in this subsection have been arbitrarily chosen for demonstration purposes.)

$$s_{t+1} = s_t + .01s_t + .01p_t + .005k_t$$

The s_t term means that they keep whatever social skills they had at the previous time step. And there is, for instance, a 0.5% “return” on knowledge as the benefit pertains to social skills.

Similar equations could be written for p_{t+1} and k_{t+1} . Combining equations for all three variables



might give us the following causal network and matrix equation.

Figure 8: Example matrix of interactions between social skills, professional network strength, and knowledge. Numbers are for demonstration only.

$$\begin{pmatrix} s_{t+1} \\ p_{t+1} \\ k_{t+1} \end{pmatrix} = \begin{pmatrix} 1.01 & .01 & .005 \\ .03 & 1.005 & .015 \\ .002 & .012 & 1.005 \end{pmatrix} \begin{pmatrix} s_t \\ p_t \\ k_t \end{pmatrix}$$

If we call this matrix A , then the long term behavior of these traits is determined by A^t .

In fact, for large t the Perron-Frobenius Theorem, which we discuss later, tells us that, in this model, social skills, professional network, and knowledge will grow according to:

$$\begin{pmatrix} s_t \\ p_t \\ k_t \end{pmatrix} \sim C \begin{pmatrix} 2.9 \\ 4.7 \\ 2.4 \end{pmatrix} (1.031)^t$$

The constant C is dependent on each person's initial values $\{s_0, h_0, k_0\}$. For instance, in this simple model, people's professional networks will be roughly three times the value of their social skills over the long term.

We could then put these variables into a regression model predicting, say, income.

$$i_t = b_0 + b_s s_t + b_p p_t + b_k k_t$$

This would tell us both how income grows exponentially over time and how inequality changes. If we know the distribution of traits at time t_1 then we can easily predict the distribution at time t_2 .

A key idea of this paper is this approach to reducing complex causal relationships to low-dimensional representations.

The model so far is deterministic, in that late life outcomes are completely determined by one's initial state. This is not realistic. We might, by chance, get an excellent teacher who is good at imparting knowledge. Or we might stumble onto a best friend who is a professional hub. So we will tweak this model by assuming that, at each time step, everyone receives a non-negative amount added to each factor. The balance between this additive growth and the multiplicative growth from the matrix will be our primary source of interesting behavior.

An alternative framing might be as an accumulation of deficits, like that used in aging & frailty research (Rockwood & Howlett, 2019). Using this viewpoint, s, p, t represent *problems*, with bigger numbers referring to worse situations. A person who has trouble exercising due to arthritis will have worse heart health. Lower heart health may, in turn, increase cognitive decay. This is a *cumulative disadvantage* approach, whereas the social skills/health/knowledge viewpoint corresponds to *cumulative advantage* (DiPrete & Eirich, 2006). The model works either way, but we will focus on the positive framing for this paper.

4.2.2 Model Limitations

Before delving deeper into the details of the model, it is worth pointing out some limitations to this approach. Since individual traits do not typically grow exponentially throughout the lifespan

(Lam, 1997), I will interpret time periods flexibly. Some time periods will be shorter than others. For instance, while children may grow quickly, working adults spend much of their potential learning time working to earn an income. So an adult may have fewer time points per year than a child. Also, it's clear that the reinvestment strategies of children and adults will differ, as will the strategies of adults in significantly different professions. However, there are quantities that grow exponentially on average, such as income within a country or frailty (Shi et al., 2011).

In addition, this approach assumes the population consists of a cohort with relatively homogeneous goals and environments. An individual's tendency to reinvest, say, social skills into professional network strength will depend on a variety of factors. Rural and urban workers may use those social skills differently, as might professional salespeople use social skills differently than bus drivers. So the matrix A will be different for very different groups of people. The causal network model proposed here is a model of cumulative (dis)advantage. It assumes all positive relationships, which might not be the case. For instance, it may be that having strong social skills increases both your economic mobility and improves ties with your family. However, if your family is lower income, then those family ties may tend to decrease mobility. Thus social skills could have an increasing *and* decreasing effect on mobility. Similarly, we could imagine a case with a negative feedback loop, where social skills indirectly and causally leads to lower social skills. Some of these issues can be resolved by flipping the sign of a particular variable, or by making the following argument: If people *can* invest a resource in their well-being, then *some* people will. So the model may be valid for cohorts more than for individuals. Exploring these cases is outside the scope of this present work.

This model does not take into account exogenous variables that are not influenced by other variables as shown in Figure 9. Parental socioeconomic status, genetics, or inherited epigenetics (Torche, 2018) are examples of traits which might be minimally influenced by other traits in someone's life. In the context of the work below, this will mean that the matrix A is not irreducible. Instead of the adjacency matrix A being a primitive matrix, exogenous variables would lead it to be block triangular matrix.




Figure 9: Exogenous variables (blue) may be uninfluenceable (or minimally influenceable) by other variables in the model.

4.2.3 The Accumulation Model

To describe the model, we modify a standard linear regression equation.

$$y = b_0 + b_1x_1 + b_2x_2 + \dots + b_Mx_M + \epsilon$$

We will not worry, for the moment, about whether each independent variable is directly measurable. This allows us to assume that the set of $\{x_i\}$ is large enough to encompass both the constant and error terms. We also center y around its mean, if necessary. The regression model can then be written in vector notation as a dot product or in terms of a matrix transpose.

$$y = b_1x_1 + b_2x_2 + \dots + b_Mx_M = \vec{b} \cdot \vec{x} = \vec{b}^T \vec{x}$$

Each individual, which could be a person or a larger discrete community, then has a set of M non-negative, real-valued variables $\{x_{1,t}, x_{2,t}, \dots, x_{M,t}\}$ at time t . Each x_i provides a resource, which individuals desire and can be reinvested. At each time step, two things happen:

- 1) Each variable x_i gains an amount proportional to each other variable x_j , with an “interest rate” of a_{ji} which adds up to $x_{i,t+1} \rightarrow a_{i1}x_{1,t} + a_{i2}x_{2,t} + \dots + a_{iM}x_{M,t}$
- 2) The individual gains a random amount in each variable defined by a random vector $\vec{X} = \langle X_1, X_2, \dots, X_M \rangle$. We assume the distribution of X_i is time independent.

Overall this means that:

$$x_{i,t+1} = a_{i1}x_{1,t} + a_{i2}x_{2,t} + \dots + a_{iM}x_{M,t} + X_i$$

Writing with matrix notation:

$$\vec{x}_{t+1} = A\vec{x}_t + \vec{X}$$

Steps (1) and (2) can be written more briefly as a multiplicative process:

$$\vec{x} \rightarrow A\vec{x}$$

And an additive process:

$$\vec{x} \rightarrow \vec{x} + \vec{X}$$

Let \vec{X}_t be random vectors identically distributed to \vec{X} , which represent an individual's realization of \vec{X} at time t . If $\vec{x}_0 = 0$, then the value of the output \vec{x}_T is itself a random vector:

$$\begin{aligned} \vec{Z}_T = \vec{x}_T &= X_T + AX_{T-1} + A^2X_{T-2} + \dots + A^{T-1}\vec{X}_1 \\ \vec{Z}_T &= \sum_{t=1}^T A^{T-t}\vec{X}_t \end{aligned}$$

Equation 1

For a fixed set of coefficients \vec{b} , the outcome of the regression y at time τ is also a random variable.

$$W_T = \vec{b}^\top \sum_{t=1}^T A^{T-t}\vec{X}_t$$

Equation 2

To summarize:

Definition: An **accumulation model** is defined by (\vec{X}, A, T) , where \vec{X} is a non-negative real-valued random vector of length M , A is an $M \times M$ primitive matrix, and T is a positive integer representing time. Let $X_t \sim X$ be a set of independent, identically distributed (iid) variables. Then an accumulation model produces a random vector:

$$\vec{Z}_T = \sum_{t=1}^T A^{T-t}\vec{X}_t$$

Which describes the distribution of variable values at time T . In addition, if a vector of regression coefficients \vec{b} is specified, then we have the random output variable:

$$W_T = \vec{b}^\top \sum_{t=1}^T A^{T-t} \vec{X}_t$$

In this case, the distribution of W_T will be called an **accumulated distribution**.

Alternatively, we can view an accumulation model in terms of a weighted, directed causal network whose adjacency matrix is A and a random vector \vec{X} on the nodes of the matrix. This format is quite flexible. For instance, we can extract x_i by setting $b_i = 1$ and $b_{j \neq i} = 0$. So statements about W_T can be used to reconstruct results about \vec{Z}_T .

In an empirical setting, data that arises from a multifactor cumulative advantage process can be thought of as repeated draws of the random variable W_T . So the distribution of W_T should match distributions of empirical data.

Since $\{\vec{X}_t\}$ are iid variables, accumulation models have similarities to the Central Limit Theorem. Indeed, if A is the identity matrix and $\vec{b} = \left\langle \frac{1}{M}, \frac{1}{M}, \dots, \frac{1}{M} \right\rangle$, then \vec{Z}_T is the mean of iid random variables and therefore distributed normally for large T due to the Central Limit Theorem.

4.2.4 Basic Properties of the Accumulation Model

The means and variances of the outcomes in an accumulation model can be written in terms of $E(\vec{X})$ and $Var(\vec{X})$.

$$E(\vec{Z}_T) = \sum_{t=1}^T A^{T-t} E(\vec{X})$$

$$E(W_T) = \vec{b}^\top \sum_{t=1}^T A^{T-t} E(\vec{X})$$

$$Var(\vec{Z}_T) = \sum_{t=1}^T A^{T-t} Var(\vec{X}) (A^{T-t})^\top$$

$$Var(W_T) = \vec{b}^\top \left[\sum_{t=1}^T A^{T-t} Var(\vec{X}) (A^{T-t})^\top \right] \vec{b}$$

4.2.5 Stories All the Way Down

One issue with using potentially unmeasurable variables is that no one need determine the precision with which the variable gets defined. Should we use a single variable for social skills, or should there be individual variables for written and verbal social skills? Or should it be further fine-grained, like “social skills when interacting with a teacher in a classroom context”? One strength of accumulation models is that they are consistent with finer-grained substitution of variables, as long as the relationships between these sets of variables obey some intuitive constraints. The next definition sets up these constraints in the context of a set of variables \mathcal{X}' which is finer-grained and “fits inside” another set of variables \mathcal{X} .

Definition: Assume that we have two sets of variables $\mathcal{X} = \{x_1, x_2, \dots, x_M\}$ and $\mathcal{X}' = \{x_1, x_2, \dots, x_{M'}\}$ with $M' > M$, and that there is a surjective map $\phi: \mathbb{R}^{M'} \rightarrow \mathbb{R}^M$. Further assume that (\vec{X}, A, T, \vec{b}) and $(\vec{X}', A', T, \vec{b}')$ describe accumulation models in \mathcal{X} and \mathcal{X}' , respectively. We say that $(\vec{X}', A', T, \vec{b}')$ is **finer than** (\vec{X}, A, T, \vec{b}) if:

- $\phi\vec{X}' = \vec{X}$ as random vectors,
- $\phi A' = A\phi$ as matrices, and
- $\vec{b}' \cdot \vec{v} = \vec{b} \cdot \phi\vec{v}$ for any $\vec{v} \in \mathbb{R}^{M'}$

The three conditions are equivalent to intuitive notions for how this map should behave. For instance, part of the map ϕ could be

$$\phi(\text{written skills}, \text{verbal skills}) = \text{written skills} + \text{verbal skills} = \text{communication skills}$$

If your written skills increase by 1 and your verbal skills increase by 2, then your communication skills should increase by 3. Similarly, both the accumulation processes A, A' and the regression function given by \vec{b} and \vec{b}' should be consistent with the relationship between variables.

Note: The proof of the following theorem and most other theorems are in the appendix.

Theorem (Stories All the Way Down): If one accumulation model is finer than another, then their outputs are equivalent. In other words $W_T = W'_T$.

4.2.6 Inequality in an Accumulation Model

We can estimate the amount of inequality in the accumulation model when T is large. To measure inequality, we will use the coefficient of variation, which is given by:

$$cv(X) = \frac{sd(X)}{E(X)} = \frac{\sqrt{Var(X)}}{E(X)}$$

The coefficient of variation is convenient for measuring inequality, since the mean and variance behave nicely with regards to sums and scalar multiplication. Like the Gini coefficient, the coefficient of variation is scale-independent. If we double everyone's wealth, then both the Gini coefficient and the coefficient of variation remain unchanged. We will see later how a non-scale-independent measure of inequality, the variance, grows almost exponentially over time.

Accumulated Inequality Theorem: Consider the accumulated random variable

$$W_T = \vec{b} \cdot \vec{Z}_T = \vec{b}^\top \sum_{t=1}^T A^{T-t} \vec{X}_t$$

Where W_T is a linear function of variables in the accumulation model. In this case, \vec{b} can be considered a vector representing regression coefficients, A corresponds to the primitive accumulation matrix, and $\vec{X}_t \sim \vec{X}$ are iid random vectors. Further, let \vec{u} and \vec{w} be the dominant left and right eigenvectors and a the dominant eigenvalue of A . Then for large T :

$$cv(W_T) \sim \frac{\sqrt{a-1}}{\sqrt{a+1}} cv(\vec{u} \cdot \vec{X})$$

Or, in terms of the variable $\psi = \frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} \vec{u} \cdot \vec{X}$ from the Accumulation Reduction Theorem.

$$cv(W_T) = \frac{\sqrt{a-1}}{\sqrt{a+1}} cv(\psi_T)$$

Interestingly, the inequality in the accumulated variable W_T is independent of the choice of \vec{b} . In other words, *any* linear combination of the variables in the model will have the same amount of inequality. The inequality in, say, x_1 will be the same as in $x_1 + x_2 + x_3 + \dots + x_M$.

The inequality in the accumulated variable W_T is a product of two quantities. One quantity, $cv(\vec{u} \cdot \vec{X})$, represents the inequality embedded in how the additive benefits \vec{X} (schooling, etc.) are distributed unequally, weighted by their usefulness in gaining other skills, traits, and resources.

The other quantity, $\frac{\sqrt{a-1}}{\sqrt{a+1}}$, is a monotone function of the multiplier a as shown in Figure 10 below. When $a = 1$, this factor is zero. In this case, Z_T is the sum of iid random variables ψ_t . For large T , the mean is $E(\psi) T$, and the standard deviation grows like $\sqrt{T} \sigma_\psi$. So $cv(Z_T) \sim \frac{\sqrt{T} \sigma_\psi}{TE(\psi)}$ which goes to zero as T gets large. When a is large, this factor goes to one. So the inequality in W_T is bounded above by the inequality in $\vec{u} \cdot \vec{X}$. In other words, *the inequality in an accumulated distribution will always be less than the inequality embedded in the process that generates it.*




Figure 10: The relationship between a and its effect on inequality.

While it may seem intuitive that cumulative advantage processes increase inequality, this need not be the case (Allison et al., 2018). In this model, the decrease in inequality for small a is caused by the weighted sums. Adding additional random variables decreases the variance relative to the mean, just as in the Central Limit Theorem.

4.3 Reduction to One Dimension

The accumulation model's generality allows it to describe a broad set of situations at the cost of specificity. While it can describe situations where 3 or 1,000 factors are involved in success, accurately inferring all the parameters in a 1,000 variable model is nigh impossible. Luckily, for large values of time, all accumulation models behave like simpler, one-dimensional accumulation models. This allows us to infer the behavior of the high-dimensional accumulation models by understanding how one-dimensional accumulated models behave.

A one-dimensional accumulation distribution is of the form:

$$Z = \sum_{t=1}^T a^{T-t} X_t$$

Where $a > 1$ is a real number and X_t are iid copies of a non-negative random variable X .

This section will use an important result from linear algebra, the Perron-Frobenius Theorem, to show that all accumulated distributions behave like a 1-dimensional accumulated distribution for large T . However, we first must justify the assumption of primitivity.

Recall that each factor x_i represents a resource, or a positive trait that individuals find useful. It's reasonable to assume that, to the extent possible, all of those resources will be reinvested in other resources and that having more of any resource will *eventually* help individuals in every other resource. For instance, being a fast runner may not directly help a student in school. However, being a good runner may encourage someone to run, which increases their overall health, which increases their performance in school. In terms of the causal network, this means there is a path from the node (running ability) to (academic achievement) with all positive coefficients.

So we assume that each variable causatively, positively, and perhaps indirectly influences the other variables over time. In terms of the matrix A , *direct* positive influence of i on j corresponds to the condition $A_{ji} > 0$. *Indirect* positive influence means that $A_{ji}^k > 0$ for large k . More specifically, we will assume that A is a primitive matrix.

Definition: A matrix A is primitive if there exists a $K > 0$, such that for every $k > K$, the matrix A^k has all positive entries.

Perron-Frobenius Theorem (C. R. Johnson & Tarazaga, 2004): Let A be a primitive matrix.

Then:

- a) A has a unique dominant real, positive eigenvalue a which is larger than the magnitude of all other eigenvalues.
- b) Up to scalar multiplication, there exists a unique right eigenvector \vec{w} corresponding to a (so that $A\vec{w} = a\vec{w}$) which can be chosen to have all positive entries.
- c) Up to scalar multiplication, there exists a unique left eigenvector \vec{u} corresponding to a (so that $\vec{u}^T A = a\vec{u}^T$) which can be chosen to have all positive entries

$$d) \lim_{T \rightarrow \infty} \frac{A^T}{a^T} = \frac{1}{\vec{w} \cdot \vec{u}} \vec{w} \vec{u}^\top$$

The Perron-Frobenius Theorem tells us that any primitive matrix A has a largest eigenvalue a and eigenvectors \vec{u} and \vec{w} , and that for large T , A^T is a map (almost completely) into the vector space spanned by \vec{w} . In other words, if there is a time-based process that involves multiplying a vector \vec{v} by A at each time step, then for large T , $A^T \vec{v} \sim a^T \left(\frac{\vec{u} \cdot \vec{v}}{\vec{w} \cdot \vec{u}} \right) \vec{w}$.

The next theorem uses this idea, but generalizes it to the sums inherent in the accumulation model. It says that, after controlling for exponential growth, accumulation models all eventually behave as a one-dimensional model.

Accumulation Reduction Theorem: All accumulated distributions of the random variable W_T converge in mean to a 1-dimensional accumulated distribution, after controlling for exponential growth. More precisely: Let \vec{X} be a non-negative random M -dimensional vector, A be a primitive matrix, and $\vec{b} \in \mathbb{R}^M$ a nonzero vector. So that

$$W_T = \vec{b}^\top \sum_{t=1}^T A^{T-t} X_t$$

Let a, \vec{u}, \vec{w} be the eigenvalue and eigenvectors of A described in the Perron-Frobenius Theorem.

Define the random variables $\psi_t = \frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} \vec{u} \cdot \vec{X}_t$ so that all ψ_t are independent, identically distributed random variables. Then we can create the one-dimensional accumulated random variable

$$Z_T = \sum_{t=1}^T a^{T-t} \psi_t$$

Under these conditions:

$$\lim_{T \rightarrow \infty} E \left(\frac{1}{a^T} |W_T - Z_T| \right) = 0$$

Corollary: For large T , the variables below grow according to the relationships:

$$\vec{Z}_T \sim \frac{1}{\vec{u} \cdot \vec{w}} \left[\sum_{t=1}^T a^{T-t} (\vec{u} \cdot \vec{X}_t) \right] \vec{w}$$

$$x_{i,T} \sim \frac{1}{\vec{u} \cdot \vec{w}} \left[\sum_{t=1}^T a^{T-t} (\vec{u} \cdot \vec{X}_t) \right] w_i$$

$$E(\vec{Z}_T) \sim \frac{1}{\vec{u} \cdot \vec{w}} E(\vec{u} \cdot \vec{X}) \left(\frac{a^T - 1}{a - 1} \right) \vec{w}$$

Specifically, for large T :

- Each x_i will tend to grow roughly exponentially over time, with a growth multiplier of a .
- An individual's growth will be larger if their random draws of \vec{X}_t lead to large dot products $\vec{u} \cdot \vec{X}_t$. In other words, \vec{u} gives the “weights” of each factor in causally influencing the long-term growth of \vec{Z}_T .
- The factors will grow proportionally to each other over time, with outcomes weighted according to the components of \vec{w} . I.e., for any two factors i, j , we know that $\frac{E(x_{i,T})}{E(x_{j,T})} \sim \frac{w_i}{w_j}$.

The Accumulation Reduction Theorem suggests that over time, the portfolio of skills, traits, and resources of people will tend to become similar to other people in the same circumstances.

Though some people may have more or less of these traits, the *ratio* between any two traits will tend toward the same value. Of course, this only applies for individuals who stay in the same situation, with similar values. In reality, people adapt their behavior and values to cultural, societal and economic forces in a complex, many-player game (Fe & Sanfelice, 2022; Kahneman, 2011; Markus & Hamedani, 2019).

4.4 One-Dimensional Accumulated Distributions

I have shown that M -dimensional accumulated models reduce to 1-dimensional accumulated distributions as time goes by. This section digs deeper into those **one-dimensional accumulated distributions**, providing tools to understand accumulated models more broadly. The random variable Z_T associated with a one-dimensional accumulated model (X, a, T) has the form:

$$Z_T = \sum_{t=1}^T a^{T-t} X_t$$

Where we assume $a > 1$. The goal here is to more closely connect distributions of data (represented by Z_T) with the processes that generate it (given by a and X). So I prove a number

of theorems about the *shape* and *mean* of a one-dimensional accumulated distribution. In particular, I show that:

- For large T , one-dimensional accumulated distributions can be uniquely characterized by their *mean* and their *shape*.
- The mean of a one-dimensional accumulated distribution grows roughly exponentially with time.
- The accumulation process $\vec{X} \rightarrow W_T \rightarrow Z_T$ preserves heavy tails (or the lack thereof), so that an accumulated distributions is heavy-tailed if and only if one of the components in the random vector \vec{X} is heavy-tailed.
- For a given $a > 1$ and T , there is a one-to-one relationships between the cumulants of X and the cumulants of its accumulated distribution Z_T .
 - If X is not heavy-tailed, then this creates a one-to-one relationship between X and Z_T .
- For a given $a > 1$ and random variable X , the shape of the accumulated distribution converges to a “shape” distribution given by Y_∞ . This shape is unique to X . In other words, two datasets generated by an accumulation process with the same shape and the same a will come from the same distribution X .
 - If X is not heavy-tailed, then this creates a one-to-one relationship between X and Y_∞ .
- Only normal distributions create exactly normal accumulated distributions.
 - Despite this, many accumulated distributions tend toward normality.

4.4.1 General Features of 1-Dimensional Accumulated Distributions

The formula for a one-dimensional accumulated distribution is quite similar to the formula for regular investments in an account with a constant interest rate.

$$Z_T = \sum_{t=1}^T a^{T-t} X_t$$

The value of X_t is the random amount an individual can invest at the beginning of each time period with an interest rate of $r = a - 1$. The constant can then be decomposed into $a = 1 + r$,

where the 1 ensures the initial investments remains in the account. This makes sense when thinking about accumulation models as describing investment in oneself.

It will be useful to create a few additional random variables, which will help us understand the shape of Z_T .

Definition: The random variables Y_T and Y_∞ associated with a one-dimensional accumulation model (X, a, T) are given by:

$$Y_T = \frac{a-1}{a^T-1} Z_T = \frac{a-1}{a^T-1} \sum_{t=1}^T a^{T-t} X_t$$

$$Y_\infty = (a-1) \sum_{t=1}^{\infty} \frac{1}{a^t} X_t$$

Theorem: The mean and variance of the accumulated distributions are:

$$\begin{aligned} \mu_{Z_T} &= \frac{a^T-1}{a-1} \mu_X & \sigma_{Z_T}^2 &= \frac{a^{2T}-1}{a^2-1} \sigma_X^2 \\ \mu_{Y_T} &= \mu_X & \sigma_{Y_T}^2 &= \frac{(a-1)(a^T+1)}{(a+1)(a^T-1)} \sigma_X^2 \\ \mu_{Y_\infty} &= \mu_X & \sigma_{Y_\infty}^2 &= \frac{a-1}{a+1} \sigma_X^2 \end{aligned}$$

These can be shown by using properties of the mean and variance: $\mu_{X_1+X_2} = \mu_{X_1} + \mu_{X_2}$, $\mu_{bX_1} = b\mu_{X_1}$, $\sigma_{X_1+X_2}^2 = \sigma_{X_1}^2 + \sigma_{X_2}^2$, and $\sigma_{bX}^2 = b^2\sigma_X^2$.

Stable Shape Theorem: If X has a finite mean and $a > 1$, then the sequence of random variables $\{Y_T\}$ converges in distribution to Y_∞ . More generally, the shape of Z_T stabilizes for large T .

These theorems explain why Y_T and Y_∞ correspond to the shape of an accumulated distribution.

We can rewrite $Z_T = \frac{1}{\mu_X} \mu_{Z_T} Y_T$. So that, given μ_X , Z_T is uniquely characterized by its mean μ_{Z_T} and its shape Y_T . Alternatively, we could have defined Y_T to have a mean of one, but that would make later calculations more complicated.

This decomposition allows us to see that the mean of Z_T grows almost exponentially with time.

By dividing by this mean, we can see that the shape of an accumulated distribution actually

converges to the distribution Y_∞ . So we can think of Y_∞ as describing the shape of Z_T after a long time has progressed – a situation which is common in research on systems which have been around for a while. Splitting up the mean and shape will let us develop tools for understanding accumulated distributions without having to repeatedly account for values that grow exponentially over time.

4.4.2 Characterizing 1-D Accumulated Distributions

If a researcher has data for a variable which they think represents multifactor cumulative advantage, they might want to understand more about the process that generated the data. In the language of our one-dimensional model: If a researcher has Z_T , what can they say about a and X ?

To build intuition, consider what happens as we vary a along the interval $[1, \infty)$. When a is very large, then $Z_T = \sum_{t=1}^T a^{T-t} X_t$ is dominated by the first term X_1 . For example, if $a \geq 2$, then the first draw, X_1 , is more influential, on average, than the rest of the draws combined. To see this, note that if $a = 2$:

$$Z_T = 2^{T-1}X_1 + \sum_{t=1}^{T-2} 2^t X_i$$

The right sum has a total average value of:

$$E\left(\sum_{t=1}^{T-2} 2^t X_i\right) = E(X) \sum_{t=1}^{T-2} 2^t = E(X)(2^{T-1} - 1)$$

So the average contribution of the first draw alone is $2^{T-1}E(X)$, while the remaining draws have a smaller total average contribution of $(2^{T-1} - 1)E(X)$. In this case, the distribution of outcome data will look very much like the variable that generated it, and individuals' outcomes will largely be determined by their start in life.

In contrast, when a is almost equal to 1, then Z_T is quite close to the sum of iid variables. In this case, the Central Limit Theorem says that Z_T will be distributed close to a normal distribution. Individuals' outcomes will be the result of a number of factors, and early-life circumstances will not have a significant effect on later outcomes.

We can think of a as a “dial” with a normal distribution on the $a = 1$ end, and a distribution similar to that of X on the $a > 2$ end. At both extremes, the multifactor cumulative advantage

process has little to add. However, some of the more interesting behavior happens, both mathematically and practically, when a is close to 1. If a is near 1, then individuals can reinvest their skills, traits, and abilities a little bit, but no one has an amazing advantage at birth.

To be more precise, this section will use the values of a and T as a cipher to generate a “dictionary” which relates X with Z_T , and with the shape variables Y_T , and Y_∞ . This dictionary is given by the Accumulated Moment Theorem, which first requires some machinery. Luckily, much of that machinery already exists.

Definition (Characteristic Function): The characteristic function of a random variable is $\varphi_X = \int_{\mathbb{R}} e^{itx} f_X(x) dx = E(e^{itX})$.

Properties of Characteristic Functions:

- a) $\varphi_{b_1X_1+b_2X_2}(s) = \varphi_{X_1}(b_1s)\varphi_{X_2}(b_2s)$
- b) $\varphi_X(0) = 1$
- c) $\varphi_X^{(k)}(0) = i^k E(X^k)$ where it exists
- d) The characteristic function $\varphi_X(s)$ uniquely determines the distribution of X

Definition (Cumulant Generating Function): We define the cumulant generating function of X in a neighborhood of $s = 0$ to be

$$G_X(s) = \log(\varphi_X(s))$$

The output of φ will be a complex number, so we will have to choose a branch of the complex plane to define this logarithm on. The only use this function will get in this manuscript is around $s = 0$. Since $\varphi_X(0) = 1$ and φ_X is defined in a neighborhood $\mathcal{U} \subset \mathbb{R}$ of $s = 0$, we can take this branch of the logarithm to be defined on the complex plane minus the negative real line.

$$G_X: \mathcal{U} \xrightarrow{\varphi_X} \mathbb{C} \setminus \{\mathbb{R}^{\leq 0}\} \xrightarrow{\log} \mathbb{C}$$

This may lead to $G_X(s)$ having a smaller domain than $\varphi_X(s)$, but it will give us the results we need in a neighborhood of $s = 0$.

What we are calling the cumulant generating function is called other things elsewhere, such as the “second characteristic” (Lukacs, 1970).

Definition (Cumulants): The cumulants of a distribution X are given in terms of the distribution's cumulant generating function:

$$\kappa_n(X) = -(-i)^n G_X^{(n)}(0)$$

Cumulants do not always exist. When they do, they can be written as polynomial functions of the central moments of X .

$$\kappa_1 = \mu_X$$

$$\kappa_2 = \text{Var}(X)$$

$$\kappa_3 = E((X - \mu_X)^3)$$

$$\kappa_4 = E((X - \mu_X)^4) - 3[\text{var}(X)]^2$$

Higher-order cumulants cannot generally be written elegantly in terms of central moments of X .

Properties of Cumulants: Let X_1, X_2 be real-valued random variables, and let b be a real number. Then:

- $\kappa_n(X_1 + X_2) = \kappa_n(X_1) + \kappa_n(X_2)$
- $\kappa_n(bX_1) = b^n \kappa_n(X_1)$

Properties of the Cumulant Generating Function: Let b_1, b_2 be positive real numbers, and let X, X_1 and X_2 be real-valued random variables. Then:

- $G_{b_1X_1+b_2X_2}(s) = G_{X_1}(b_1s) + G_{X_2}(b_2s)$
- $G_X(0) = 0$

The next property gives a nice way of finding the characteristic function of X if we know the characteristic function of Y_∞ .

Properties of Y_∞ : For a given a and X , if \tilde{Y}_∞ is independent of and identically distributed to Y_∞ , then:

$$\text{a) } Y_\infty \sim \frac{1}{a} \tilde{Y}_\infty + \frac{a-1}{a} X$$

$$\text{b) } \varphi_X((a-1)s) = \frac{\varphi_{Y_\infty}(as)}{\varphi_{Y_\infty}(s)}$$

Proof: The proof of (a) is a direct calculation.

$$\begin{aligned} \frac{a-1}{a} X + \frac{1}{a} \tilde{Y}_\infty &= \frac{a-1}{a} X + \frac{a-1}{a} \sum_{t=1}^T \frac{1}{a^t} X_t \\ &= (a-1) \left[\frac{1}{a} X + \frac{1}{a^2} X_1 + \frac{1}{a^3} X_2 + \dots \right] \sim (a-1) \sum_{t=1}^{\infty} \frac{1}{a^t} X_t = Y_\infty \end{aligned}$$

The proof of (b) follows from (a) and the properties of characteristic functions.

Corollary: If X_1 and X_2 are different distributions and $a > 1$, then the shape of the accumulated distribution Y_∞ associated with (X_1, a) and (X_2, a) are also different.

Alternatively: For a fixed $a > 1$, there is a one-to-one relationship between the distributions X and Y_∞ .

We can see the corollary is true, because Y_∞ uniquely determines the characteristic function of X which in turn uniquely determines X .

Example (Distributions that Keep their Shape when Accumulated): Consider the case where the accumulated distribution is distributed normally. Note that this case does not satisfy one of our assumptions, which is that X and therefore Y_∞ are positive random variables. However, this case will be illustrative regardless. The characteristic function of $Y_\infty \sim N(\mu, \sigma^2)$ is $\varphi_{Y_\infty}(s) = e^{-i\mu s + \frac{1}{2}\sigma^2 s^2}$. For a given a , we can then calculate the characteristic function of X .

$$\begin{aligned} \varphi_X((a-1)s) &= \frac{\varphi_{Y_\infty}(a s)}{\varphi_{Y_\infty}(s)} \\ &= \frac{e^{-i\mu a s + \frac{1}{2}\sigma^2 a^2 s^2}}{e^{-i\mu s + \frac{1}{2}\sigma^2 s^2}} \\ &= e^{-i\mu s(a-1) + \frac{1}{2}\sigma^2(a^2-1)s^2} \end{aligned}$$

Using the change of variables $(a-1)s = r$, we get:

$$\varphi_X(r) = e^{-i\mu r + \frac{1}{2}\sigma^2 \left(\frac{a^2-1}{(a-1)^2}\right) r^2}$$

This is the characteristic function of a normal distribution, so $X \sim N\left(\mu, \frac{a^2-1}{(a-1)^2} \sigma^2\right)$.

In retrospect, this shouldn't be surprising, since the sum of normal random variables is normal, and $Y_\infty = \lim_{T \rightarrow \infty} \sum_{t=1}^T \frac{a-1}{a^t} X_t$. If the original distribution is *exactly* normal, then the shape of the

accumulated distribution will also be normal. However as we'll see later, if a is *close* to 1, then the accumulated distribution of many distributions will be *close* to normal. But it will never become normal.

A similar argument shows that the Lévy distribution, with characteristic function $\varphi_X = e^{i\mu - \sqrt{-2ict}}$ also has this property: X is Lévy if and only if Y_∞ is Lévy. More generally, a random variable X is called strictly stable if a linear combination of two copies of X is equivalent in distribution to a constant times X (Nolan, 2020). Strictly stable variables have the property that they keep their shape after being accumulated.

Accumulated Moment Theorem: Let X be a non-negative real-valued random variable, $a > 1$, and T, Z_T, Y_T, Y_∞ as above. Further let $M_n(X)$ be the n th raw moment of X and $\kappa_n(X)$ be the n th cumulant. For a given n , all of the following exist if and only if any one of them does:

- | | |
|---------------------|--------------------------|
| (1) $M_n(X)$ | (2) $\kappa_n(X)$ |
| (3) $M_n(Z_T)$ | (4) $\kappa_n(Z_T)$ |
| (5) $M_n(Y_T)$ | (6) $\kappa_n(Y_T)$ |
| (7) $M_n(Y_\infty)$ | (8) $\kappa_n(Y_\infty)$ |

Furthermore, if these exist for a given n , then all lower-order moments and cumulants exist.

Also, the cumulants are related as follows:

$$\begin{aligned}\kappa_n(Z_T) &= \frac{a^{nT} - 1}{a^n - 1} \kappa_n(X) \\ \kappa_n(Y_T) &= \frac{(a^{nT} - 1)(a - 1)^n}{(a^T - 1)^n (a^n - 1)} \kappa_n(X) \\ \kappa_n(Y_\infty) &= \frac{(a - 1)^n}{a^n - 1} \kappa_n(X)\end{aligned}$$

While there is some disagreement over the definition of a heavy-tailed distribution, if we take the definition of a heavy-tailed function as one which has some non-existent moments, then we come to the following conclusion.

Corollary: An accumulated distribution is heavy-tailed if and only if the distribution that generated it is heavy-tailed.

A similar result holds more generally for M -dimensional accumulation models (\vec{X}, A, T) , where the additive amount at each time step is a vector $\vec{X} = \langle X_1, X_2, \dots, X_M \rangle$. The accumulated distribution given by W_T is heavy-tailed if and only if one of the components of \vec{X} is heavy-tailed. To see this, consider how the M -dimensional model reduces to a one-dimensional model, whose shape can then be described in terms of Y_∞ .

$$\sum_{t=1}^T A^{T-t} \vec{X}_t \rightarrow \sum_{t=1}^T a^{T-t} \psi_t \rightarrow Y_\infty$$

Where $\psi_t = \frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} \vec{u} \cdot \vec{X}_t$ is a weighted sum of the components of \vec{X}_t . Note that the sum of non-negative random variables has a heavy-tail if and only if *one* of the random variables has a heavy tail. So ψ_t is heavy-tailed if and only if one or more of the terms in $u_1 X_{1,t} + u_2 X_{2,t} + \dots + u_M X_{M,t}$ is heavy-tailed. Furthermore, since A is primitive, each component of \vec{u} is nonzero. So ψ_t (and hence Y_∞) is heavy-tailed iff one or more of the components of \vec{X}_t is a heavy-tailed variable. In other words, for large T , heavy tails in outcome data represent either a heavy tail in the distribution of some resources that individuals receive (X_i), or model breakdown.

4.4.3 Dictionary of Distributions

We now have a “dictionary” of cumulants to understand the relationship between the 1-dimensional accumulated distributions that correspond to distributions of outcome data, and the processes that generate them. For a given $a > 1$ and T , the following table gives the cumulants of X, Z_T, Y_T , and Y_∞ .

Table 4: Cumulants of Accumulated Distributions

n	X	Z_T	Y_T	Y_∞
1	μ_X	$\frac{a^T - 1}{a - 1} \mu_X$	μ_X	μ_X
2	σ_X^2	$\frac{a^{2T} - 1}{a^2 - 1} \sigma_X^2$	$\frac{(a^T + 1)(a - 1)}{(a^T - 1)(a + 1)} \sigma_X^2$	$\frac{a - 1}{a + 1} \sigma_X^2$
3	$\kappa_3(X)$	$\frac{a^{3T} - 1}{a^3 - 1} \kappa_3(X)$	$\frac{(a^{3T} - 1)(a - 1)^3}{(a^T - 1)^3 (a^3 - 1)} \kappa_n(X)$	$\frac{(a - 1)^2}{a^2 + a + 1} \kappa_n(X)$
n	$\kappa_n(X)$	$\frac{a^{nT} - 1}{a^n - 1} \kappa_n(X)$	$\frac{(a^{nT} - 1)(a - 1)^n}{(a^T - 1)^n (a^n - 1)} \kappa_n(X)$	$\frac{(a - 1)^n}{a^n - 1} \kappa_n(X)$

4.4.4 Dictionary of Distributions for $a \sim 1$

The formulae in Table 4 are fairly complicated. However, when a is close to one (though it will never be less than one), we can find approximations that are more interpretable. We will do this by expanding each coefficient in terms of $(a - 1)$, which will give us accurate behavior for a close to one.

The expression $\frac{1}{a^n - 1}$ shows up repeatedly in the formulas in Table 1. Since $a^n - 1 = (a - 1)(a^{n-1} + a^{n-2} + \dots + a + 1)$, we know that $\frac{1}{a^n - 1}$ has a pole of order one at 1. So we find a Laurent expansion:

$$\frac{1}{a^n - 1} = \frac{1}{n} \frac{1}{(a - 1)} + \left(\frac{1}{2n} - \frac{1}{2} \right) + o(a - 1)$$

Then we can expand the expressions in Table 4:

$$\begin{aligned} \kappa_n(Y_\infty) &= \frac{(a - 1)^n}{a^n - 1} \kappa_n(X) \\ &= (a - 1)^n \left[\frac{1}{n} \frac{1}{(a - 1)} + \left(\frac{1}{2n} - \frac{1}{2} \right) + o(a - 1) \right] \kappa_n(X) \\ &= \left[\frac{1}{n} (a - 1)^{n-1} + \left(\frac{1}{2n} - \frac{1}{2} \right) (a - 1)^n + o((a - 1)^{n+1}) \right] \kappa_n(X) \end{aligned}$$

To smallest order in $(a - 1)$:

$$\kappa_n(Y_\infty) = \frac{1}{n} (a - 1)^{n-1} \kappa_n(X)$$

We can create a similar approximation for Z_T by expanding $a^{nT} - 1$ around $a = 1$

$$a^{nT} - 1 = nT(a - 1) + \frac{1}{2}nT(nT - 1)(a - 1)^2 + o((a - 1)^3)$$

So that

$$\begin{aligned} \kappa_n(Z_T) &= (a^{nT} - 1) \left(\frac{1}{a^n - 1} \right) \kappa_n(X) \\ &= \left[nT(a - 1) + \frac{1}{2}nT(nT - 1)(a - 1)^2 + o((a - 1)^3) \right] \left[\frac{1}{n(a - 1)} + \left(\frac{1}{2n} - \frac{1}{2} \right) \right. \\ &\quad \left. + o(a - 1) \right] \kappa_n(X) \\ &= T + \left[nT \left(\frac{1}{2n} - \frac{1}{2} \right) + \frac{1}{2}nT(nT - 1) \left(\frac{1}{n} \right) \right] (a - 1) + o((a - 1)^2) \kappa_n(X) \\ \kappa_n(Z_T) &= \left[T + \frac{1}{2}nT(T - 1)(a - 1) + o((a - 1)^2) \right] \kappa_n(X) \end{aligned}$$

The first order relationship $\kappa_n(Z_T) \sim T\kappa_n(X)$ is not incredibly surprising, at least for the first three cumulants. For $n \leq 3$, it's true that $\kappa_n(X_1 + X_2) = \kappa_n(X_1) + \kappa_n(X_2)$. And when a is very close to 1, so that $a^t \sim 1$:

$$\begin{aligned} \kappa_n(Z_T) &= \sum_{t=1}^T \kappa_n(a^t X_t) \\ &\sim \sum_{t=1}^T \kappa_n(X_t) \\ &= \sum_{t=1}^T \kappa_n(X) \\ &= T\kappa_n(X) \end{aligned}$$

A more precise approximation might be

$$\kappa_n(Z_T) = \left[T + \frac{1}{2}nT(T - 1)(a - 1) \right] \kappa_n(X)$$

A similar calculation shows that, to lowest order, $\kappa_n(Y_T) = \frac{1}{T^{n-1}} \kappa_n(X)$.

Table 5: Approximate Cumulants of Accumulated Distribution for a~1

n	X	Z_T	Y_∞
1	μ_X	μ_X	μ_X
2	σ_X^2	$[T + T(T - 1)(a - 1)]\sigma_X^2$	$\frac{1}{2}(a - 1)\sigma_X^2$
3	$\kappa_3(X)$	$\left[T + \frac{3}{2}T(T - 1)(a - 1)\right]\kappa_3(X)$	$\frac{1}{3}(a - 1)^2\kappa_3(X)$
n	$\kappa_n(X)$	$\left[T + \frac{1}{2}nT(T - 1)(a - 1)\right]\kappa_n(X)$	$\frac{1}{n}(a - 1)^n\kappa_n(X)$

We can understand more about accumulated distributions by comparing with the normal distribution, which has only two nonzero cumulants $\kappa_1 = \mu$ and $\kappa_2 = \sigma^2$.

- If the cumulants of X grow slowly with n , then the accumulated distribution will be close to normal. To see this, consider $a = 1 + \epsilon$ where $\epsilon < 1$ is reasonably small. Then the n th cumulant of Y_∞ is $\frac{1}{n}\epsilon^n\kappa_n(X)$. If the cumulants of X do not grow too fast, then this expression quickly gets close to zero for large n . So the distribution of Y_∞ will be close to normal.
- The $n \geq 3$ cumulants of a normal distribution are all zero. Since the cumulants $\kappa_n(Y_\infty)$ are multiples of $\kappa_n(X)$, this tells us again that X is normal if and only if Y_∞ is normal.
- More generally, the table shows that $\kappa_n(Y_\infty) < \kappa_n(X)$ for $n > 1$. So the shape of the accumulated distribution is “more normal” than the shape of the original distribution X .

4.5 Examples

Another way to explore the relationship between generating distributions X , which might represent training or opportunities, and accumulated distributions Z_T, Y_T, Y_∞ , which correspond to outcome data, is through numerical simulations. This section gives the results of numerical simulations for some accumulated distributions, along with the algebraic derivation of the accumulated exponential distribution functions. Figure 11 gives the US income distribution in 2019, which is similar to income distributions in developed countries around the world (Tao et al., 2019). Though this work does not directly address distribution-fitting, such as through

maximum likelihood estimation, some of the simulation results seem to qualitatively match distributions of income.




Figure 11: Income distribution in the United States for individuals between ages 25 and 65. Data from 2019 American Community Survey.

4.5.1 Numerical Simulation: Accumulated Bernoulli

Imagine a scenario where, at each time point, an individual either receives a benefit of 1 with probability p , and 0 the rest of the time. The benefit could be the occurrence of a professional connection, a published research paper, or other just a lucky break. In this case, X is a Bernoulli random variable.

$$f_X(x) = \begin{cases} 1 & \text{if } 0 \leq x \leq p \\ 0 & \text{if } x > p \end{cases}$$

The accumulated Bernoulli distribution is practically interesting, because it models the situation where the additive amounts at each time step are discrete. This case is mathematically interesting, because it demonstrates how the accumulation process behaves for various values of a . In the examples here, $p = 0.1$. Unless explicitly labeled, vertical axes are not comparable.




Figure 12: Accumulated Bernoulli Y_T for various values of T , with $a = 1.01$ and $p = 0.1$.

The line on the left of each graph for $T \leq 50$ is not an artifact. This corresponds to people who got no successes on their draws. The graph shows the normalized distribution Y_T , rather than the non-normalized Z_T . So, in the $T = 5$ graph, the peak at 0.2 corresponds to people who've gotten one "success" in their five draws of the random variable X . The peaks for $T \leq 30$ get wider as T increases, which arises from the fact that individuals who got a success earlier get more "interest". Since $a = 1.01$ is close to 1, the distribution eventually becomes close to normal. Also notice that the distribution stabilizes for large T , as predicted by the Stable Shape Theorem.

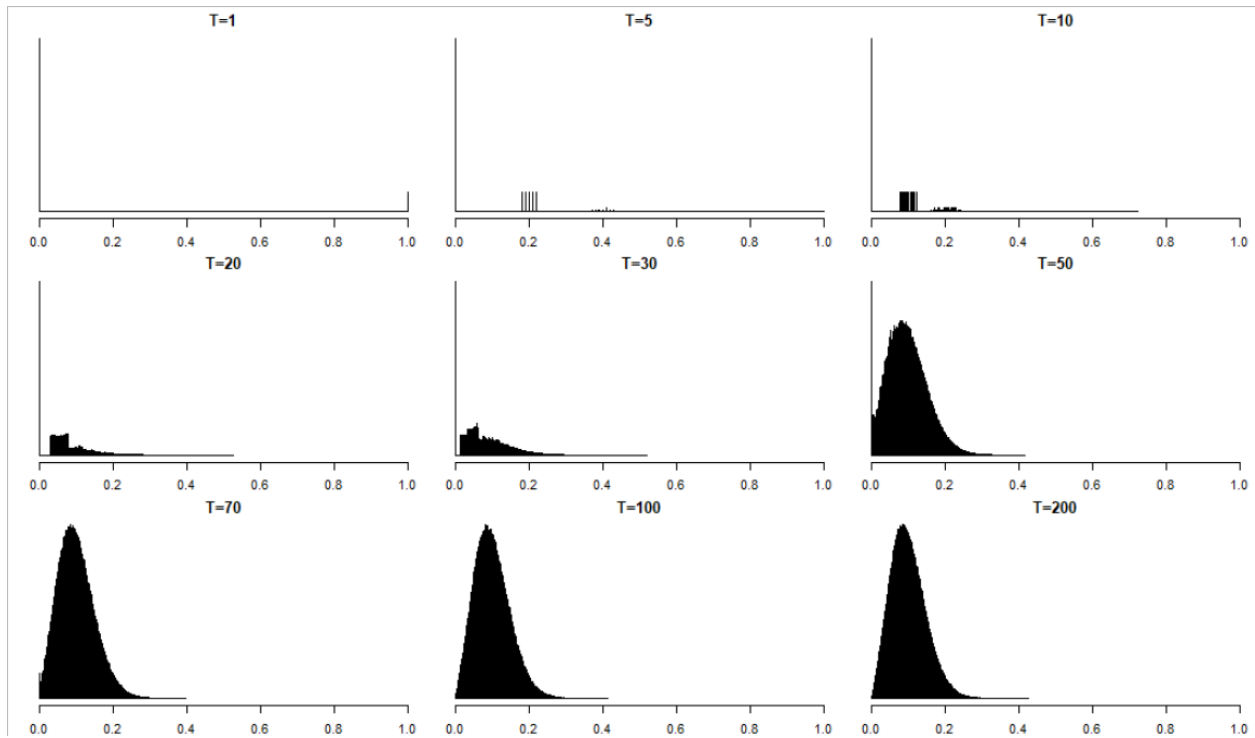


Figure 13: Accumulated Bernoulli Y_T for various values of T , with $a = 1.05$ and $p = 0.1$

Again, note the marked peaks at $x = 0$. The trends here are similar to the $a = 1.01$ case, except there is more variation around the peaks in $T = 5$. This arises because $a = 1.05$ gives a larger relative benefit to early successes in X than $a = 1.01$ does.




Figure 14: Accumulated Bernoulli Y_T for various values of T , with $a = 1.2$ and $p = 0.1$

We can see here the effects of $a = 1.2$ by noticing relatively large proportion of individuals near zero, even for large T . When $T = 200$, 15% of individuals have values below the first percentile, and all of them have at least one success. With a 20% interest rate, individuals with early successes end up being much more successful than those who don't. The peaks in the later graphs are stable, and each corresponds to individuals who got a success at a given early timepoint.




Figure 15: Accumulated Bernoulli Y_∞ for various values of a , with $p = 0.1$. Plotted with a simulated probability distribution function.

Recall that the mean in each graph is 0.1. The values of Z_T grow near-exponentially with T . As a increases, the normal distribution gives way to a distribution where most are proportionally close to zero. In the $a = 1.01$ regime, an individual's value of Y_T arises primarily from the *additive* process $x \rightarrow x + X_t$. In the large a regime, an individual's relative value comes from the *multiplicative* process $x \rightarrow ax$.

4.5.2 Numerical Simulation: Accumulated Normal Distribution

If the additive bonus X at each time step comes from a large number of potentially small opportunities, then X could be normally distributed. For instance, children might receive an increase in knowledge through schooling every year, with the amount gained being normally distributed. Technically, the normal distribution is excluded by our assumption that X only take on non-negative values. Instead, we can use a truncated normal distribution with probability distribution function

$$f_X(x) = \begin{cases} 0 & \text{if } x < 0 \\ C f_W(x) & \text{if } x \geq 0 \end{cases}$$

Here $W \sim \text{Norm}(1, 0.04)$ and C is chosen to ensure that the integral is one. The random variable X is non-negative and matches the normal distribution W on .9999997 of its mass.

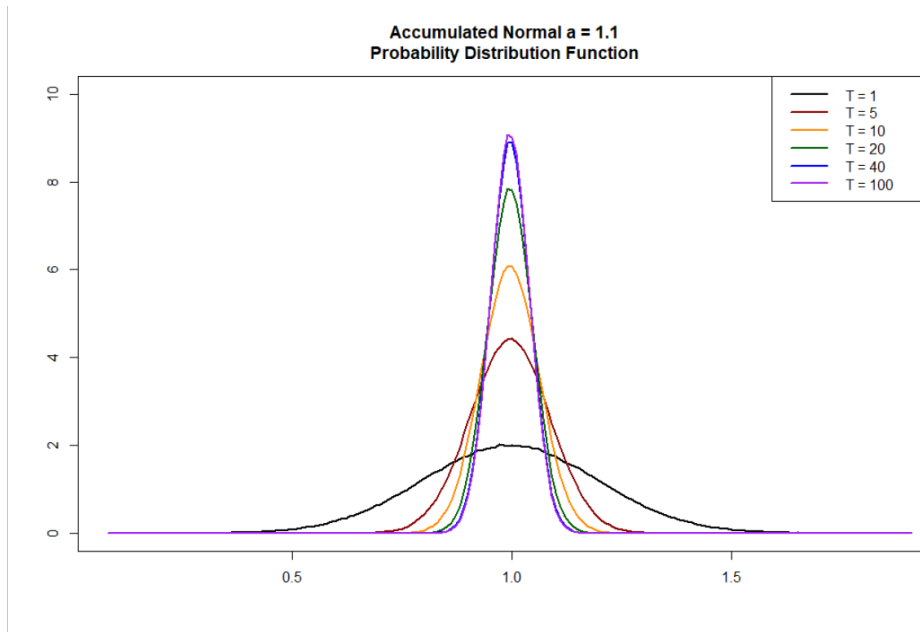


Figure 16: Accumulated normal distribution Y_T for various values of T with $a = 1.1$, $\mu = 1$, $\sigma = 0.2$

The variance is also the second cumulant. As predicted by the cumulant table, the variance decreases as T increases until it reaches a limiting value.

4.5.3 Numerical Simulation: Accumulated Pareto

It might be the case that most individuals receive only a little bonus at each time step, while a small number get very lucky. This would lead to heavy-tailed distributions, which are quite common in economic data (Davies & Shorrocks, 2000; Heinrich Mora et al., 2021) and measures of popularity (Myers et al., 2014; Price, 1976) To understand heavy-tailed accumulated distributions this simulation uses a Pareto distribution $P(X < x) = \frac{C}{(x+1)^2}$ which has infinite variance.

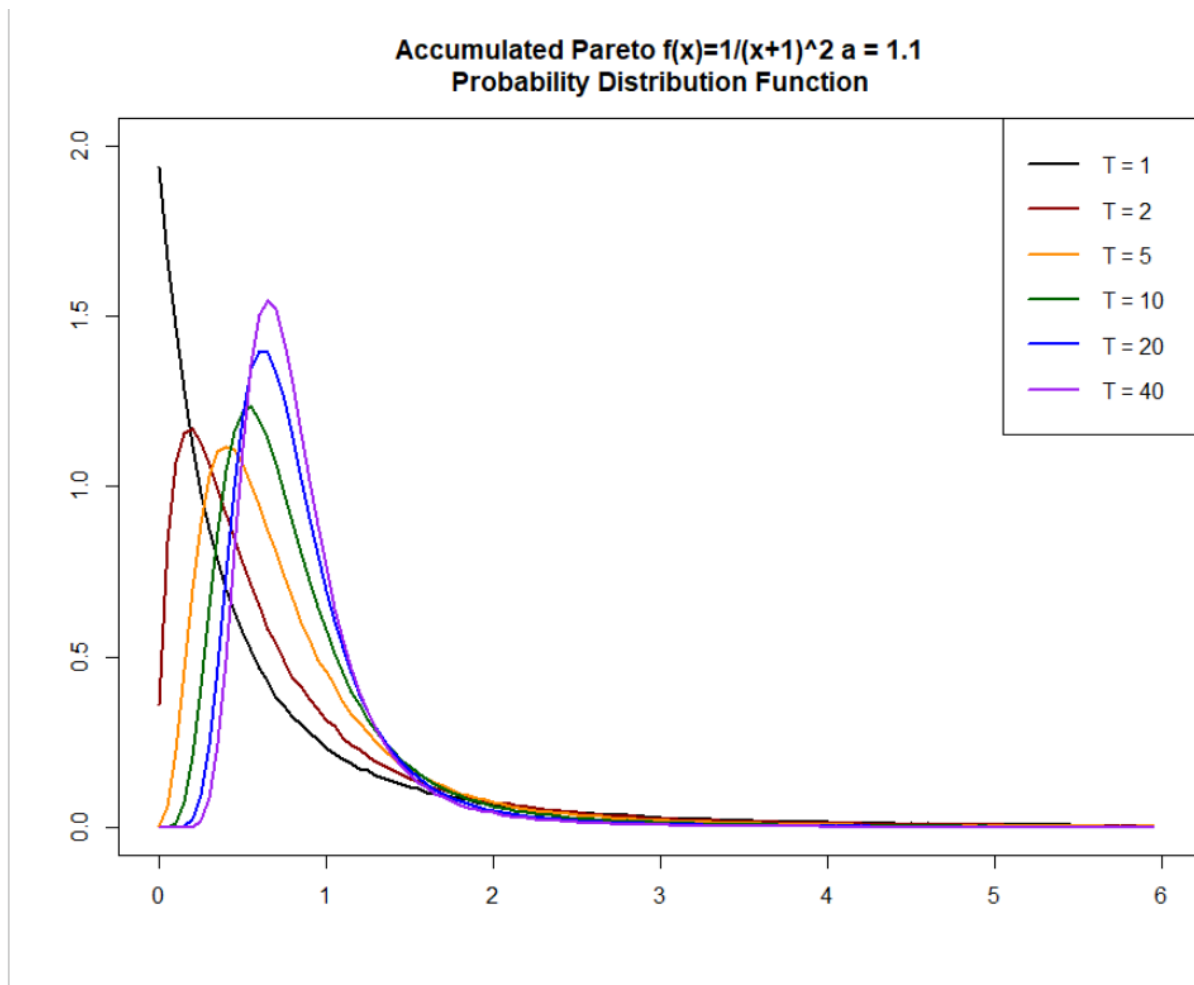


Figure 17: Accumulated Pareto distribution Y_T for various values of T with exponent 2 and $a = 1.1$. Plotted with a simulated probability distribution function.

Notice that the mode gradually moves right as T increases, despite the mean remaining constant. This is due to the long tail of the Pareto distribution contracting as T increases. Regardless of T , however, the accumulated Pareto remains a heavy-tailed distribution.

4.5.4 Numerical Simulation: Accumulated Exponential

In many cases, the probability of gaining a large bonus will be less than the probability of gaining a small bonus. Pareto distributions have this property, and have heavy tails. The exponential distribution $f_X(x) = se^{-sx}$ is a good example of a light-tailed random variable with a monotonically decreasing probability distribution function.




Figure 18: Accumulated exponential distribution Y_∞ for various values of a with constant $s = 1$. Plotted with a simulated probability distribution function.

4.5.5 Explicit Distribution Functions of Accumulated Exponential

The accumulated exponential provides one of the few distributions which allows an explicit formula for its distribution functions. In fact, the accumulated exponential distribution has a special property: Recall that $Z_T = \sum_{t=1}^T a^{T-t} X_t$, so that $f_{Z_T}(x) = f_X * f_{aX} * f_{a^2 X} * \dots * f_{a^{T-1} X}(x)$, where $*$ is the convolution operator.

$$(f * g)(x) = \int_{-\infty}^{\infty} f(z)g(x - z)dz$$

I will show that, for the accumulated exponential, we can write this convolution as a sum in terms of the component distribution functions $f_{Z_T}(x) = \sum_{k=0}^{T-1} c_{k,T} f_{a^k X}(x)$ for some constants $c_{k,T}$. This is quite an elegant result, since it allows us to write the convolution of exponential functions as a sum of those functions. Convolutions are very difficult to calculate, while sums are straightforward.

To show this, we need some machinery. Accumulated random variables are defined as linear combinations of random variables, which means we need to turn our attention to some well-known properties of convolutions.

Proposition: Let X and Y be independent random variables, $b \in \mathbb{R}$, and $*$ the convolution operator. Then:

$$1) f_{X+Y}(x) = (f_X * f_Y)(x) = \int_{\Omega} f_X(z) f_Y(x-z) dz, \text{ and}$$

$$2) f_{bX} = \frac{1}{b} f\left(\frac{x}{b}\right)$$

Where, for a fixed value of x , Ω is the shared domain of z -values of $f_X(z)$ and $f_Y(x-z)$.

Definition: The q -Pochhammer symbols are given by:

$$(b; q)_n = \prod_{k=0}^{n-1} (1 - bq^k)$$

$$(b; q)_{\infty} = \prod_{k=0}^{\infty} (1 - bq^k)$$

Note that q must be less than one in order for the product to converge. We will use a specific instance of the q -Pochhammer symbol, where $b = q$.

$$(q; q)_n = \prod_{k=1}^n (1 - q^k)$$

$$(q; q)_{\infty} = \prod_{k=1}^{\infty} (1 - q^k)$$

We also assume the convention that $(q; q)_0 = 1$.

Theorem (PDF of Accumulated Exponential): Let X be an exponentially distributed random variable so that

$$f_X(x) = \begin{cases} se^{-sx} & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases}$$

Then the probability distribution functions Z_T, Y_T, Y_{∞} for the accumulated exponential are

$$f_{Z_T}(x) = \sum_{k=0}^{T-1} \frac{s}{a^k} c_{k,T} e^{-\frac{s}{a^k}x} = \sum_{k=0}^{T-1} c_{k,T} f_{a^k X}(x)$$

$$f_{Y_T}(x) = \sum_{k=0}^{T-1} \frac{s\gamma_T}{a^k} c_{k,T} e^{-\frac{s\gamma_T x}{a^k}} = \sum_{k=0}^{T-1} c_{k,T} f_{\frac{a^k}{\gamma_T} X}(x)$$

$$f_{Y_\infty} = \sum_{k=0}^{\infty} \left(\frac{a}{a-1}\right) a^k s d_k e^{-\left(\frac{a}{a-1}\right) a^k s x} = \sum_{k=0}^{\infty} d_k f_{\frac{a-1}{a^{k+1}} X}(x)$$

Where $\gamma_T = \frac{a^T - 1}{a - 1}$ and

$$c_{k,T} = \frac{(-1)^{T-k-1}}{\left(\frac{1}{a}; \frac{1}{a}\right)_k \left(\frac{1}{a}; \frac{1}{a}\right)_{T-k-1}} a^{-\frac{(T-k)(T-k-1)}{2}}$$

$$d_k = \frac{(-1)^k}{\left(\frac{1}{a}; \frac{1}{a}\right)_k \left(\frac{1}{a}; \frac{1}{a}\right)_\infty} a^{-\frac{k(k+1)}{2}}$$

With appropriate assumptions when k is 0 or $T - 1$, we can also write:

$$c_{k,T} = a^{\frac{k(k+1)}{2}} \frac{1}{(a^k - 1)} \frac{1}{(a^{k-2} - 1)} \cdots \frac{1}{(a - 1)} \cdot 1 \cdot \frac{1}{(1 - a^1)} \frac{1}{(1 - a^2)} \cdots \frac{1}{(1 - a^{T-k-1})}$$

While the first expression for $c_{k,T}$ is more elegant, the second expression is potentially more computable since both the numerator and denominator are bounded for large T .

The next theorem follows directly from the previous one, and the fact that $\int_x^\infty e^{-bz} dz = \frac{1}{b} e^{-bx}$.

Theorem (CCDF of Accumulated Exponential): Let X be an exponentially distributed random variable with parameter s . And let $\bar{F}_{Z_T}(x) = P(Z_T > x) = \int_x^\infty f_{Z_T}(x) dx$ denote the complementary cumulative distribution functions (ccdf). Then

$$\bar{F}_{Z_T}(x) = \sum_{k=0}^{T-1} c_{k,T} e^{-\frac{s}{a^k} x} = \sum_{k=0}^{T-1} c_{k,T} \bar{F}_{a^k X}(x)$$

$$\bar{F}_{Y_T}(x) = \sum_{k=0}^{T-1} c_{k,T} e^{-\frac{s(a^T - 1)}{a^k(a-1)} x} = \sum_{k=0}^{T-1} c_{k,T} \bar{F}_{\frac{a^k}{\gamma_T} X}(x)$$

$$\bar{F}_{Y_\infty}(x) = \sum_{k=0}^{\infty} d_k e^{-\left(\frac{a}{a-1}\right) a^k s x} = \sum_{k=0}^{\infty} d_k \bar{F}_{\frac{a-1}{a^{k+1}} X}(x)$$

The cumulative distribution functions $F(x) = 1 - \bar{F}(x)$ can be calculated using the ccdf formulae.

4.6 Conclusion

In this chapter, I have developed a model of a high-dimensional cumulative advantage (or disadvantage) process. I calculated basic properties of this process and its outcome distributions, and showed that the high-dimensional model can be approximated by a one-dimensional model.

Specifically I found:

- Formulas for the output variables in the model, including the mean and variance.
- The model is consistent with respect to intuitive “changes of variables” to accommodate the idea of stories all the way down.
- For large time, multifactor cumulative advantage processes can be approximated by a one-dimensional process, with provides a single “systemic interest rate” for the many-dimensional process.
- Each factor influencing success has a weight which determines its causal strength in affecting well-being (determined by the left eigenvector of A). Individuals who increase the heavily-weighted factors early on tend to be the most well-off.
- Another set of weights (given by the dominant right eigenvector of A) describes how the factors grow in proportion to each other over time, and how strong of an indicator each factor is of overall well-being.
- For a given systemic interest rate a and time T , there is a one-to-one relationship between accumulated distributions corresponding to outcomes and the distribution of the additive process that generated the data.

I also gave examples of various one-dimensional accumulated distributions, and explicitly derived the functional form of accumulated exponential distributions.

Notably, the examples of accumulated distributions tend to have distribution functions which are skewed right with either a single peak, or decreasing from zero. These are similar to distributions of data that show up in student capital (Quarles et al., 2020), income (Bandourian et al., 2002), frailty (Rockwood et al., 2004), and mental health (Tomitaka et al., 2018). While an explicit method for data-fitting using a technique like maximum likelihood estimation is outside the

scope of this paper, these connections do suggest promise in using the accumulation model to understand the processes generating standard metrics of well-being.

4.6.1 Future Work

The current chapter provides a number of evidence-based, conceptual, and statistical opportunities for future research. A method could be developed to connect the model to distributions of data. Such a model would, for a given one-dimensional variable, infer a family of distributions X and a parameter a . This could generate a deeper understanding of how well-being (or negative processes like aging) accumulate in a population. For instance, a distribution-fitting method might provide a most reasonable systemic interest rate, helping us understand where the rich are getting richer fastest.

In the context of socioeconomic variables, the fact that the high-dimensional accumulation model reduces to one-dimension is interesting. It suggests that there might be a single-dimensional construct embedded in a high-dimensional socioeconomic dataset, which could be found using PCA or some other dimension reduction technique. This construct would likely represent socioeconomic status (SES). Interestingly, the model suggests that SES arises as a useful tool *because those are the traits that people value and can reinvest in*. For instance, income and health are correlated because (a) people value both traits, (b) people can turn their money toward better health, and (c) people with better health can earn more money. Repeat this argument for every pair of positively construed variables, and there is likely to be a latent dimension of well-being which is socially constructed by individuals' choices. Future work could examine socioeconomic status in the context of this model and high-dimensional data. In addition, there is some statistical refinement that could be done on the model itself. To generate example distribution, I primarily relied on numerical simulations. However, it is not clear how to numerically simulate Y_∞ , which is given by an infinite sum. So I used Y_T to approximate Y_∞ , since $Y_T \rightarrow Y_\infty$ for large T . In practice, this convergence seems to happen quite quickly. This is good, because simulations of Y_T can take a long time, which grows linearly with T . Future work could examine *how quickly* the convergence happens, to better enable simulations of the stable distribution Y_∞ while minimizing computational time.

I've used cumulants to show that, when all the cumulants exist, there is a one-to-one relationship between X and Y_∞ for a given value of a . I hypothesize that this will also be true when all the cumulants do not exist (i.e. when X and Y_∞ are heavy-tailed). Since many economic variables are heavy-tailed, this would be an interesting area for future examination.

4.7 Proofs of Major Theorems

Proof of the Stories All the Way Down Theorem: We want to show that, if $(\vec{X}', A', T, \vec{b}')$ is finer than (\vec{X}, A, T, \vec{b}) , then the two expressions below are equivalent.

$$W_T = \vec{b} \cdot \sum_{t=1}^T A^{T-t} \vec{X}_t$$

$$W'_T = \vec{b}' \cdot \sum_{t=1}^T (A')^{T-t} \vec{X}'_t$$

The calculation is straightforward. We start by using $\phi \vec{X}' = \vec{X}$.

$$W_T = \vec{b} \cdot \sum_{t=1}^T A^{T-t} \vec{X}_t$$

$$= \vec{b} \cdot \sum_{t=1}^T A^{T-t} \phi \vec{X}'_t$$

Since $\phi A' = A \phi$, we also know that $\phi (A')^{T-t} = A^{T-t} \phi$.

$$= \vec{b} \cdot \sum_{t=1}^T \phi (A')^{T-t} \vec{X}'_t$$

$$= \vec{b} \cdot \phi \sum_{t=1}^T (A')^{T-t} \vec{X}'_t$$

By the third property, $\vec{b}' \cdot \vec{v} = \vec{b} \cdot \phi \vec{v}$ for any $\vec{v} \in \mathbb{R}^{M'}$. So

$$= \vec{b}' \cdot \sum_{t=1}^T (A')^{T-t} \vec{X}'_t = W'_T$$

This concludes the proof.

Proof of Accumulated Inequality Theorem:

From the Accumulation Reduction Theorem, for large T ,

$$W_T \sim \frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} \left[\sum_{t=1}^T a^{T-t} (\vec{u} \cdot \vec{X}_t) \right]$$

To calculate the coefficient of variation, we first look at the mean, variance, and standard deviation of Z_T .

$$\begin{aligned} E(W_T) &\sim \frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} E(\vec{u} \cdot \vec{X}) \left(\frac{a^T - 1}{a - 1} \right) \\ \text{Var}(W_T) &= \left(\frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} \right)^2 \text{Var} \left[\sum_{t=1}^T a^{T-t} (\vec{u} \cdot \vec{X}_t) \right] \\ &= \left(\frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} \right)^2 \left[\sum_{t=1}^T a^{2(T-t)} \text{Var}(\vec{u} \cdot \vec{X}_t) \right] \\ &= \left(\frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} \right)^2 \left[\sum_{t=1}^T a^{2(T-t)} \right] \text{Var}(\vec{u} \cdot \vec{X}) \\ \text{Var}(W_T) &= \left(\frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} \right)^2 \left[\frac{a^{2T} - 1}{a^2 - 1} \right] \text{Var}(\vec{u} \cdot \vec{X}) \\ \sigma_{W_T} &= \sqrt{\text{Var}(W_T)} \\ \sigma_{W_T} &= \left(\frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} \right) \sqrt{\frac{a^{2T} - 1}{a^2 - 1}} \sigma_{\vec{u} \cdot \vec{X}} \end{aligned}$$

This means that the coefficient of variation is:

$$\begin{aligned} cv(W_T) &= \frac{\sqrt{\text{Var}(Z_T)}}{E(Z_T)} \\ &= \frac{\left(\frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} \right) \sqrt{\frac{a^{2T} - 1}{a^2 - 1}} \sqrt{\text{Var}(\vec{u} \cdot \vec{X})}}{\frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} E(\vec{u} \cdot \vec{X}) \left(\frac{a^T - 1}{a - 1} \right)} \\ &= \frac{\sqrt{\frac{a^{2T} - 1}{a^2 - 1}} \sqrt{\text{Var}(\vec{u} \cdot \vec{X})}}{\left(\frac{a^T - 1}{a - 1} \right) E(\vec{u} \cdot \vec{X})} \\ &= \frac{\sqrt{\frac{a^{2T} - 1}{a^2 - 1}} \sqrt{\text{Var}(\vec{u} \cdot \vec{X})}}{\left(\frac{a^T - 1}{a - 1} \right) E(\vec{u} \cdot \vec{X})} \end{aligned}$$

We can reduce the first fraction.

$$\begin{aligned}
\frac{\sqrt{\frac{a^{2T}-1}{a^2-1}}}{\left(\frac{a^T-1}{a-1}\right)} &= \left(\frac{a-1}{a^T-1}\right) \sqrt{\frac{a^{2T}-1}{a^2-1}} \\
&= \frac{(a-1)\sqrt{a^T-1}\sqrt{a^T+1}}{(a^T-1)\sqrt{a-1}\sqrt{a+1}} \\
&= \frac{\sqrt{a-1}\sqrt{a^T+1}}{\sqrt{a+1}\sqrt{a^T-1}} \\
&= \frac{\sqrt{a-1}\sqrt{1+\frac{1}{a^T}}}{\sqrt{a+1}\sqrt{1-\frac{1}{a^T}}}
\end{aligned}$$

For large T this goes to $\frac{\sqrt{a-1}}{\sqrt{a+1}}$. And $\frac{\sqrt{\text{Var}(\vec{u} \cdot \vec{X})}}{E(\vec{u} \cdot \vec{X})} = cv(\vec{u} \cdot \vec{X})$, which gives us

$$cv(Z_T) \sim \frac{\sqrt{a-1}}{\sqrt{a+1}} cv(\vec{u} \cdot \vec{X})$$

Since the coefficient of variation is unitless, we can also write this in terms of $\psi = \frac{\vec{b} \cdot \vec{w}}{\vec{u} \cdot \vec{w}} \vec{u} \cdot \vec{X}$.

$$cv(Z_T) = \frac{\sqrt{a-1}}{\sqrt{a+1}} cv(\psi_T)$$

This concludes the proof.

Proof of the Accumulation Reduction Theorem

To prove the Accumulation Reduction Theorem, we need a few lemmas.

Lemma: Let A be a matrix, with a left eigenpair (\vec{u}, a) and right eigenpair (\vec{v}, λ) , where $\lambda \neq a$. Then $\vec{u} \cdot \vec{v} = 0$.

Proof: Because \vec{u} is a right eigenvector with eigenvalue a , $\vec{u}^\top A \vec{v} = \vec{u}^\top a \vec{v} = a \vec{u} \cdot \vec{v}$. But also $\vec{u}^\top A \vec{v} = \vec{u}^\top \lambda \vec{v} = \lambda \vec{u} \cdot \vec{v}$. Therefore $a \vec{u} \cdot \vec{v} = \lambda \vec{u} \cdot \vec{v}$. Since $\lambda \neq a$, we know that $\vec{u} \cdot \vec{v} = 0$.

Lemma: Assume A is a primitive matrix, and a, \vec{u}, \vec{w} as in the Perron-Frobenius Theorem. Then the matrix $B = A - \frac{a}{\vec{u} \cdot \vec{w}} \vec{w} \vec{u}^\top$ has all the same eigenvectors and eigenvalues of A , except that the eigenvalue corresponding to \vec{u} and \vec{w} is 0 instead of a . Furthermore, the spectral radius of B , $\rho(B)$, (which is the maximum modulus of B 's eigenvalues) is strictly less than a . In other words, if λ is an eigenvalue of B , then $|\lambda| < a$.

Proof: We will first show that \vec{w} is a right eigenvector of B with eigenvalue 0.

$$B\vec{w} = \left(A - \frac{a}{\vec{u} \cdot \vec{w}} \vec{w} \vec{u}^\top\right) \vec{w}$$

$$\begin{aligned}
&= A\vec{w} - \frac{a}{\vec{u} \cdot \vec{w}} \vec{w}\vec{u}^\top \vec{w} \\
&= a\vec{w} - \frac{a}{\vec{u} \cdot \vec{w}} (\vec{u}^\top \vec{w}) \vec{w} \\
&= a\vec{w} - a\vec{w} = \mathbf{0}\vec{w}
\end{aligned}$$

If (\vec{v}, λ) is any other right eigenpair, then:

$$\begin{aligned}
B\vec{v} &= \left(A - \frac{a}{\vec{u} \cdot \vec{w}} \vec{w}\vec{u}^\top \right) \vec{v} \\
&= A\vec{v} - \frac{a}{\vec{u} \cdot \vec{w}} \vec{w}\vec{u}^\top \vec{v}
\end{aligned}$$

By the previous lemma, $\vec{u}^\top \vec{v} = 0$. So this is

$$\begin{aligned}
&= A\vec{v} - \mathbf{0} \\
&= A\vec{v} = \lambda\vec{v}
\end{aligned}$$

The proofs for the left eigenvector are similar. Since the spectral radius of a matrix is the magnitude of its largest eigenvalue the Perron-Frobenius Theorem says that $\rho(A) = a$. However, since a is not an eigenvalue of B , the spectral radius of B is the magnitude of the *second largest* eigenvalue of A . The Perron-Frobenius Theorem tells us that this is strictly smaller than a . So $\rho(B) < a$.

Proof of Accumulation Reduction Theorem: We want to show that

$$\lim_{T \rightarrow \infty} E \left(\frac{1}{a^T} |W_T - Z_T| \right) = 0$$

We can rewrite the statement inside parentheses:

$$\begin{aligned}
\frac{1}{a^T} |W_T - Z_T| &= \frac{1}{a^T} \left| \vec{b}^\top \sum_{t=1}^T A^{T-t} \vec{X}_t - \sum_{t=1}^T a^{T-t} \psi_t \right| \\
&= \frac{1}{a^T} \left| \vec{b}^\top \sum_{t=1}^T A^{T-t} \vec{X}_t - \sum_{t=1}^T a^{T-t} \frac{1}{\vec{u} \cdot \vec{w}} \vec{w} \vec{u}^\top \vec{X}_t \right| \\
&= \frac{1}{a^T} \left| \vec{b}^\top \sum_{t=1}^T \left(A^{T-t} - \frac{a^{T-t}}{\vec{u} \cdot \vec{w}} \vec{w}\vec{u}^\top \right) \vec{X}_t \right|
\end{aligned}$$

Note that, if $B = A - \frac{a}{\vec{u} \cdot \vec{w}} \vec{w}\vec{u}^\top$, then $(A^j - \frac{a^j}{\vec{u} \cdot \vec{w}} \vec{w}\vec{u}^\top) = B^j$. So we can rewrite this as:

$$= \frac{1}{a^T} \left| \vec{b}^\top \sum_{t=1}^T B^{T-t} \vec{X}_t \right|$$

It will be helpful to reindex $j = T - t$.

$$= \frac{1}{a^T} \left| \vec{b}^\top \sum_{j=0}^{T-1} B^j \vec{X}_{T-j} \right|$$

We can now rephrase our goal in terms of B . We want to show that:

$$E \left(\frac{1}{a^T} \left| \vec{b}^\top \sum_{j=0}^{T-1} B^j \vec{X}_{T-j} \right| \right) \rightarrow 0$$

Now let $\{(\vec{u}_k, \lambda_k)\}$ be a complete set of left eigenvectors of A , chosen so that $\vec{u}_1 = \vec{u}$ and $\lambda_1 = a$. Since $\{\vec{u}_k\}$ spans \mathbb{R}^M , we can write $\vec{b}^\top = \sum_{k=1}^M b_k \vec{u}_k^\top$ as a linear combination of the left eigenvectors.

$$\begin{aligned} E \left(\frac{1}{a^T} \left| \vec{b}^\top \sum_{j=0}^{T-1} B^j \vec{X}_{T-j} \right| \right) &= E \left(\frac{1}{a^T} \left| \sum_{k=1}^M b_k \vec{u}_k^\top \sum_{j=0}^{T-1} B^j \vec{X}_{T-j} \right| \right) \\ &= \sum_{k=1}^M b_k E \left(\frac{1}{a^T} \left| \vec{u}_k^\top \sum_{j=0}^{T-1} B^j \vec{X}_{T-j} \right| \right) \end{aligned}$$

We will show each term of this sum goes to zero, by looking at each eigenvector \vec{u}_k . There are two cases.

The first case is when $\vec{u} = \vec{u}_1$. Recall that $\vec{u}^\top B = 0 \vec{u}^\top$. So this case is:

$$\begin{aligned} E \left(\frac{1}{a^T} \left| \vec{u}^\top \sum_{j=0}^{T-1} B^j \vec{X}_{T-j} \right| \right) &= E \left(\frac{1}{a^T} \left| \sum_{j=0}^{T-1} \vec{u}^\top B^j \vec{X}_{T-j} \right| \right) \\ &= E \left(\frac{1}{a^T} \left| \sum_{j=0}^{T-1} 0 \vec{u}^\top \vec{X}_{T-j} \right| \right) = 0 \end{aligned}$$

For $k > 1$, $\vec{u}_k^\top B = \lambda_k \vec{u}_k^\top$ where $|\lambda_k| < a$. So we get:

$$\begin{aligned} E \left(\frac{1}{a^T} \left| \vec{u}_k^\top \sum_{j=0}^{T-1} B^j \vec{X}_{T-j} \right| \right) &= E \left(\frac{1}{a^T} \left| \sum_{j=0}^{T-1} \vec{u}_k^\top B^j \vec{X}_{T-j} \right| \right) \\ &= E \left(\frac{1}{a^T} \left| \sum_{j=0}^{T-1} \lambda_k^j \vec{u}_k^\top \vec{X}_{T-j} \right| \right) \\ &\leq \frac{1}{a^T} \sum_{j=0}^{T-1} |\lambda_k|^j E(|\vec{u}_k^\top \vec{X}_{T-j}|) \end{aligned}$$

But all the \vec{X}_{T-j} are identical variables, so $E(|\vec{u}_k^\top \vec{X}_{T-j}|) = E(|\vec{u}_k^\top \vec{X}|)$, which we can pull out of the sum.

$$\begin{aligned} &= \frac{1}{a^T} E(|\vec{u}_k^\top \vec{X}|) \sum_{j=0}^{T-1} |\lambda_k|^j \\ &= \frac{1}{a^T} E(|\vec{u}_k^\top \vec{X}|) \left(\frac{|\lambda_k|^T - 1}{|\lambda_k| - 1} \right) \end{aligned}$$

$$= \left(\frac{\lambda_k}{a}\right)^T \left(\frac{1 - 1/\lambda_k^T}{\lambda - 1}\right) E(|\vec{u}_k^\top \vec{X}|)$$

Note that $\left(\frac{\lambda_k}{a}\right) < 1$, so $\left(\frac{\lambda_k}{a}\right)^T \rightarrow 0$ for large T . The other factors in the expression remain bounded, so each term of the expansion in b_k goes to zero. This concludes the proof.

Proof of the Stable Shape Theorem

To prove the Stable Shape Theorem, we also need a theorem attributed to Lévy.

Lévy's Continuity Theorem (Fristedt & Gray, 1997)

A sequence of random variables $\{W_T\}$ with characteristic functions $\varphi_{W_T}(s)$ converges in distribution to a random variable W_∞ if and only if $\varphi_{W_T}(s)$ converges pointwise to a function $\varphi(s)$ which is continuous at 0. If so, then $\varphi = \varphi_{W_\infty}$.

Proof of the Stable Shape Theorem: Lévy's Continuity Theorem tells us that if (a) the characteristic functions $\varphi_{Y_T}(s)$ converge pointwise to φ_{Y_∞} in a neighborhood of the origin, and (b) that the function φ_{Y_∞} is continuous at the origin, then Y_T converges to Y_∞ . It will help for us to work with the function $G_X(s) = \log(\varphi_X(s))$. Because the logarithm function is analytic in a neighborhood of 1, and $\varphi_X(s)$ is continuous and differentiable in a neighborhood of 0, we need only show that (a) $G_{Y_T}(s)$ converges pointwise to $G_{Y_\infty}(s)$ in a neighborhood of 0 and (b) $G_{Y_\infty}(s)$ is continuous at the origin. Statement (b) is true for free, since the characteristic function of any random variables is continuous in a neighborhood of the origin, and $\log()$ is analytic. So we need only show (a).

We can rewrite the characteristic function of $Y_T = \frac{a-1}{a^T-1} \sum_{t=1}^T a^{T-t} X_t$ using properties of characteristic functions.

$$\begin{aligned} \varphi_{Y_T}(s) &= \varphi_{\frac{a-1}{a^T-1} \sum_{t=1}^T a^{T-t} X_t}(s) = \prod_{t=1}^T \varphi_X\left(\frac{a-1}{a^T-1} a^{T-t} s\right) \\ &= \prod_{t=1}^T \varphi_X\left(\frac{a-1}{a^t - \frac{1}{a^{T-t}}} s\right) \end{aligned}$$

Then, taking the logarithm gives us:

$$G_{Y_T}(s) = \log \left[\prod_{t=1}^T \varphi_X\left(\frac{a-1}{a^t - \frac{1}{a^{T-t}}} s\right) \right]$$

$$\begin{aligned}
&= \sum_{t=1}^T \log \left[\varphi_X \left(\frac{a-1}{a^t - \frac{1}{a^{T-t}}} s \right) \right] \\
&= \sum_{t=1}^T G_X \left(\frac{a-1}{a^t - \frac{1}{a^{T-t}}} s \right)
\end{aligned}$$

A similar argument shows that

$$G_{Y_\infty}(s) = \sum_{t=1}^{\infty} G_X \left(\frac{a-1}{a^t} s \right)$$

We want to show that $G_{Y_T}(s)$ converges to $G_{Y_\infty}(s)$ pointwise for every s in a neighborhood of the origin. At a given s , we can bound the difference of the functions using the triangle inequality.

$$\begin{aligned}
\lim_{T \rightarrow \infty} |G_{Y_T}(s) - G_{Y_\infty}(s)| &= \lim_{T \rightarrow \infty} \left| \sum_{t=1}^T G_X \left(\frac{a-1}{a^t - \frac{1}{a^{T-t}}} s \right) - \sum_{t=1}^T G_X \left(\frac{a-1}{a^t} s \right) \right| \\
&\leq \lim_{T \rightarrow \infty} \sum_{t=1}^T \left| G_X \left(\frac{a-1}{a^t - \frac{1}{a^{T-t}}} s \right) - G_X \left(\frac{a-1}{a^t} s \right) \right|
\end{aligned}$$

Since X has a finite mean, we know φ_X and G_X are continuous and differentiable in a neighborhood $\mathcal{V} \subset \mathbb{R}$ properly containing the origin. Therefore, given a finite-length interval $U \in \mathcal{V}$, there is an $h > 0$ such that, for any two points $s_1, s_2 \in U$, the difference of G_X is bounded: $|G_X(s_1) - G_X(s_2)| \leq h|s_2 - s_1|$. We can use this to bound the summands in the expression above.

Choose $U = [-s, s]$ (or $[s, -s]$ if s is negative) and the h that corresponds to U . Since $a > 1$, the two points $\frac{a-1}{a^t - \frac{1}{a^{T-t}}} s$ and $\frac{a-1}{a^t} s$ fall within U for all $T > 1$ and $t \in [1, T]$. So

$$\left| G_X \left(\frac{a-1}{a^t - \frac{1}{a^{T-t}}} s \right) - G_X \left(\frac{a-1}{a^t} s \right) \right| \leq h \left| \left(\frac{a-1}{a^t - \frac{1}{a^{T-t}}} s \right) - \left(\frac{a-1}{a^t} s \right) \right|$$

Which means

$$\begin{aligned}
\lim_{T \rightarrow \infty} |G_{Y_T}(s) - G_{Y_\infty}(s)| &\leq \lim_{T \rightarrow \infty} \sum_{t=1}^T h \left| \left(\frac{a-1}{a^t - \frac{1}{a^{T-t}}} s \right) - \left(\frac{a-1}{a^t} s \right) \right| \\
&= \lim_{T \rightarrow \infty} \sum_{t=1}^T h|s|(a-1) \left| \left(\frac{1}{a^t - \frac{1}{a^{T-t}}} \right) - \left(\frac{1}{a^t} \right) \right| \\
&= h|s|(a-1) \lim_{T \rightarrow \infty} \sum_{t=1}^T \left| \frac{a^{T-t}}{a^T - 1} - \frac{1}{a^t} \right|
\end{aligned}$$

Finding a common denominator for the part inside the sum gives us:

$$\begin{aligned}
&= h|s|(a-1) \lim_{T \rightarrow \infty} \sum_{t=1}^T \left| \frac{a^t}{a^t(a^T-1)} - \frac{(a^T-1)}{a^t(a^T-1)} \right| \\
&= h|s|(a-1) \lim_{T \rightarrow \infty} \sum_{t=1}^T \frac{1}{a^t(a^T-1)} \\
&= h|s|(a-1) \lim_{T \rightarrow \infty} \left(\frac{1}{a^T-1} \sum_{t=1}^T \frac{1}{a^t} \right) \\
&= h|s|(a-1) \lim_{T \rightarrow \infty} \left(\frac{1}{a^T-1} \frac{1 - \left(\frac{1}{a}\right)^T}{a-1} \right) = 0
\end{aligned}$$

This actually shows a stronger result: That $G_{Y_T}(s)$ converges uniformly to $G_{Y_\infty}(s)$ on any finite-length interval in \mathcal{V} . This concludes the proof.

Proof of the Accumulated Moment Theorem

Lemma LM: If the n th moment $M_n = E(X^n)$ of a non-negative real-valued random variable X exists, then so does M_j for all $j < n$.

Proof: If f_X is the probability distribution function of X , and $j < n$ then

$$E(X^j) = \int_0^\infty x^j f_X(x) dx = \int_0^1 x^j f_X(x) dx + \int_1^\infty x^j f_X(x) dx$$

The left integral always exists, since it's bounded on a finite interval. Since $x^j \leq x^n$ on $(1, \infty)$,

$$\int_1^\infty x^j f_X(x) dx \leq \int_1^\infty x^n f_X(x) dx$$

We know the right integral exists by assumption. So the integral $\int_0^\infty x^j f_X(x) dx = E(X^j)$ exists.

Lemma MK: The n th moment of a random variable exists if and only if the n th cumulant exists. The relationship between the moment and cumulant are given by Faà di Bruno's formula, which can be found, for instance in Lukacs (1970), section 2.4.

Proof of Accumulated Moment Theorem: If we can show the equivalence of statements (1)-(8), the Lemma LM gives us the existence of lower-order cumulants and moments for free. To prove the equivalence of (1)-(8), first note that, by Lemma MK, the n th cumulant exists if and only if the n th moment also exists. This gives us the horizontal relationships in the Accumulated Moment Theorem. We only need show the vertical relationships. For each constructed distribution (Z_T, Y_T, Y_∞) , we will show that the existence of its n th cumulant is equivalent to the existence of the n th cumulant of X .

$\kappa_n(X)$ exists iff $\kappa_n(Z_T)$ exists: This follows directly from the properties of cumulants. We know that $Z_T = \sum_{t=1}^T a^{T-t} X_t$ is a finite sum of iid variables. So using the properties:

$$\begin{aligned}
\kappa_n(Z_T) &= \kappa_n\left(\sum_{t=1}^T a^{T-t} X_t\right) \\
&= \sum_{t=1}^T \kappa_n(a^{T-t} X_t) \\
&= \sum_{t=1}^T a^{n(T-t)} \kappa_n(X_t) \\
&= \kappa_n(X) \sum_{t=1}^T a^{n(T-t)} \\
&= \kappa_n(X) \sum_{j=0}^{T-1} (a^n)^j \\
&= \kappa_n(X) \frac{a^{nT} - 1}{a^n - 1}
\end{aligned}$$

This shows that $\kappa_n(X)$ exists iff $\kappa_n(Z_T)$ does, and gives an explicit formula.

$\kappa_n(X)$ exists iff $\kappa_n(Y_T)$ exists: The proof of this, and the derivation of the relationship between these quantities is similar to that for Z_T .

$\kappa_n(X)$ exists iff $\kappa_n(Y_\infty)$ exists: For this, we can use the relationship

$$Y_\infty \sim \frac{1}{a} Y_\infty + \frac{a-1}{a} X$$

Taking cumulants gives

$$\begin{aligned}
\kappa_n(Y_\infty) &= \kappa_n\left(\frac{1}{a} Y_\infty + \frac{a-1}{a} X\right) \\
&= \frac{1}{a^n} \kappa_n(Y_\infty) + \left(\frac{a-1}{a}\right)^n \kappa_n(X)
\end{aligned}$$

Solving for $\kappa_n(Y_\infty)$ algebraically then gives the relationship $\kappa_n(Y_\infty) = \frac{(a-1)^n}{a^n - 1} \kappa_n(X)$, which also shows the existence statement.

Proof for the Distribution Formulas of Accumulated Exponential

Proving the PDF theorem will be rather involved, and will require some lemmas.

Lemma Exp1: If r, s are distinct positive numbers, then

$$\int_0^x e^{-rz} e^{-s(x-z)} dz = \frac{1}{s-r} e^{-rx} + \frac{1}{r-s} e^{-sx}$$

Proof:

$$\begin{aligned}
\int_0^x e^{-rz} e^{-s(x-z)} dz &= e^{-sx} \int_0^x e^{-rz} e^{sz} dz \\
&= e^{-sx} \int_0^x e^{(s-r)z} dz \\
&= e^{-sx} \left[\left(\frac{1}{s-r} \right) e^{(s-r)z} \right]_0^x \\
&= e^{-sx} \left[\left(\frac{1}{s-r} \right) e^{(s-r)x} - \frac{1}{s-r} \right] \\
&= \frac{1}{s-r} e^{-rx} + \frac{1}{r-s} e^{-sx}
\end{aligned}$$

Lemma Exp2: Let r_1, r_2, \dots, r_n be a sequence of distinct positive numbers.

$$\sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} \cdot \frac{1}{(r_i - r_n)} = \frac{1}{\prod_{j=1}^{n-1} (r_j - r_n)}$$

Proof: By induction on n . The base case is $n = 2$. In this case, the left hand side becomes:

$$\sum_{i=1}^1 \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^1 (r_j - r_i)} \cdot \frac{1}{(r_i - r_n)} = \frac{1}{r_1 - r_2} = \frac{1}{\prod_{j=1}^1 (r_j - r_2)}$$

For the inductive step, assume the result is true for any set of n numbers. We will show it's true for any set of $n + 1$ numbers. Specifically, we want to show:

$$\sum_{i=1}^n \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} \cdot \frac{1}{(r_i - r_{n+1})} = \frac{1}{\prod_{j=1}^n (r_j - r_{n+1})}$$

Start by splitting the $i = n$ term out of the left hand side:

$$\begin{aligned}
&\sum_{i=1}^n \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} \cdot \frac{1}{(r_i - r_{n+1})} \\
&= \left(\sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} \cdot \frac{1}{(r_i - r_{n+1})} \right) + \frac{1}{\prod_{j=1}^{n-1} (r_j - r_n)} \cdot \frac{1}{(r_n - r_{n+1})}
\end{aligned}$$

Now we use the fact that the result is true for n , to replace part of the right term:

$$= \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} \cdot \frac{1}{(r_i - r_{n+1})} + \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} \cdot \frac{1}{(r_i - r_n)} \cdot \frac{1}{(r_n - r_{n+1})}$$

$$\begin{aligned}
&= \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} \cdot \frac{1}{(r_i - r_{n+1})} + \sum_{i=1}^{n-1} \frac{-1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} \cdot \frac{1}{(r_n - r_{n+1})} \\
&= \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} \cdot \left[\frac{1}{(r_i - r_{n+1})} - \frac{1}{(r_n - r_{n+1})} \right] \\
&= \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} \cdot \left[\frac{(r_n - r_{n+1}) - (r_i - r_{n+1})}{(r_i - r_{n+1})(r_n - r_{n+1})} \right] \\
&= \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} \cdot \left[\frac{r_n - r_i}{(r_i - r_{n+1})(r_n - r_{n+1})} \right]
\end{aligned}$$

Now notice that the $r_n - r_i$ on the top will cancel out the factor with $j = n$ in the product.

$$\begin{aligned}
&= \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} \cdot \frac{1}{(r_n - r_i)} \left[\frac{r_n - r_i}{(r_i - r_{n+1})(r_n - r_{n+1})} \right] \\
&= \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} \cdot \frac{1}{(r_i - r_{n+1})} \cdot \frac{1}{(r_n - r_{n+1})} \\
&= \frac{1}{(r_n - r_{n+1})} \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} \cdot \frac{1}{(r_i - r_{n+1})}
\end{aligned}$$

Where we pulled out the factor without any i or j in it. Now we can use induction again, since the sum is exactly the left side of the equation with $n - 1$. Except in this case, r_n is replaced with r_{n+1} . That's ok, since the result is true regardless of what we label the numbers. Replacing the sum, we get:

$$\begin{aligned}
&= \frac{1}{(r_n - r_{n+1})} \frac{1}{\prod_{j=1}^{n-1} (r_j - r_{n+1})} \\
&= \frac{1}{\prod_{j=1}^n (r_j - r_{n+1})}
\end{aligned}$$

This proves the inductive step, and the lemma.

Lemma Exp3: Let r_1, r_2, \dots, r_n be a sequence of distinct positive numbers, and set

$$f_i(x) = \begin{cases} r_i e^{-r_i x} & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \quad \text{Then}$$

$$f_1 * f_2 * \dots * f_n(x) = r_1 r_2 \dots r_n \left[\sum_{i=1}^n \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} e^{-r_i x} \right]$$

Proof: By induction on n . If $n = 1$, then this is trivially true.

Now assume the result is true for $n - 1$, so that:

$$f_1 * f_2 * \cdots * f_{n-1}(x) = r_1 r_2 \cdots r_{n-1} \left[\sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} e^{-r_i x} \right]$$

Then the convolution with n functions becomes:

$$(f_1 * f_2 * \cdots * f_{n-1}) * f_n = \int_0^x r_1 r_2 \cdots r_{n-1} \left[\sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} e^{-r_i z} \right] r_n e^{-r_n(x-z)} dz$$

Note that the upper limit is x . This is because $f_i(y)$ is zero for all negative values of y . So $f_n(x - z) = 0$ for all $x - z < 0$, or equivalently for all $z > x$. Pulling out the constants, we get:

$$= r_1 r_2 \cdots r_{n-1} r_n \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} \int_0^x e^{-r_i z} e^{-r_n(x-z)} dz$$

We can rewrite the integral using Lemma Exp1:

$$= r_1 r_2 \cdots r_{n-1} r_n \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} \left[\frac{1}{r_n - r_i} e^{-r_i x} + \frac{1}{r_i - r_n} e^{-r_n x} \right]$$

We can distribute across the brackets to get:

$$= r_1 r_2 \cdots r_{n-1} r_n \sum_{i=1}^{n-1} \left[\frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} \frac{1}{r_n - r_i} e^{-r_i x} + \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} \frac{1}{r_i - r_n} e^{-r_n x} \right]$$

Splitting up the sum gives us:

$$= r_1 r_2 \cdots r_{n-1} r_n \left[\sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} \frac{1}{r_n - r_i} e^{-r_i x} + \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} \frac{1}{r_i - r_n} e^{-r_n x} \right]$$

The denominator in the first sum just becomes the product over all r_i with $i \neq j$ including r_n :

$$= r_1 r_2 \cdots r_{n-1} r_n \left[\sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} e^{-r_i x} + \sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^{n-1} (r_j - r_i)} \frac{1}{r_i - r_n} e^{-r_n x} \right]$$

We can rewrite the right sum using Lemma Exp2:

$$= r_1 r_2 \cdots r_{n-1} r_n \left[\sum_{i=1}^{n-1} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} e^{-r_i x} + \frac{1}{\prod_{j=1}^{n-1} (r_j - r_n)} e^{-r_n x} \right]$$

The rightmost term is just the $i = n$ term in the sum. So we can combine and get:

$$= r_1 r_2 \cdots r_{n-1} r_n \sum_{i=1}^n \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (r_j - r_i)} e^{-r_i x}$$

Which is what we want.

Lemma Exp4: Let $b < 1$ and $c \in \mathbb{R}$, and set

$$f_i(x) = \begin{cases} cb^i e^{-cb^i x} & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \quad \text{Then}$$

$$f_1 * f_2 * \cdots * f_n(x) = c \sum_{i=1}^n \left[\frac{(-1)^{n-i}}{(b; b)_{i-1} (b; b)_{n-i}} b^{\frac{1}{2}[(n-i)^2 + (n+i)]} \right] e^{-cb^i x}$$

Proof: Setting $r_i = cb^i$ allows us to use Lemma Exp3.

$$f_1 * f_2 * \cdots * f_n(x) = cb^1 cb^2 \cdots cb^n \left[\sum_{i=1}^n \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (cb^j - cb^i)} e^{-cb^i x} \right]$$

Since the product inside the sum has $n - 1$ factors, there are $n - 1$ c 's in each term. We can cancel these with all but one of the c 's outside the sum.

$$\begin{aligned} &= c^n b^1 b^2 \cdots b^n \left[\sum_{i=1}^n \frac{1}{c^{n-1}} \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (b^j - b^i)} e^{-cb^i x} \right] \\ &= cb^1 b^2 \cdots b^n \left[\sum_{i=1}^n \frac{1}{\prod_{\substack{j=1 \\ j \neq i}}^n (b^j - b^i)} e^{-cb^i x} \right] \end{aligned}$$

Rewriting this:

$$= c \sum_{i=1}^n \left[\frac{b^1 b^2 \cdots b^n}{\prod_{\substack{j=1 \\ j \neq i}}^n (b^j - b^i)} \right] e^{-cb^i x}$$

Let us focus on the part inside the brackets, and let's temporarily invert it for convenience.

$$\frac{1}{b^1 b^2 \cdots b^n} \prod_{\substack{j=1 \\ j \neq i}}^n (b^j - b^i) = \frac{1}{b^1 b^2 \cdots b^n} \left[\prod_{j=1}^{i-1} (b^j - b^i) \right] \left[\prod_{j=i+1}^n (b^j - b^i) \right]$$

Recall that $b < 1$. In the left product j is always less than i , while the opposite is true in the right product. We'll rewrite these so they look like q-Pochhammer symbols.

$$= \frac{1}{b^1 b^2 \cdots b^n} \left[\prod_{j=1}^{i-1} b^j (1 - b^{i-j}) \right] \left[\prod_{j=i+1}^n b^i (b^{j-i} - 1) \right]$$

$$\begin{aligned}
&= \frac{1}{b^1 b^2 \dots b^n} [b^1(1 - b^{i-1})b^2(1 - b^{i-2}) \dots b^{i-1}(1 - b^1)] [b^i(b-1)b^i(b^2-1) \dots b^i(b^{n-i} \\
&\quad - 1)] \\
&= \frac{1}{b^1 b^2 \dots b^n} \left[b^1 b^2 b^3 \dots b^{i-1} \prod_{k=1}^{i-1} (1 - b^k) \right] \left[(b^i)^{n-i} (-1)^{n-i} \prod_{k=1}^{n-i} (1 - b^k) \right] \\
&= \frac{b^1 b^2 b^3 \dots b^{i-1} \cdot (b^i)^{n-i}}{b^1 b^2 \dots b^n} (-1)^{n-i} (b; b)_{i-1} (b; b)_{n-i}
\end{aligned}$$

Focusing on that mass of b 's

$$\begin{aligned}
\frac{b^1 b^2 b^3 \dots b^{i-1} \cdot (b^i)^{n-i}}{b^1 b^2 \dots b^n} &= \left(\frac{b^1}{b^1} \cdot \frac{b^2}{b^2} \cdot \frac{b^3}{b^3} \dots \frac{b^{i-1}}{b^{i-1}} \right) \cdot \left(\frac{b^i}{b^i} \cdot \frac{b^i}{b^{i+1}} \cdot \frac{b^i}{b^{i+2}} \dots \frac{b^i}{b^{n-1}} \right) \cdot \frac{1}{b^n} \\
&= 1 \cdot \frac{1}{b} \cdot \frac{1}{b^2} \dots \frac{1}{b^{n-i-1}} \cdot \frac{1}{b^n} \\
&= \frac{1}{b^{\frac{(n-i)(n-i-1)}{2} + n}} \\
&= \frac{1}{b^{\frac{1}{2}[(n-i)^2 + (n+i)]}}
\end{aligned}$$

Putting this all together gives the inverse of each coefficient as:

$$\frac{1}{b^1 b^2 \dots b^n} \prod_{\substack{j=1 \\ j \neq i}}^n (b^j - b^i) = (-1)^{n-i} b^{-\frac{1}{2}[(n-i)^2 + (n+i)]} (b; b)_{i-1} (b; b)_{n-i}$$

Which means the coefficient is:

$$\frac{(-1)^{n-i} b^{\frac{1}{2}[(n-i)^2 + (n+i)]}}{(b; b)_{i-1} (b; b)_{n-i}}$$

And

$$f_1 * f_2 * \dots * f_n(x) = c \sum_{i=1}^n \left[\frac{(-1)^{n-i}}{(b; b)_{i-1} (b; b)_{n-i}} b^{\frac{1}{2}[(n-i)^2 + (n+i)]} \right] e^{-cb^i x}$$

Lemma Exp5: Let $b > 1$ and $c \in \mathbb{R}$, and set

$$f_i(x) = \begin{cases} cb^i e^{-cb^i x} & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \quad \text{Then}$$

$$f_1 * f_2 * \dots * f_n(x) = c \sum_{i=1}^n \left[\frac{(-1)^{i-1}}{\left(\frac{1}{b}; \frac{1}{b}\right)_{i-1} \left(\frac{1}{b}; \frac{1}{b}\right)_{n-i}} b^{-\frac{1}{2}i^2 + \frac{3}{2}i} \right] e^{-cb^i x}$$

Proof: The difference between Lemmas Exp4 and Exp5 is that $b > 1$ here, so we'll need to work with $\left(1 - \left(\frac{1}{b}\right)^k\right)$. From the proof of Lemma Exp4

$$f_1 * f_2 * \dots * f_n(x) = c \sum_{i=1}^n \left[\frac{b^1 b^2 \dots b^n}{\prod_{\substack{j=1 \\ j \neq i}}^n (b^j - b^i)} \right] e^{-cb^i x}$$

Where the inverse of the coefficient is again

$$\frac{1}{b^1 b^2 \dots b^n} \prod_{\substack{j=1 \\ j \neq i}}^n (b^j - b^i) = \frac{1}{b^1 b^2 \dots b^n} \left[\prod_{j=1}^{i-1} (b^j - b^i) \right] \left[\prod_{j=i+1}^n (b^j - b^i) \right]$$

Rewriting this in terms of $\left(1 - \frac{1}{b^k}\right)$:

$$\begin{aligned} &= \frac{1}{b^1 b^2 \dots b^n} \left[\prod_{j=1}^{i-1} b^i (b^{j-i} - 1) \right] \left[\prod_{j=i+1}^n b^j (1 - b^{i-j}) \right] \\ &= \frac{1}{b^1 b^2 \dots b^n} \left[\prod_{j=1}^{i-1} b^i \left(\frac{1}{b^{i-j}} - 1 \right) \right] \left[\prod_{j=i+1}^n b^j \left(1 - \frac{1}{b^{j-i}} \right) \right] \\ &= \frac{1}{b^1 b^2 \dots b^n} \left[(b^i)^{i-1} (-1)^{i-1} \left(1 - \frac{1}{b^{i-1}} \right) \left(1 - \frac{1}{b^{i-2}} \right) \dots \left(1 - \frac{1}{b^1} \right) \right] \\ &\quad \left[b^{i+1} b^{i+2} \dots b^n \left(1 - \frac{1}{b^1} \right) \left(1 - \frac{1}{b^2} \right) \dots \left(1 - \frac{1}{b^{n-i}} \right) \right] \\ &= \frac{(b^i)^{i-1} b^{i+1} b^{i+2} \dots b^n}{b^1 b^2 \dots b^n} (-1)^{i-1} \left[\prod_{k=1}^{i-1} \left(1 - \left(\frac{1}{b} \right)^k \right) \right] \left[\prod_{k=1}^{n-i} \left(1 - \left(\frac{1}{b} \right)^k \right) \right] \\ &= \frac{(b^i)^{i-1} b^{i+1} b^{i+2} \dots b^n}{b^1 b^2 \dots b^n} (-1)^{i-1} \left(\frac{1}{b}; \frac{1}{b} \right)_{i-1} \left(\frac{1}{b}; \frac{1}{b} \right)_{n-i} \end{aligned}$$

Simplifying the portion in the front gives

$$\begin{aligned} &= \frac{(b^i)^{i-1}}{b^1 b^2 \dots b^i} (-1)^{i-1} \left(\frac{1}{b}; \frac{1}{b} \right)_{i-1} \left(\frac{1}{b}; \frac{1}{b} \right)_{n-i} \\ &= \frac{1}{b^i} \cdot (b^1 \cdot b^2 \dots b^{i-1}) (-1)^{i-1} \left(\frac{1}{b}; \frac{1}{b} \right)_{i-1} \left(\frac{1}{b}; \frac{1}{b} \right)_{n-i} \\ &= \frac{1}{b^i} \cdot b^{1+2+\dots+(i-1)} (-1)^{i-1} \left(\frac{1}{b}; \frac{1}{b} \right)_{i-1} \left(\frac{1}{b}; \frac{1}{b} \right)_{n-i} \\ &= b^{-i} \cdot b^{\frac{1}{2}i^2 - \frac{1}{2}i} (-1)^{i-1} \left(\frac{1}{b}; \frac{1}{b} \right)_{i-1} \left(\frac{1}{b}; \frac{1}{b} \right)_{n-i} \\ &= b^{\frac{1}{2}i^2 - \frac{3}{2}i} (-1)^{i-1} \left(\frac{1}{b}; \frac{1}{b} \right)_{i-1} \left(\frac{1}{b}; \frac{1}{b} \right)_{n-i} \end{aligned}$$

Plugging this inverse coefficient back into the formula gives us the conclusion:

$$f_1 * f_2 * \cdots * f_n(x) = c \sum_{i=1}^n \left[\frac{(-1)^{i-1}}{\left(\frac{1}{b}; \frac{1}{b}\right)_{i-1} \left(\frac{1}{b}; \frac{1}{b}\right)_{n-i}} b^{-\frac{1}{2}i^2 + \frac{3}{2}i} \right] e^{-cb^i x}$$

Proof of PDF Theorem: We'll start by calculating the probability distribution function of Z_T .

Now $Z_T = \sum_{t=1}^T a^{T-t} X_t$. This means that

$$\begin{aligned} f_{Z_T}(x) &= f_{X_T + aX_{T-1} + \cdots + a^{T-1}X_1}(x) \\ &= (f_{X_T} * f_{aX_{T-1}} * f_{a^2X_{T-2}} \cdots * f_{a^{T-1}X_1})(x) \end{aligned}$$

And since X is exponential with parameter s ,

$$f_X(x) = \begin{cases} se^{-sx} & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases}$$

And $f_{bX}(x) = \frac{1}{b} f_X\left(\frac{1}{b}x\right)$, so

$$f_{a^i X}(x) = \begin{cases} \frac{s}{a^i} e^{-sx/a^i} & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases}$$

So the expression $(f_{X_T} * f_{aX_{T-1}} * f_{a^2X_{T-2}} \cdots * f_{a^{T-1}X_1})(x)$ is exactly like in Lemma Exp4 where $c = sa$ and $b = \frac{1}{a}$. So that $r_i = cb^i = \frac{sa}{a^i} = \frac{s}{a^{i-1}}$. Lemma Exp3 says that:

$$(f_{X_T} * f_{aX_{T-1}} * f_{a^2X_{T-2}} \cdots * f_{a^{T-1}X_1})(x) = \sum_{i=1}^T \left[\frac{r_1 r_2 \cdots r_T}{\prod_{\substack{j=1 \\ j \neq i}}^T (r_j - r_i)} \right] e^{-r_i x}$$

Let's focus on the inverse of the coefficient in brackets:

$$\frac{1}{r_1 r_2 \cdots r_T} \prod_{\substack{j=1 \\ j \neq i}}^T (r_j - r_i) = \frac{1 \cdot a \cdot a^2 \cdots a^{T-1}}{s^T} \prod_{\substack{j=1 \\ j \neq i}}^T \left(\frac{s}{a^{j-1}} - \frac{s}{a^{i-1}} \right)$$

There are $T - 1$ factors of s in the numerator, since the $i = j$ term is omitted. This leaves only one $1/s$.

$$= \frac{1 \cdot a \cdot a^2 \cdots a^{T-1}}{s} \prod_{\substack{j=1 \\ j \neq i}}^T \left(\frac{1}{a^{j-1}} - \frac{1}{a^{i-1}} \right)$$

We can clear the fractions by multiplying by the highest power of a in each factor. For $j < i$, this means multiplying through by a^{i-1} . For $j > i$, it means multiplying through by a^{j-1} .

$$\begin{aligned} &= \frac{1 \cdot a \cdot a^2 \cdots a^{T-1}}{s} \left[\prod_{j=1}^{i-1} \left(\frac{1}{a^{j-1}} - \frac{1}{a^{i-1}} \right) \right] \left[\prod_{j=i+1}^T \left(\frac{1}{a^{j-1}} - \frac{1}{a^{i-1}} \right) \right] \\ &= \frac{1 \cdot a \cdot a^2 \cdots a^{T-1}}{s} \left[\prod_{j=1}^{i-1} \frac{1}{a^{i-1}} (a^{i-j} - 1) \right] \left[\prod_{j=i+1}^T \frac{1}{a^{j-1}} (1 - a^{j-i}) \right] \end{aligned}$$

$$= \frac{1 \cdot a \cdot a^2 \dots a^{T-1}}{s} \frac{1}{a^{(i-1)(i-1)}} \frac{1}{a^i a^{i+1} \dots a^{T-1}} \left[\prod_{j=1}^{i-1} (a^{i-j} - 1) \right] \left[\prod_{j=i+1}^T (1 - a^{j-i}) \right]$$

There is an $a^i a^{i+1} \dots a^{T-1}$ in both the numerator and denominator. Canceling these gives

$$= \frac{1}{s} \left(\frac{a^1 \cdot a^2 \dots a^{i-1}}{a^{(i-1)(i-1)}} \right) \left[\prod_{j=1}^{i-1} (a^{i-j} - 1) \right] \left[\prod_{j=i+1}^T (1 - a^{j-i}) \right]$$

The exponent of the numerator is $a^{\frac{(i-1)i}{2}}$. Simplifying then gives:

$$= \frac{1}{sa^{\frac{(i-1)(i-2)}{2}}} \left[\prod_{j=1}^{i-1} (a^{i-j} - 1) \right] \left[\prod_{j=i+1}^T (1 - a^{j-i}) \right]$$

The products can be reindexed.

$$= \frac{1}{sa^{\frac{(i-1)(i-2)}{2}}} \left[\prod_{j=1}^{i-1} (a^j - 1) \right] \left[\prod_{j=1}^{T-i} (1 - a^j) \right]$$

Inverting this gives:

$$sa^{\frac{(i-1)(i-2)}{2}} \left[\prod_{j=1}^{i-1} \frac{1}{(a^j - 1)} \right] \left[\prod_{j=1}^{T-i} \frac{1}{(1 - a^j)} \right]$$

Plugging this back into the equation, we get:

$$\begin{aligned} (f_{X_T} * f_{aX_T} * f_{a^2X_T} \dots * f_{a^{T-1}X_T})(x) &= \sum_{i=1}^T sa^{\frac{(i-1)(i-2)}{2}} \left[\prod_{j=1}^{i-1} \frac{1}{(a^j - 1)} \right] \left[\prod_{j=1}^{T-i} \frac{1}{(1 - a^j)} \right] e^{-\frac{s}{a^{i-1}}x} \\ &= \sum_{i=1}^T \frac{s}{a^{i-1}} a^{\frac{(i-1)i}{2}} \left[\prod_{j=1}^{i-1} \frac{1}{(a^j - 1)} \right] \left[\prod_{j=1}^{T-i} \frac{1}{(1 - a^j)} \right] e^{-\frac{s}{a^{i-1}}x} \end{aligned}$$

We then reindex $k = i - 1$.

$$= \sum_{k=0}^{T-1} \frac{s}{a^k} \left(a^{\frac{k(k+1)}{2}} \left[\prod_{j=1}^k \frac{1}{(a^j - 1)} \right] \left[\prod_{j=1}^{T-k-1} \frac{1}{(1 - a^j)} \right] \right) e^{-\frac{s}{a^k}x}$$

The expression in parentheses is the first expression for $c_{k,T}$. To get the second form, we use Lemma Exp4.

$$\begin{aligned} (f_{X_T} * f_{aX_T} * f_{a^2X_T} \dots * f_{a^{T-1}X_T})(x) &= sa \sum_{i=1}^T \left[\frac{(-1)^{T-i}}{\left(\frac{1}{a}; \frac{1}{a}\right)_{i-1} \left(\frac{1}{a}; \frac{1}{a}\right)_{T-i}} \left(\frac{1}{a}\right)^{\frac{1}{2}[(T-i)^2 + (T+i)]} \right] e^{-\frac{s}{a^{i-1}}x} \\ &= s \sum_{i=1}^T \left[\frac{(-1)^{T-i}}{\left(\frac{1}{a}; \frac{1}{a}\right)_{i-1} \left(\frac{1}{a}; \frac{1}{a}\right)_{T-i}} a^{-\frac{1}{2}[(T-i)^2 + (T+i)] + 1} \right] e^{-\frac{s}{a^{i-1}}x} \end{aligned}$$

We can reindex: $k = i - 1$. The exponent of the a becomes

$$\begin{aligned}
& -\frac{1}{2}[(T-i)^2 + (T+i)] + 1 = -\frac{1}{2}[T^2 - 2iT + i^2 + T + i - 2] \\
& = -\frac{1}{2}[T^2 - 2(k+1)T + (k+1)^2 + T + (k+1) - 2] \\
& = -\frac{1}{2}[T^2 - 2kT - 2T + k^2 + 2k + 1 + T + k - 1] \\
& = -\frac{1}{2}[T^2 - 2kT - T + k^2 + 3k] \\
& = -\frac{1}{2}[(T-k)(T-k-1) + 2k] \\
& = \frac{-(T-k)(T-k-1) - 2k}{2} \\
& = \sum_{k=0}^{T-1} \left[\frac{s}{a^k} \frac{(-1)^{T-k-1}}{\left(\frac{1}{a}; \frac{1}{a}\right)_k \left(\frac{1}{a}; \frac{1}{a}\right)_{T-k-1}} a^{-\frac{(T-k)(T-k-1)}{2}} \right] e^{-\frac{s}{a^k}x}
\end{aligned}$$

To show the formula for Y_T , we need only note that $Y_T = \frac{a-1}{a^T-1} Z_T$. So if we set $\gamma_T = \frac{a^T-1}{a-1}$ then we get:

$$\begin{aligned}
f_{Y_T}(x) &= f_{\frac{1}{\gamma_T} Z_T}(x) = \gamma_T f_{Z_T}(\gamma_T x) \\
&= \frac{a^T - 1}{a - 1} \sum_{k=0}^{T-1} \frac{s}{a^k} c_{k,T} e^{-\frac{s(a^T-1)}{a^k(a-1)}x}
\end{aligned}$$

This gives the formula for Y_T . To get the formula for Y_∞ , we first find the pdf of the partial sum variable

$$\bar{Y}_T = (a-1) \sum_{t=1}^T \frac{1}{a^t} X_t$$

The sequence $\{\bar{Y}_T\}$ converges to Y_∞ pointwise. So $\lim_{T \rightarrow \infty} f_{\bar{Y}_T}(x) = f_{Y_\infty}(x)$ for all $x \in [0, \infty)$. We need only calculate $f_{\bar{Y}_T}(x)$ and examine its behavior in the limit.

$$f_{\bar{Y}_T}(x) = f_{\frac{a-1}{a} X}(x) * f_{\frac{a-1}{a^2} X}(x) * f_{\frac{a-1}{a^3} X}(x) * \cdots * f_{\frac{a-1}{a^T} X}(x)$$

Now we use Lemma Exp5 with $b = a$ and $c = \frac{s}{a-1}$ so that $r_i = \frac{sa^i}{a-1}$.

$$\begin{aligned}
f_i(x) &= f_{\frac{a-1}{a^i} X}(x) = \begin{cases} \frac{a^i s}{(a-1)} e^{-\frac{a^i}{(a-1)}sx} & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \\
f_1 * f_2 * \cdots * f_T(x) &= \frac{s}{a-1} \sum_{i=1}^T \left[\frac{(-1)^{i-1}}{\left(\frac{1}{a}; \frac{1}{a}\right)_{i-1} \left(\frac{1}{a}; \frac{1}{a}\right)_{T-i}} a^{-\frac{1}{2}i^2 + \frac{3}{2}i} \right] e^{-\frac{a^i}{(a-1)}sx}
\end{aligned}$$

Rescaling $k = i - 1$ gives us

$$\begin{aligned}
&= \frac{s}{a-1} \sum_{k=0}^{T-1} \left[\frac{(-1)^k}{\left(\frac{1}{a}; \frac{1}{a}\right)_k \left(\frac{1}{a}; \frac{1}{a}\right)_{T-k-1}} a^{\frac{-k(k-1)+2}{2}} \right] e^{-\left(\frac{a}{a-1}\right)a^k s x} \\
&= \sum_{k=0}^{T-1} \frac{sa}{a-1} a^k \left[\frac{(-1)^k}{\left(\frac{1}{a}; \frac{1}{a}\right)_k \left(\frac{1}{a}; \frac{1}{a}\right)_{T-k-1}} a^{\frac{-k(k+1)}{2}} \right] e^{-\left(\frac{a}{a-1}\right)a^k s x}
\end{aligned}$$

For a given $a > 1$, the constant $\left(\frac{1}{a}; \frac{1}{a}\right)_{T-k-1}$ is bounded as T gets larger, and converges to $\left(\frac{1}{a}; \frac{1}{a}\right)_{T-k-1}$. The higher order terms die off quickly, so this series converges as $T \rightarrow \infty$.

So $\lim_{T \rightarrow \infty} f_{\bar{y}_T}(x)$ converges to

$$= \frac{s}{a-1} \sum_{k=0}^{\infty} \left[\frac{(-1)^k}{\left(\frac{1}{a}; \frac{1}{a}\right)_k \left(\frac{1}{a}; \frac{1}{a}\right)_{\infty}} a^{\frac{-k(k-1)+2}{2}} \right] e^{-\left(\frac{a}{a-1}\right)a^k s x}$$

This concludes the proof.

Chapter 5

Systemic Inequality

5.1 Introduction

5.1.1 Understanding Systemic Inequality

Entrenched problems of inequality are a target for billions of dollars of intervention investment. Despite this, certain problems remain seemingly intractable. In the United States, pre-tax income of the bottom 60% has grown less than 1% per year on average since 1979, while the incomes of the top 1 percent have grown at 3% per year (Stone et al., 2020). The high school dropout rate has held steady at about 5% (McFarland et al., 2018). Inequality in physical and mental health have also been stubbornly persistent over time (Cook et al., 2017; Olfson et al., 2015; Zimmerman & Anderson, 2019). These types of problems all fall under the category of *systemic inequality*.

In the research literature, the term *systemic inequality* is often not clearly defined, or it is taken for granted that the reader understands its meaning. In addition, the terms “systemic”, “structural” and “systematic” are often used to refer to similar forces. Various authors have referred to systemic inequality in terms of norms, social structures, and formal institutions (Lashitew et al., 2023), or in terms of disadvantages to certain groups that manifest in multiple ways simultaneously (Cech & Waidzunas, 2021). Some authors focus on a variety of component systems of society that might influence individual opportunity, such as “political, legal, economic, health care, school, and criminal justice systems” (Braveman et al., 2022). However, these systems themselves interact, and are part of a larger societal system in which every factor is endogenous (El-Sayed & Galea, 2017; Page & Zelner, 2020).

I define **systemic inequality** as the unequal distribution of well-being that arises from *a large number of heterogeneous* factors which are often causally related, many of which are hard or practically impossible for humans to measure, understand, and/or influence on a large scale.

Systemic inequality typically arises over time from complex systems in which individuals make decisions and aim to improve their well-being, though often non-optimally.

This framework includes all definitions from previous research by considering each individual in a space which not only represents their personal abilities and resources, but also their capacity to navigate the norms, structures, and institutions that are necessary for success. While the definition focuses on individuals, this does not discount the role that institutions play. In this approach, institutional decisions designed to increase equity can be represented as improvements in the ability of previously disadvantaged individuals to navigate those institutions. For instance, Amis, Mair & Munir found five organizational practices, such as hiring and promotion, that play a role in reproduction of inequality (Amis et al., 2020). An intervention that encouraged more equitable promotion practices could be construed as an individual-level improvement in some individuals' ability to navigate the business.

One example of systemic inequality is income inequality, either within a population or between populations such as racial groups. Research has shown a large number of factors related to income which plausibly have a causal influence. Some of these factors include: food insecurity (Wight et al., 2014), education (Boshara et al., 2015), self-perception of socioeconomic status (Tan et al., 2020), marital status (McLanahan & Percheski, 2008), residence location (S. F. Reardon & Bischoff, 2011), parenthood (Jones & Tertilt, 2008), race (Semega et al., 2021), skin color (Keith & Herring, 1991), the government policies that one lives under (Tazhitdinova, 2022), wealth (Killewald et al., 2017), skills (Edin et al., 2022), and genetics (Bowles & Gintis, 2002). Teasing out the complex causal web connecting these factors, as many have tried and failed to work their way up the economic ladder.

5.1.2 Hidden Variables

An astute reader will note that my list of factors related to income includes relatively easy-to-measure variables. For good reason, researchers and the media will often focus their attention on available data. But this creates a streetlight problem (Freedman, 2010) where the seeming solution set to a problem consists only of the easily viewable causes. However, any adult can explain with personal experience how subtle traits or experiences have affected their own life path. These might include the presence of a helpful mentor (Campbell & Campbell, 1997), parents who promoted certain mindsets (Yeager & Walton, 2011), how straight our teeth are

(Hamermesh, 2011), how strongly we react to stress (Boyce, 2019), whether we had a mathematically-oriented sibling which consequently influenced our academic interests (Joensen & Nielsen, 2018). While each of the listed cases has been studied, their complete and collective effects on individuals are out of the reach of researchers. Likely, there are many more factors yet unstudied.

The causes of our circumstances may be subtle, idiosyncratic, and hard for us to explain even within our own life path. Trying to understand many of these factors at scale is nigh impossible. A given factor may only have a small effect on each individual, or a large effect but only influence a few individuals. Each of these unmeasured causes is likely to only have a small effect overall, often strongly influenced by random chance. But since there are so many of them, they collectively create a significant amount of variance unexplainable by standard data sets (Salganik et al., 2020). And individuals move through various states of well-being throughout their lives, for instance by moving in and out of poverty (Fritzell & Henz, 2021), which means the factors influencing well-being are also idiosyncratic throughout time.

5.1.3 Reinforcing Networks

In addition, we know that these causal factors are often causally related to each other, creating a multidimensional cumulative advantage process (DiPrete & Eirich, 2006). Having a reliable car might increase someone's income in multiple ways: Showing up to work on time increases job stability. Time saved in transit allows more time for training. Increased geographic mobility leads to social opportunities, which can create a professional network. In turn, job stability, free time, and a robust social network may reinforce each other.

These mutually reinforcing relationships might also be framed in terms of negative traits, in a network of cumulative disadvantage. For example, chronic joint pain can reduce the likelihood of exercise (Breivik et al., 2006), which can contribute to poorer health outcomes (Pinckard et al., 2019). In reality, each individual is subject to forces that both promote and detract from well-being and success.

There are dimensions of well-being which are shared among almost all people (UN General Assembly, 1948). However, different individuals define well-being or success differently. Skills and values which are useful in one environment may be maladaptive in others (Anderson, 2000; Uskul et al., 2019; Yosso, 2005). Researchers often get past this complication by being specific

in the quantities being measured and allowing the reader to decide whether those quantities are important. This chapter takes a different approach, assuming that individuals are members of a cohort with similar goals and constraints.

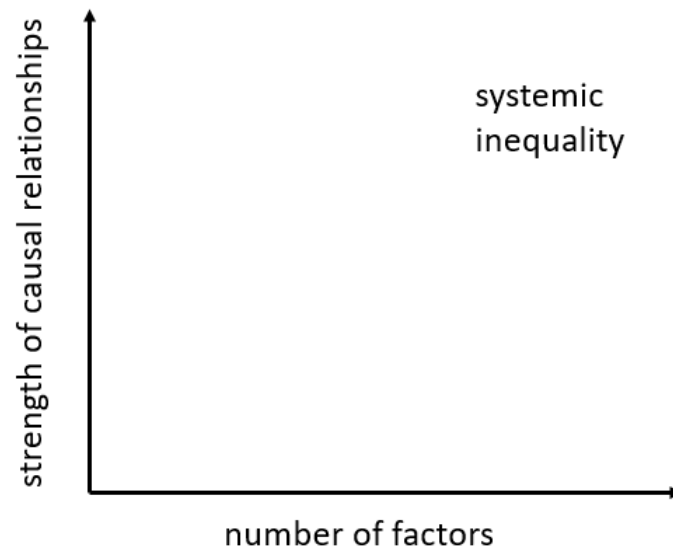


Figure 19: Graphical representation of systemic inequality

Systemic inequality arises in other domains such as obesity (Vandenbroeck et al., 2007), education (Quarles et al., 2018), GDP between countries (Hidalgo, 2015), physical health (Cech & Waidzunas, 2021), and mental health (Alegría et al., 2018). Unsurprisingly, income, education, and health are correlated, as people well-off in one factor use that advantage to help them become well-off in others.

While this chapter focuses primarily on social inequality, there is also evidence that similarly interrelated, multifactor processes happen on a biological level (Boyle et al., 2017). So it may be possible to generalize these results to non-societal contexts.

5.1.4 Progression of Inequality over Time

The most long-lasting forms of inequality do seem to be systemic, likely because the many, mutually reinforcing causes make decreasing gaps difficult. However, there are forms of inequality that are not systemic. Some situations, such as policy changes or natural disasters, create conditions of inequality that have a single cause. Indeed, researchers typically try to use

these natural experiments to isolate simple causal effects (Dunning, 2008). The separation of Germany into East and West after World War II was, in some, sense a natural experiment in the short term. Germany was suddenly divided into two pieces with very different political and economic systems. This created, at the time of separation, non-systemic inequality in many dimensions with a single cause. However, after 40 years of separate cultural and economic systems, those differences became systemic as people adapted to their new situation. At the time of reunification, East & West Germany had significant gaps in income, unemployment, life satisfaction (Petrunyk & Pfeifer, 2016), female workforce participation (Matysiak & Steinmetz, 2008), and culture (Krüger & Degel, 2022; Meier & Mutz, 2016). Over time, there has been convergence in these factors. However gaps between the East and West remain due to the embedded, systemic nature of those differences. If history is any indication, differences between East & West Germany will remain for a long time (Uskul et al., 2019). The multi-factor, mutually reinforcing nature of systems then (a) perpetuates and magnifies these average differences and (b) causes inequality to propagate to traits where they might not have been historically, such as life satisfaction. Similar systemization effects can be seen in other areas: Children in utero during the Dutch Hunger Winter in World War II suffered decreased labor outcomes 50 years later (Scholte et al., 2015). Two generations after the abolition of slavery in the US, economic indicators for descendants of slaves were the same as for descendants of free blacks (Sacerdote, 2005) creating enduring black-white success gaps that persist to this day. Unfortunately, experiments, natural or otherwise, cannot detangle the complex causality embedded in systemic inequality. Experimentation is designed to isolate a single factor to determine its causal effect. However, when causality is indirect, perhaps mediated by many variables, or when there are nonlinear interaction effects between many variables, experimentation misses the complexity of interactions.

5.1.5 Inequality Between Groups

Inequality can be considered within a population or between populations. These two cases are worth discussing briefly, because of the way inequality is measured differently in each of them. Within-population inequality is often measured in ways relative to the spread of the distribution, using measures such as the Gini index or the coefficient of variation. Figure 20A gives a demonstration of how we might visualize within population inequality by looking at how spread

out the probability distribution function is. In contrast, between-population inequality is often examined using measures of center, such as median or mean. In Figure 20B, we can see how there is between-group inequality between whites and blacks, since these groups' median earnings are \$40,000 and \$30,000, respectively. However, there is greater *within-group* inequality among whites than there is among blacks, because the distribution of white earnings is more spread out while black earnings are bunched up closer to zero. This graph also shows the significant amounts of heterogeneity that are hidden within any large population. Looking at median or mean differences is useful. However, using any single summary statistic to describe complex situations can draw the reader's mind toward overly simple explanations. Racial inequality and other group differences do not typically arise because one group is advantaged in a single dimension, but because many heterogeneous factors interact to create average differences. The many-factor approach I take here aims to partially address that heterogeneity.

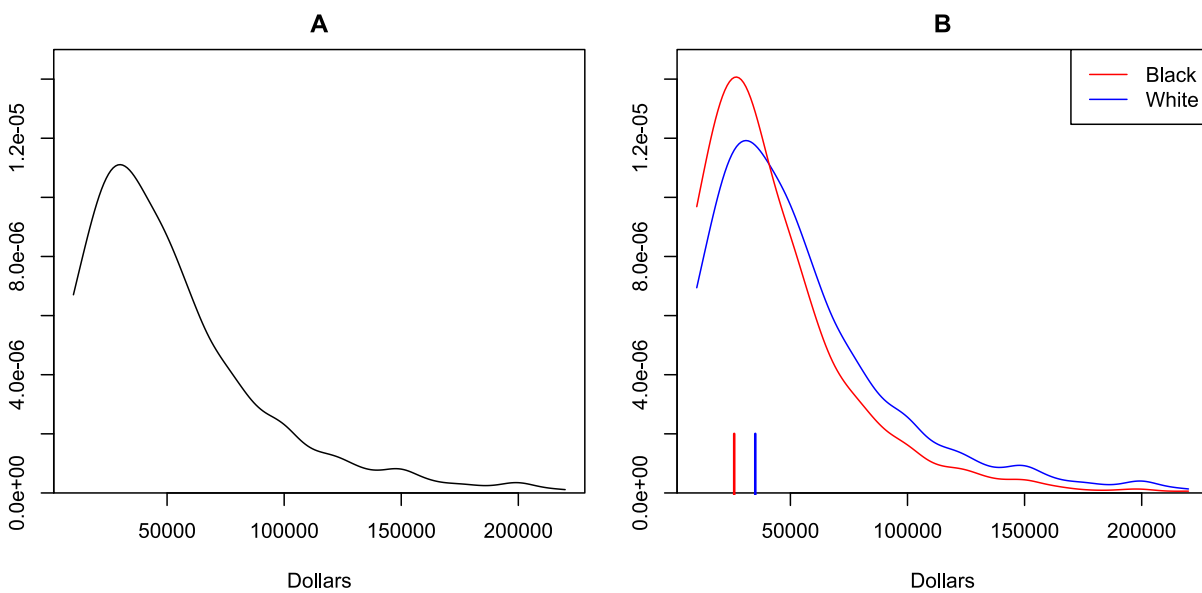


Figure 20: Distribution of Income (A) Earnings distribution for adult Americans. (B) Earnings distributions and medians for white and black Americans. Data taken from 2019 American Community Survey, and includes all people age 25-65 with positive earnings.

5.1.6 Models of Inequality

As this chapter uses two models to study inequality, it is worth reviewing some of the many models that have been used to study inequality. There is a long tradition dating back to Pareto

(1897) of using a small number of simple principles to model distributions of inequality, like those in Figure 20. Dagum (1977) used differential equations to generate his eponymous distribution. Others have followed his approach with, for instance, trickle-up (Henle et al., 2008), trickle-down (Sarabia et al., 2017), and conservation of money-based first principles (Drăgulescu & Yakovenko, 2000). These single-variable approaches achieve surprisingly good descriptions of static distributions, given that they do not account for the many factors contributing to economic success. The findings, and the universality of income distributions across countries (Tao et al., 2019), points to potential underlying universal principles governing income. One approach to modeling static inequality caused by multiple factors is to look for natural groups or latent variables in the data. Landale (2017) used latent class analysis to group whites and Hispanics in the LA area, and compare their perceptions of discrimination. Wilson & Urick (2022) took a similar approach to examining the opportunity gap in science. The economic complexity literature analyzes the economics of geographic regions by looking at exports, patent data (Hidalgo, 2021), or added value (Koch, 2021). The complexity of these metrics, as measured through a dimensional reduction technique, correlates with measures such as GDP and wage inequality (Sbardella et al., 2017).

Other researchers have built time-based models to predict longitudinal outcomes along many dimensions simultaneously. Sawhill & Reeves (2016) developed a simulation around a series of linear regressions, with one for each time point. Each regression took into account both circumstances of birth and previous life outcomes. Bloome (2015) built a Markov chain model, grouping individuals along family structure, income, and age, and then examined inequality dynamics as individuals progressed through their lives. Farrell, et al. (2018) used a scale-free network to examine health deficits in aging. In their model, each node represents a potential health deficit which gets repaired by homeostatic forces at a rate dependent on the state of neighbor nodes. Farrell's model is most similar to the one in Chapter Four and this chapter's Section 5.2. My model differs from Farrell's in that it examines temporally increasing behavior created by cumulative advantage rather than the balancing forces of homeostasis. In addition, my model delves more deeply into the conceptual and mathematical underpinnings related to growth and inequality.

The analysis in this chapter is split into two sections. Section 5.2 examines the growth of systemic inequality over time within a population by using the accumulation model from Chapter

Four. Section 5.3 examines the effects of interventions on systemic inequality using a utility function which maps multiple factors into a single dimension of well-being or success.

5.2 Systemic Inequality Arising Over Time

To examine how systemic inequality builds up over time, I use the accumulation model from Chapter Four. I assume every causal factor is a non-negative real-valued variable x_i which could represent an individual's income, social skills, ability to navigate the educational system, or many other things. When possible, each person uses their success in x_i to improve other factors. Between any two factors i and j , there is a “reinvestment rate” a_{ji} . This leads to a causal network which describes the relationships between factors. And at each time step, an individual's value of x_j gains an amount for each i equal to $a_{ji}x_i$. Additionally, at each time step, individuals receive a random amount added to each factor corresponding to schooling or gains from interacting with the world.

An individual therefore has, at time T , a random vector of factors given by:

$$\vec{Z}_T = \sum_{t=1}^T A^{T-t} \vec{X}_t$$

Here A is the sum of the identity matrix and the matrix defined by (a_{ij}) , and $\vec{X}_t \sim \vec{X}$ is the random vector of additive gains whose probabilities are assumed to be independent of time.

Chapter Four points out the strengths and limitations of this model. The model does capture general distributional and time-based trends found in nature and society. It does not represent the effects of exogenous variables such as genetics or parental investments. It also represents a unidirectional cumulative (dis)advantage process, in that factors grow over time, but do not shrink. However, the model can handle a variety of conceptual approaches. For example, an individual's “agency” could be represented as a combination of luck (random additive effects) and strategic experience & behavioral traits (which could be represented as a variable x_i).

The results of Chapter Four show a number of results which will be relevant to our discussion:

There is a systemic growth rate which gives the multiplicative growth of all factors in the long term. The per-unit-time growth rate is $a - 1$, where a is the largest eigenvalue of the matrix A . In some sense, the value of a is the systemic “reinvestment rate” which emerges from the individual relationships between individual factors (which are given individually by a_{ij}).

There is a vector \vec{w} which describes the long-term ratio of each factor to each other. In other words, for any individual and large T , the ratio of x_i/x_j is equal to w_i/w_j . The vector \vec{w} is the right eigenvector of A corresponding to the eigenvalue a .

There is a vector \vec{u} which describes the overall causal effect of each factor on growth. In other words, if u_i is twice as large as u_j , then increasing x_i will cause twice the growth of increasing x_j . This growth happens for large T and affects all other factors. The vector \vec{u} is the left eigenvector of A corresponding to the eigenvalue a .

More precisely, for large T , the following result holds:

$$\vec{Z}_T \sim \frac{1}{\vec{u} \cdot \vec{w}} \left[\sum_{t=1}^T a^{T-t} (\vec{u} \cdot \vec{X}_t) \right] \vec{w}$$

Or, if we are interested in the results of a regression $v = b_1x_1 + b_2x_2 + \dots + b_Mx_M$, the model creates a random variable for v at a given time:

$$V_T = \vec{b} \cdot \vec{Z}_T \sim \frac{1}{\vec{u} \cdot \vec{w}} \left[\sum_{t=1}^T a^{T-t} (\vec{u} \cdot \vec{X}_t) \right] (\vec{b} \cdot \vec{w})$$

Each of the individual variables x_i can also be represented as the result of such a regression using $\vec{b} = \langle 0, 0, \dots, 1, \dots, 0, 0 \rangle$.

5.2.1 Statistics of Systemic Inequality Change

The graphs in Figure 21 show what happens to summary statistics of V_T as the value of a and T vary.

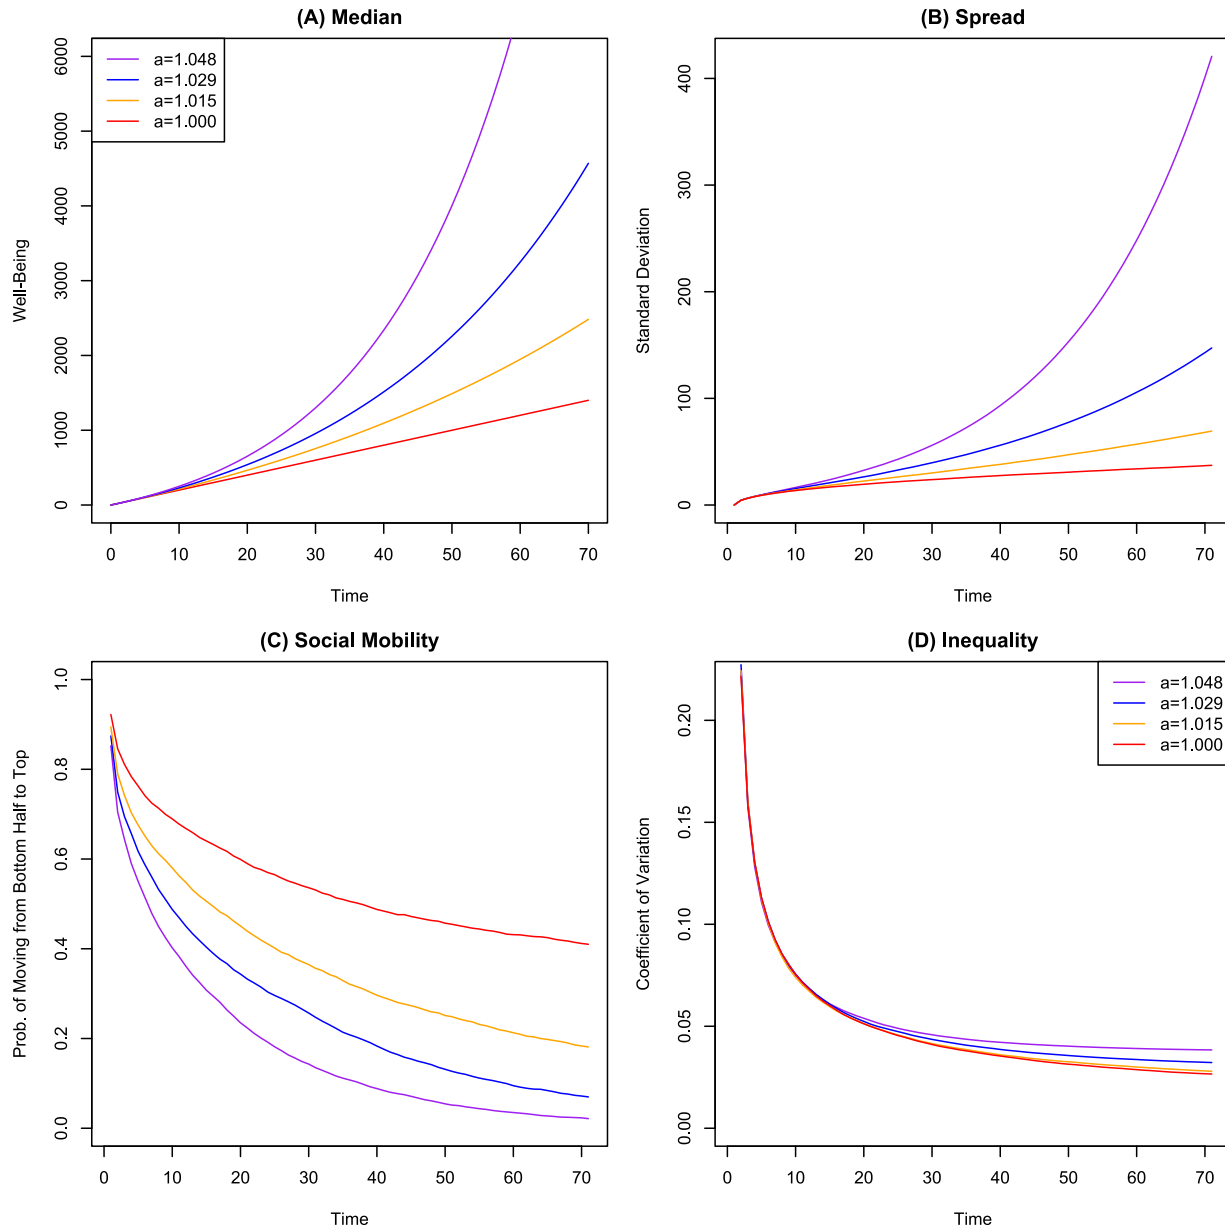


Figure 21: Summary statistics of accumulation models with different growth rates over time. (A) Median values as a varies. (B) Standard deviation of outcomes. (C) Relative social mobility, as measured by the probability that an individual below the median at time t would have a value above the median some time before $t + 50$. (D) Inequality, as measured by the coefficient of variation.

The median (Figure 21A) grows linearly at first, as the dominant source of growth comes from additive amounts \vec{X} such as schooling. However, much like an investment account, as well-being (or success) becomes larger, growth becomes exponential as it is dominated by reinvestment. This exponential growth happens more quickly for larger values of a . The red lines ($a = 1$)

correspond to a “control” case where there is no reinvestment. For instance, when $a = 1$, the median grows linearly since the only source of growth is a constant average additive amount at each time step. Graphs of means were similar to the graphs of the medians.

The gap between the advantaged and the disadvantaged, measured by the standard deviation, also grows faster for larger growth rates (Figure 21B). However, we can see that the spread grows less slowly than the mean, as inequality (Figure 21D) decreases over time. For reference, the red line corresponding to $a = 1$, goes like $cv \sim \frac{1}{\sqrt{T}}$ which can be explained by the Central Limit Theorem.

Relative social mobility also decreases over time (Figure 21C). The explanation for this is related to the old adage, “It’s harder to climb to the top of the socioeconomic ladder when the rungs are farther apart.” As gaps increase, additive effects like schooling are less effective, and growth in well-being comes more from reinvestments of pre-existing resources. People with fewer resources have less to reinvest, which makes it less likely they’ll move up the ladder. We can also see that, with a larger rate of return on those reinvestments (i.e. as the growth multiplier a increases), social mobility decreases.

When considering large-scale societal trends, these results should be interpreted with caution. The model predicts decreasing inequality over time, which does not represent inequality in the US (Stone et al., 2020). The assumed homogeneity of goals and opportunity in the simulated population does not take into account market segmentation or the effects of policies. However, we can see behavior in the model which has been observed elsewhere. The spread is significantly smaller than the observed growth, consistent with the economic observation that growth drives absolute mobility (Hout, 2015). If we think of different geographic regions having different systemic growth multipliers, a , then the medians in Figure 21A predict significant between-region inequality. This is consistent with evidence (Lakner & Milanovic, 2013; Manduca, 2019).

Methodological details for Figure 21: Each colored line represents the same simulation. For any pair of values $i, j \in \{1, 2, \dots, M\}$, the effect of factor j on factor i is generated using a combination of the outgoing/causal importance of factor j and the incoming/indicator importance of factor i . Specifically, the ij element of A is a_{ij} plus 1 if $i = j$, where:

$$a_{ij} = b_{ij} + c_{ij}$$

$$b_{ij} \in \text{Exp}\left(\frac{1}{\bar{b}_{*j}}\right)$$

$$c_{ij} \in \text{Exp}\left(\frac{1}{\bar{c}_{i*}}\right)$$

$$\bar{b}_{*j} \in \text{Exp}\left(\frac{M}{\eta}\right)$$

$$\bar{c}_{i*} \in \text{Exp}\left(\frac{M}{\eta}\right)$$

I call a matrix generated this way an exponential-exponential matrix (or exp-exp matrix, for short). I will denote the space of $M \times M$ random exp-exp matrices with parameter η as $\text{ExpExp}(\eta, M)$. The random amount added at each time step is $X_i \in \text{Exp}(1)$. Due to the randomness of the matrices, the eigenvalue a varies for the same starting parameter η values but different random seeds. The simulations above use matrices drawn using $M = 20$ and $\eta = 0, 0.01, 0.02, 0.03$, respectively. The outcome variable used to calculate summary statistics was the sum of all factors $\sum_i x_i$, which corresponds to the regression where $\vec{b} = \langle 1, 1, 1, \dots, 1 \rangle$. Other values of \vec{b} give similar results.

5.2.2 *The Causes of Success Need Not be the Indicators of It*

Recall the three-factor toy model from Chapter Four shown in Figure 22. Each individual in the model has three factors: social skills s , professional network strength p , and knowledge k . As time progresses, individuals are able to reinvest their success in one factor into the others. This leads to time progression given by the matrix equation:

$$\begin{pmatrix} s_{t+1} \\ p_{t+1} \\ k_{t+1} \end{pmatrix} = \begin{pmatrix} 1.01 & .01 & .005 \\ .03 & 1.005 & .015 \\ .002 & .012 & 1.005 \end{pmatrix} \begin{pmatrix} s_t \\ p_t \\ k_t \end{pmatrix}$$




Figure 22: Causal relationship between social skills, professional network, and knowledge in a toy model.

The results from Chapter Four showed that, in the long term, the values s, p, k go like:

$$\begin{pmatrix} s_t \\ p_t \\ k_t \end{pmatrix} \sim C \begin{pmatrix} 2.9 \\ 4.7 \\ 2.4 \end{pmatrix} (1.031)^t$$

Where C is some constant and $\vec{w} = \langle 2.9, 4.7, 2.4 \rangle$ is the right eigenvector of the causal matrix. The values of \vec{w} are the relative size of each factor in the long term. Since all of these traits can be construed as forms of success, successful people will have very large professional networks (p_t). For large t , the ratio of professional network to social skills will be $4.7/2.9 = 1.6$ in the units used to measure these quantities. So one's professional network will be a strong indicator of their success.

However, professional networks are not the main *driver* of success in this example. The long-term relative causal strength of each factor is given by the *left* eigenvector of the causal matrix, $\vec{u} = \langle 4.5, 2.9, 2.6 \rangle$. In this model, the most important factor in creating success is social skills. In fact, social skills are $\frac{4.5}{2.9} - 1 = 55\%$ more important than professional skills.

Consider the example of three people, Ayesha, Brian, and Chloe, who are given a set of skills in their childhood, but no further contributions as adults. Ayesha is given a professional network but no skills to use them, so her initial state is $\langle s_0, p_0, k_0 \rangle = \langle 0, 3, 0 \rangle$. Brian is given a balance of skills $\langle 1, 1, 1 \rangle$. Chloe has only social skills $\langle 3, 0, 0 \rangle$. The now-adults use the skills they've learned according to the process described above for 50 time steps. At $t = 50$, their social skills, professional network, and knowledge are:

Ayesha: $\langle s_{50}, p_{50}, k_{50} \rangle = \langle 2.9, 6.6, 2.9 \rangle$

Brian: $\langle s_{50}, p_{50}, k_{50} \rangle = \langle 4.0, 6.3, 3.5 \rangle$

Chloe: $\langle s_{50}, p_{50}, k_{50} \rangle = \langle 7.0, 8.0, 2.5 \rangle$

Because social skills are more influential than the other traits, Chloe now has a stronger professional network than either Ayesha or Brian. This makes sense intuitively, since social skills might be more important in generating a professional network than vice versa.

Of course, the numbers in this model are made up. However, evidence from research supports the idea that the indicators of success are not always the causes. For instance, noncognitive traits in childhood, such as self-esteem and having an internal locus of control, are shown to influence later life outcomes such as wages and criminality (Heckman et al., 2006). This is true despite the fact that adult success is not typically measured in self-esteem.

The benefit of this model is that it shows how the causal strength of a single factor accumulates primarily through indirect effects filtered through the network. This is called eigenvector centrality in some contexts (Newman, 2018). In a weighted, directed network like this, the relative size of the components of left eigenvector (causal strength) need not be related to those of the right eigenvector (outcome size). Figure 23 shows the distribution of correlations between causal strength and relative outcome size for 10,000 randomly generated matrices, where the mean correlation is effectively zero.

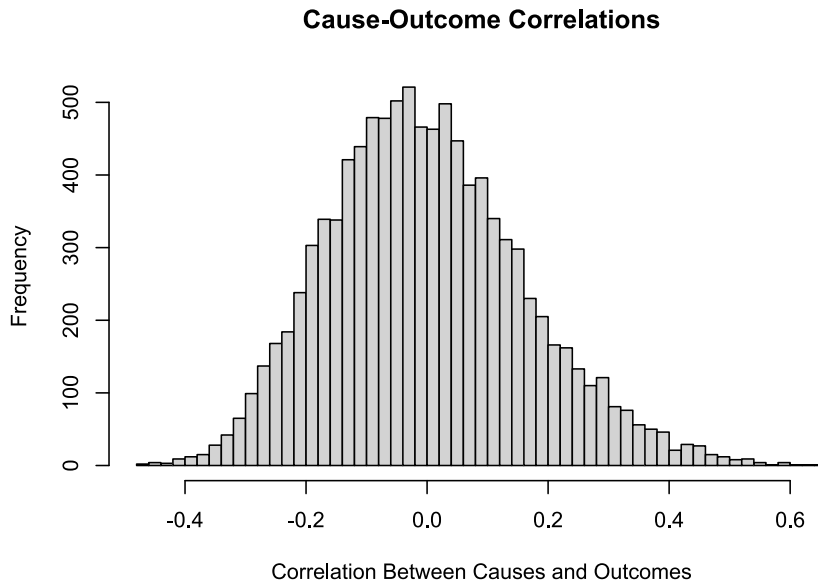


Figure 23: Distribution of correlations between causal weights \vec{u} and relative outcome size \vec{u} of randomly generated exp-exp matrices. 10,000 matrices were drawn from $ExpExp(.04, 40)$. For each matrix, the correlation between the dominant left and right eigenvector was calculated.

These matrices were generated assuming that a_{ij} , the causal effect of factor j on factor i , is independent of a_{jk} , the causal effect of factor k on factor j . However, this might not be the case, since humans make strategic decisions. Most people recognize that, say, income has a large causal effect on well-being. So, people regularly attempt to invest skills and other resources into earning a larger income. The matrix A is a result of many strategic decisions. Consequently, it's not realistic to expect *no* correlation between causes and indicators of success. However, there is a limitation to the extent that individuals *are able to* or *choose to* reinvest in any particular factor. So it's justifiable to say that the causes of success may be very different from the indicators.

5.2.3 Increasing relationship between causal factors increases inequality

One type of societal change involves allowing skills and resources to be more effectively used to increase other skills and resources without careful consideration of equity. For example, advanced degrees allow those with the resources to become more skilled. Gap-year programs allow those with parental support to gain professional connections and get into prestigious schools. Workplaces that reward long work hours benefit those without family constraints. And

states funding universities at a higher rate than community colleges give students with more privileged backgrounds the opportunity to use the skills that their (on average) more educated parents taught them.

Increasing these relationships also increases growth, as resources become used more efficiently. However, they can also increase inequality. In fact, according to the accumulation model, the primary driver of increased inequality in these scenarios is growth.

Figures 24 and 25 show what happens as the process described above happens, and the average relationship between the causes of success, $mean(a_{ij})$, increases. A larger average causal effect leads to both increased inequality and increased growth, as measured by the systemic growth factor. The results suggest that efforts to increase opportunity to use more advanced resources should keep a focus on equity as well.

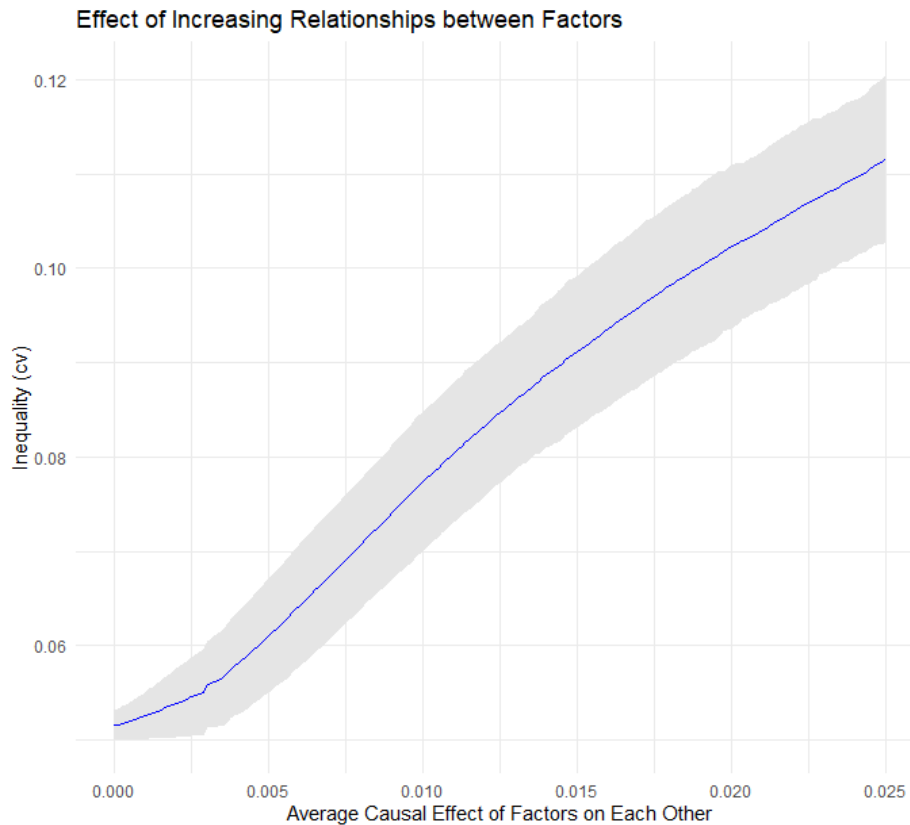


Figure 24: Relationship between the average effect of each factor on another factor and inequality. The blue line gives the mean, and the gray bar gives the standard deviation. Both mean and standard deviation were calculated using a sliding window using a total of 30,000 iterations.

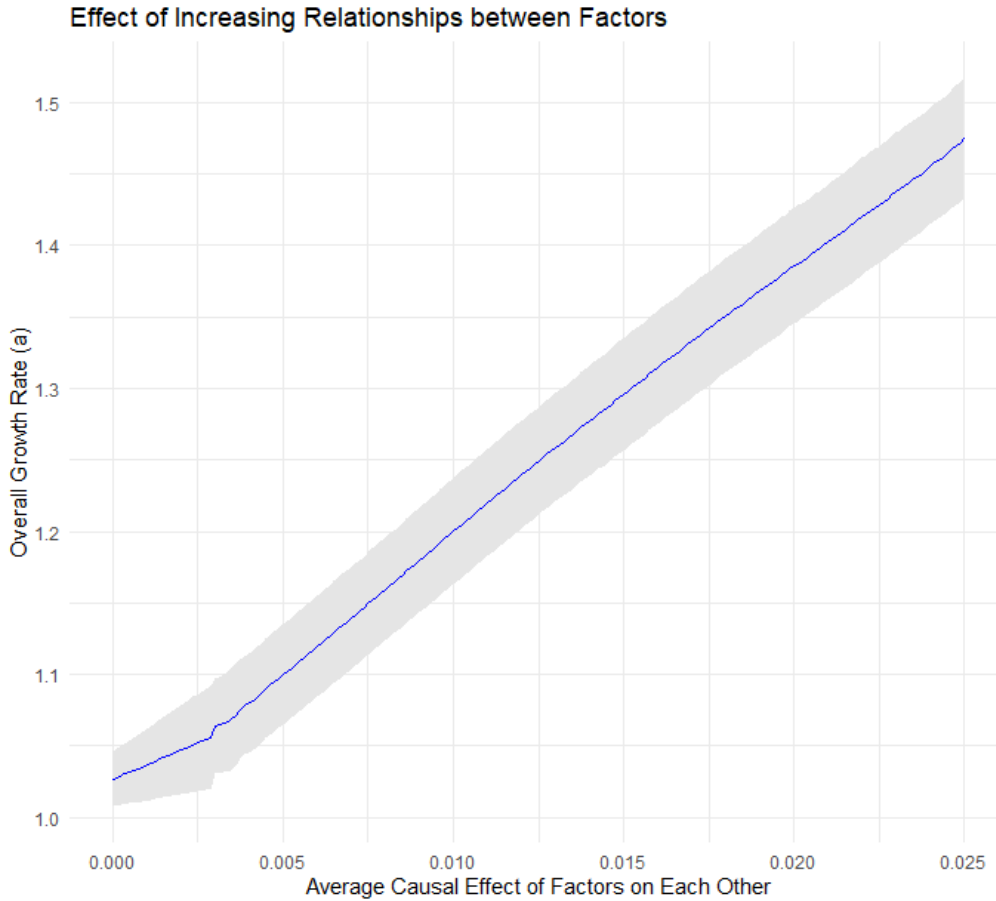


Figure 25: Relationship between the average effect of each factor on another factor and growth. The blue line gives the mean, and the gray bar gives the standard deviation. Both mean and standard deviation were calculated using a sliding window using a total of 30,000 iterations.

Simulation Details and Mathematics:

To generate these simulated data, a 20-factor accumulation model was used. At each iteration, a matrix A was drawn from $ExpExp(\eta, 20)$, where η varied from 0 to 0.2. The dominant eigenvalue a of A was calculated as the growth rate. The average causal effect of any factor on another was calculated using the formula $mean_{ij}(A_{ij}) - 1/M$, where the $1/M$ accounts for the 1's on the diagonal of M . Inequality was calculated by first generating a population of 10,000 people using the matrix, where the random amount added to each individual's factor x_i at each time step is $X_i \in Exp(1)$. Then the coefficient of variation was taken over that population.

This relationship shows up in the mathematics behind the accumulation model. Using the accumulation model, the primary mechanism through which inequality arises is by increasing the systemic interest rate, a . Inequality in the accumulation model, as measured by coefficient of variation, is:

$$inequality = \frac{\sqrt{a-1}}{\sqrt{a+1}} cv(\vec{u} \cdot \vec{X})$$

The formula $\frac{\sqrt{a-1}}{\sqrt{a+1}}$ gives an increasing function of a , which goes from zero when $a = 1$ to one as a gets large. In other words, as growth increases, so does inequality. The other component is $cv(\vec{u} \cdot \vec{X})$, which does not change as the relationship between factors increases.

5.2.4 Adding a new factor that contributes to success

We can also consider what might happen when a new factor appears, as shown in Figure 26. This is most relevant when discussing new technologies, such as the personal computer or generative AI. The introduction of a powerful new technology has the potential to improve well-being in a lot of ways also has the potential to increase inequality, as people who are more effective with that tool use the tool more effectively in a lot of ways. Meanwhile, the people who are not able to use the tools effectively won't necessarily decrease in their pre-existing skills, traits, and abilities, they just fall behind those who can adapt to the new tech.




Figure 26: Graphical representation of the accumulation network with social skills, professional network strength, knowledge, and an added new factor (red) which we can call computer skills.

The internet is a great example, since it has the potential to be an equalizer. The internet allows people to gain information about their health (Fry et al., 2015), learn about new job opportunities (Böhm, 2013), improve their professional networks (Davis et al., 2020), and find love (Bruch & Newman, 2019). If the importance of computer skills is relatively low compared with other factors, then the causes of success become diversified and inequality decreases. However, computer skills have become *very* important. Without them, people can find health misinformation (Bin Naeem & Kamel Boulos, 2021), spend far too much time playing addictive games (King & Delfabbro, 2020), feel connected to extremists through social media (Brady et al., 2021), and waste time looking for a mate in an unwelcome online market (Sparks et al., 2022). In this regime, where a single factor is very important, computer skills can increase inequality.

Figure 27 shows what happens as a new factor is introduced in the accumulation model. When the factor has a relatively low impact, inequality decreases. As the impact of the new factor becomes larger, so does inequality. Unlike with the previous simulation, where increasing inequality was only driven by growth, Figure 27D shows that inequality here is driven by both growth and by inequality in short-term investments of well-being.

This result could be construed as an argument for a diversified economy. Having a large number of relatively important traits & skills that could lead to success has the potential to keep inequality low while increasing growth. In the case where only a handful of factors are responsible for success, inequality in those factors will lead to inequality in the economy as a whole. Alternatively, this result can be viewed in the context of Goldin & Katz's race between education and technology (Goldin & Katz, 2008). If technological change happens more quickly than societies can adapt, then the small number of people who are able to use that technology will be able to benefit disproportionately.

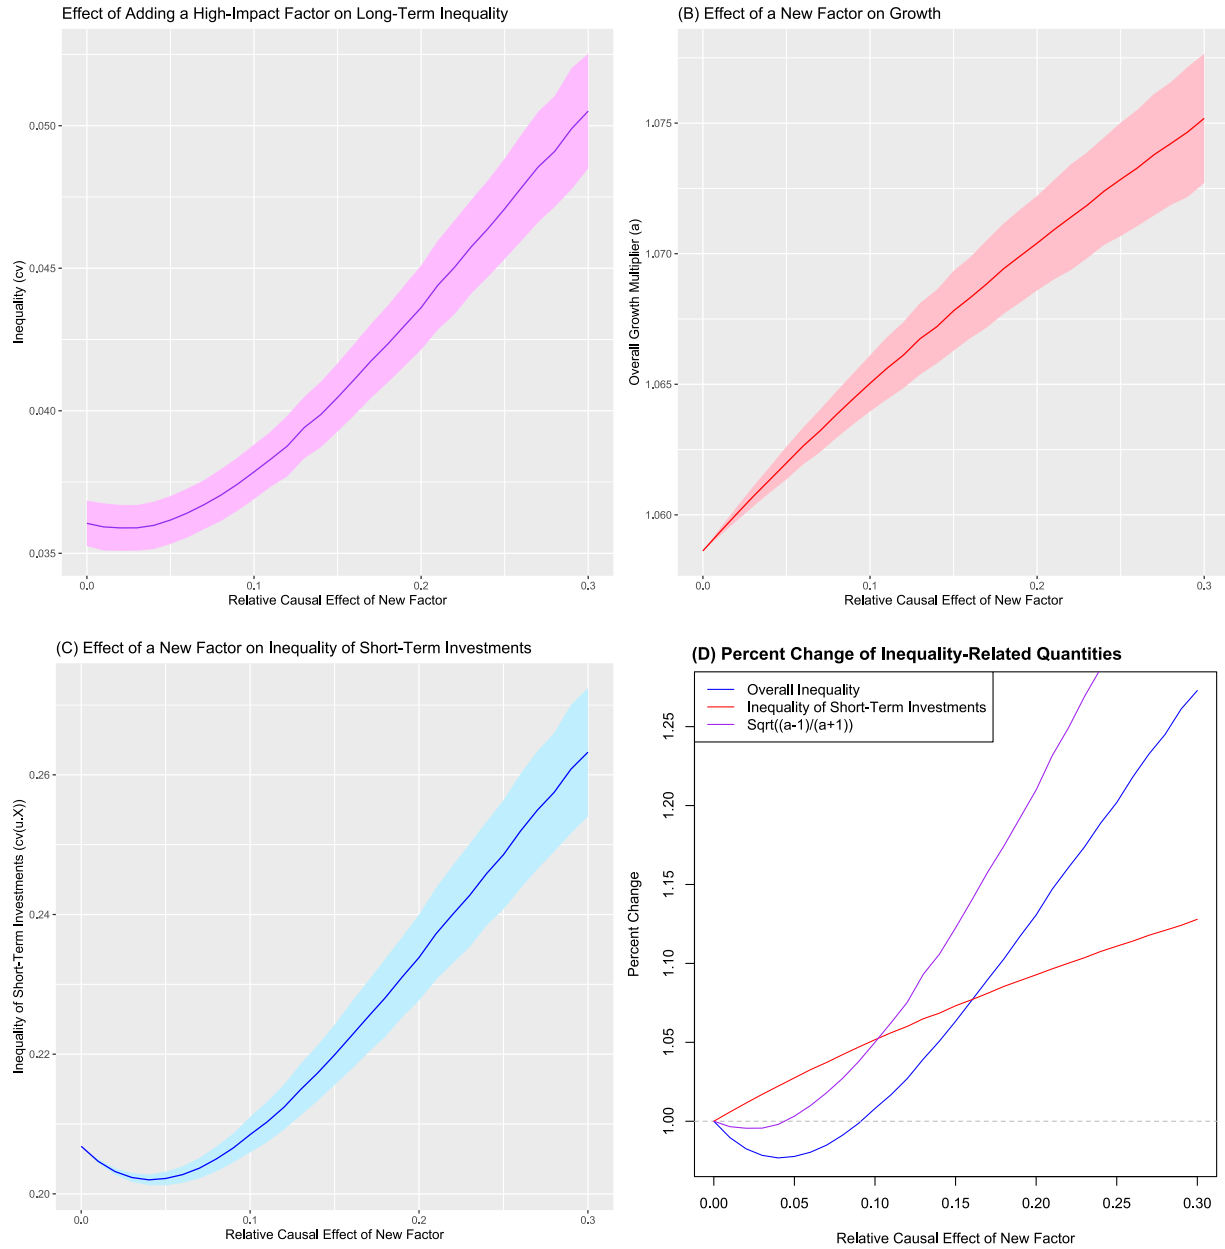


Figure 27: Effects on Inequality of a New Factor (A) Effect of new factor on overall inequality at time $t = 50$ as a function of the factor's effect as a proportion of the effects of the other factors. (B) Systemic growth multiplier a . (C) Inequality in short-term investments $cv(\vec{u} \cdot \vec{X})$. (D) Percent change of overall inequality broken up in terms of the two factors $\sqrt{\frac{a-1}{a+1}}$ and $cv(\vec{u} \cdot \vec{X})$. The y-values on the purple curve are the product of the y-values from the red and blue curves.

Mathematical Details:

Overall, inequality in this model is given by the relationship:

$$inequality = \frac{\sqrt{a-1}}{\sqrt{a+1}} cv(\vec{u} \cdot \vec{X})$$

Where a represents the overall growth multiplier (so $a - 1$ is the per-unit-time growth rate), \vec{u} is the causal strength of all the factors in determining success (including the new one), and \vec{X} is the random additive amount gained at each time (assumed to be from schooling, community support, etc.).

This means that $\vec{u} \cdot \vec{X}$ is a random variable which describes the results of training and other supports at a given time point weighted by its causal effect on well-being overall. The quantity $cv(\vec{u} \cdot \vec{X})$ describes inequality in these amounts, and this is independent of the causal effects of the matrix and more systemic growth. Figure 27C shows how this quantity changes with the relative effect of the new factor.

When there are many independent factors contributing to well-being, inequality is low as a central limit theorem-type behavior takes hold. The coefficient of variation is the standard deviation divided by the mean, and for M independent variables with roughly the same mean,

$$cv \sim \frac{sd}{mean} \sim \frac{\sqrt{M}}{M} = \frac{1}{\sqrt{M}}. \text{ So adding another variable will just decrease inequality.}$$

However, as the effects of the new factor become more important the variance of the new factor $var(X_{M+1})$ comes to dominate the variance of $\vec{u} \cdot \vec{X}$. Much of the increasing inequality we see in Figure 27 is the inequality inherent in who gets contributions to the new factor X_{M+1} .

Simulation Details for Figure 27

I used an accumulation model to generate simulated data. Each iteration used the same $M \times M$ matrix A , and then varied the effects of the new factor (factor $M + 1$). The matrix A was drawn from $ExpExp(.03, 30)$. The relative causal effect (rcf) of factor $M + 1$ was calculated as follows

$$rcf = \frac{\sum_{i=1}^M A_{i, M+1}}{(\sum_{i=1}^M \sum_{j=1}^M A_{ij}) - M}$$

The value of M was subtracted from the denominator to account for the ones on the diagonal. Then a new $(M + 1) \times (M + 1)$ matrix A' was made by putting A in the first $M \times M$ slots. The causal effect of factor $M + 1$ on each other factor i , were $A_{i, M+1}$ for $i \leq M$. These values were drawn from an exponential distribution, then normalized so that their sum was rcf . The $M + 1$ diagonal element and the causal effects of other factors on factor $M + 1$ were drawn to be

consistent with the generating process of A . Each iteration involved 1,000 randomly generated individuals. And for each value of rcf , 2000 iterations were run to make the bands in Figure 27. Different random seeds and parameters provided qualitatively similar results.

5.3 Interventions

This section focuses on designing the most effective interventions to address systemic inequality. In particular, it focuses on *direct interventions* – where an individual receives a direct benefit to some aspect(s) of their well-being. These could include a monetary donation, some form of education, socially-oriented outreach at a retirement home, or a visit to the doctor. Section 5.2 examined how inequality propagated over a span of time. Here, I examine near-term responses to interventions and how the most effective intervention might vary depending on how systemic a problem is.

5.3.1 Model

We again assume that each person has a set of M factors (continuous, positive variables) $\vec{x} = \langle x_1, x_2, \dots, x_M \rangle$. However, in this case, we are interested in how changing those factors increases a person's well-being or success. To do this, we use a positive-valued utility function $v = f(x_1, x_2, \dots, x_M)$ where v represents some form of well-being or success. We will examine the effects of direct interventions on well-being by assuming that policy makers have the ability to distribute a set amount of increases to individuals and factors to their choice of factors. To allow comparison of investments in different factors, I assume that all the x_i 's have the same units. Interventions have a constant effect on each factor x_i regardless of the current value of x_i , and all values are scaled in the units of the intervention. An intervention is assumed to be a scalar δ that describes the total contribution and a unit vector \vec{p} that describes how the contribution is distributed among causal factors, which leads to a change $\vec{x} \rightarrow \vec{x} + \delta\vec{p}$. For instance, an intervention of δ dollars could all go towards x_1 , leading to a utility (well-being) of $f(x_1 + \delta, x_2, x_3, \dots, x_M)$. Or, the δ dollars could be distributed evenly among factors, which would give a utility of $f\left(x_1 + \frac{\delta}{M}, x_2 + \frac{\delta}{M}, x_3 + \frac{\delta}{M}, \dots, x_M + \frac{\delta}{M}\right)$.

To better represent reality, this utility function should satisfy a few intuitive conditions, which can also be represented mathematically:

- **(A) Continuous and has continuous first and second derivatives.** A small investment should lead to a comparably small return.
- **(B) Positive returns.** Since each value x_i is a positively-construed resource, having *more* of x_i is always better.
 - $\frac{\partial f}{\partial x_i} > 0$ for every i
- **(C) Diminishing marginal returns within each variable.** While it might be good to have more of a good thing, the hundredth unit of the good thing does not increase your well-being as much as the first. For instance, going from 0 steps walked every day to 5000 steps provides more benefit than going from 10,000 to 15,000. However, having a nutritionist choose your meals will only be a little better than planning your own balanced meals.
 - $\frac{\partial^2 f}{\partial x_i^2} < 0$ for every i
- **(D) Eventually small returns.** For any two factors x_i and x_j , there is a point where an individual has so much of x_i that it benefits them more to gain x_j .³
 - For every i, j , increasing x_i while holding all other variables constant will eventually lead to a point where $\frac{\partial f}{\partial x_i} < \frac{\partial f}{\partial x_j}$
- **(E) Positive interactions between factors.** Factors should provide synergy, in that a larger value of one should provide a greater return on others. Having access to healthy food is more useful if one has the knowledge to know how to balance a diet and the time to cook a healthy meal.
 - $\frac{\partial^2 f}{\partial x_i \partial x_j} > 0$

As long as $0 < \alpha_i < 1$ for each i , the Cobb-Douglas utility function

$$f(\vec{x}) = x_1^{\alpha_1} x_2^{\alpha_2} \cdots x_M^{\alpha_M}$$

satisfies all of the conditions above. The value $\tilde{\alpha} = \sum_i \alpha_i$ gives the fastest possible rate of growth of this function. For instance, if $\tilde{\alpha} = 2$, then well-being can grow at most quadratically as a

³ We could also assume the stronger condition of “vanishing marginal returns” where the derivative $\frac{\partial f}{\partial x_i}$ goes to zero for large x_i , holding other variables constant.

function of the amount invested. In simulations in this section, I will use Cobb-Douglas utility functions with randomly generated exponents.

By assumption, all interventions increase well-being for their recipients. However, some interventions may be more effective than others, or cause the rich to get richer. These are the questions we explore in this section.

5.3.2 A Note about the Difference between Gaps and Inequality

Inequality metrics commonly used for income and wealth inequality, such as the Gini coefficient or the coefficient of variation, are scale-independent since they need to account for inflation. If everyone gets a 3% raise, then inequality does not increase. However, when talking about interventions, this can create counterintuitive results. For instance assume Bertie has a wealth of \$100, Slim has \$500, and we have \$300 to gift which we can split up any way we want. Intuition says that it is unfair to give Bertie \$100 and Slim \$200, since we're helping the rich get richer. However, this allocation *decreases* wealth inequality using standard inequality metrics. Since Slim has five times as much as Bertie, the allocation that would leave inequality unchanged is to give Bertie \$50 and Slim \$250. Both things can happen at the same time: scale-independent metrics inequality are decreasing, and the rich can get richer at a faster rate than the poor. Scale-independent inequality metrics are not as useful for studying direct interventions targeted at systemic inequality, for a number of reasons. The first is that interventions typically target a subset of the population rather than the economy as a whole. Giving Bertie and Slim money will not influence macroeconomic inflation. In addition, systemic inequality involves many, often non-monetary, factors. Returns to education, for instance, are not likely to scale like income. There is a statistical reason to focus on gaps as well. For the model used here, scale invariance can create overly sensitive dependence on initial values. Consider the coefficient of variation, which is the standard deviation of a group divided by the mean.

$$cv = \frac{sd}{mean}$$

This function can be very sensitive to changes in the denominator when the denominator is small. So, for instance, if the initial mean of the population is 1 and an intervention increases the mean by 1, any changes in the standard deviation will likely play a secondary role to the fact that the intervention doubled the denominator.

While it can be useful to look at inequality through multiple lenses, in this section, references to inequality refer to gaps, rather than scale-independent economic measures of inequality. Our central question here is: Which interventions cause *gaps* between the rich and the poor to decrease? Does a given intervention help the rich get richer, or do the poor get richer in comparison to the rich? Gaps can be measured using simple subtraction (in the case of two individuals), difference in means (in the case of two groups), or standard deviation (in the case of a single group).

5.3.3 Targeted Interventions

In a systemic inequality regime, blanket interventions that are not targeted at specific groups have the potential to increase inequality. To see this, we start with a simple example, and work our way up to a more complicated intervention between groups.

Consider the (overly simplistic) case where a person's health h is a product of their knowledge k and wealth w , so that⁴ $h = f(k, w) = kw$. Staying healthy requires the knowledge about what to do and the money to take advantage of the knowledge by paying for health care, gym memberships, healthy food, etc. In our scenario, two people, Gregor and Johann, have equally little knowledge about health, $k = 1$. However, Gregor has more wealth than Johann. Johann's wealth is 3, so his health is $1 \cdot 3 = 3$. Gregor's wealth is 5, so he has better health of $1 \cdot 5 = 5$. A public information intervention helps both Gregor and Johann learn about fitness, increasing their knowledge by 2. Figure 28 shows what happens to their health.

⁴ This function does not have diminishing returns, and therefore does not count as a utility function as we've defined it. However, using simple values will prove illustrative at first.




Figure 28: Effects of a uniform intervention on two individuals.

While both Gregor (brown) and Johann (blue) have the same initial and ending knowledge, Gregor’s health increases more because he has more resources to take advantage of the greater knowledge.

The result is an increase in inequality, through a rich-get-richer mechanism. The gap between Gregor and Johann increases from 2 to 8. Holding each person’s wealth constant, Gregor’s health is given by $h = 5k$, while Johann’s health is $h = 3k$. Each additional bit of knowledge gives Gregor a larger benefit, because he has the wealth to take better advantage of it. Perhaps Gregor can afford that gym membership that he has just learned he should have.

Now let’s consider a slightly different scenario with utility function⁵ $f(k, w, u) = k^{0.5}w^{0.7}u^{0.2}$. In this case, u describes a hidden factor that influence one’s health, but that might not be obvious or easy for researchers to measure. Such factors might include proficiency in communicating

⁵ This utility function does have decreasing marginal returns in each variable.

with doctors (Riedl & Schüßler, 2017), conscientiousness (Bogg & Roberts, 2004), or sensitivity to one’s environment (Boyce, 2019). Further, let us introduce Emmy who, with respect to knowledge and wealth, is identical to Johann. However, Emmy was given the confidence and skills to recognize health concerns and communicate clearly with health care workers. As shown in Table 6, Emmy has more of the hidden factor than Gregor or Johann.

Table 6: Initial values of knowledge, wealth, and the hidden factor

	Knowledge (Initial)	Wealth	Hidden Factor	Health Pre- Intervention	Health Post- Intervention	Effect of Intervention
<i>Johann</i>	1	3	1	2.2	4.3	2.1
<i>Gregor</i>	1	5	1	3.1	6.2	3.1
<i>Emmy</i>	1	5	5	4.1	8.1	4.0

As before, Gregor and Emmy, who have more wealth, benefit more from the intervention than Johann does. However, Emmy received *more* benefit from the intervention than Gregor did. The hidden variables increased the effect of the rich-get-richer mechanism. In regression models, this effect can be seen when errors are correlated with predictor variables after controlling for other available variables (and assuming the model is appropriately specified).

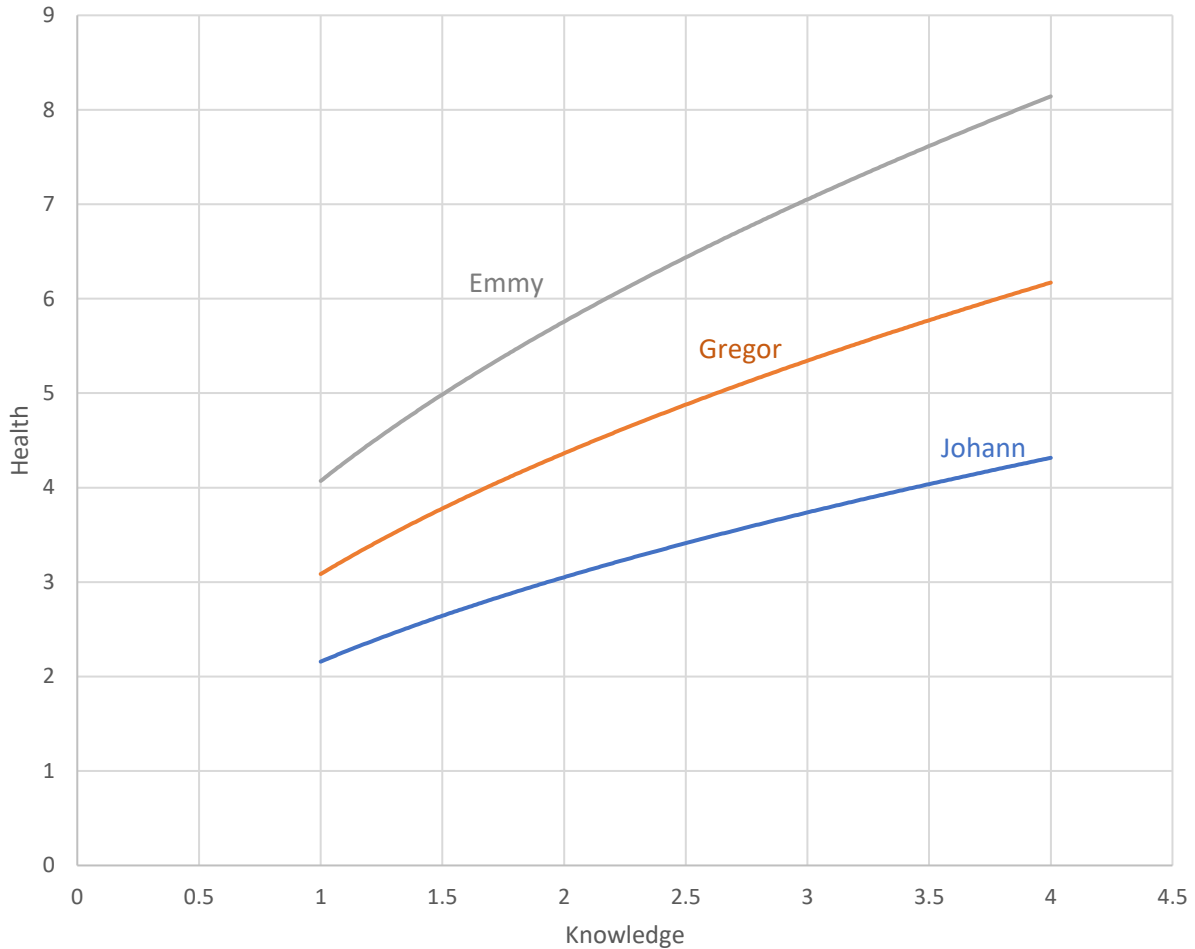


Figure 29: Effects of the health intervention on Emmy, Gregor, and Johann

One mechanism that can counteract the rich-get-richer mechanism comes from diminishing returns. Giving a poor person \$1000 will likely provide a lot more benefit than giving a millionaire \$1000. As gaps between the well-off and the not-so-well-off get larger, interventions that provide the same benefit to everyone will tend to *decrease* inequality. In practice, and in this model, there is a counterbalancing effect between the mechanisms of rich-get-richer and diminishing marginal returns. In practice, though, if we want to reduce the gaps between Emmy, Johann, and Gregor, we should focus our intervention on helping Gregor. A targeted intervention will be more helpful in this case.

Simulation

Rather than focusing on a few people, we now turn our attention to gaps between groups caused by systemic inequality. Our intuition from Emmy, Johann, and Gregor tells us that a non-targeted, universal intervention has the potential to increase inequality between groups due to the large number of (potentially hidden) factors, and that hidden factors may exacerbate that effect. The approach here is to consider two populations (call them “Group High” and “Group Low”). An individual’s well-being, regardless of their group, is a function of ten factors. In h of the factors, which I’ll call *inequality-generating variables*, Group Low is disadvantaged compared with Group High. However, the groups are similar in the remaining dimensions. We can consider these groups to be rich kids and poor kids within an economically segregated school district, racial groups within a country, populations of different countries, or just any two groups where one group has a structural advantage compared with the other.

I consider a two-by-two simulated experimental design. One treatment dimension focuses on whether the intervention is targeted or not: Either resources are split evenly among all individuals, or the resources are given only to those in the group with lower well-being. The other treatment dimension focuses on whether the inequality-generating variables are hidden: Either interventionists are aware of all variables, and can intervene in the variables where there are gaps. Or the variables that cause differences between groups are hidden, so interventions can only intervene on variables where groups do not differ substantially.

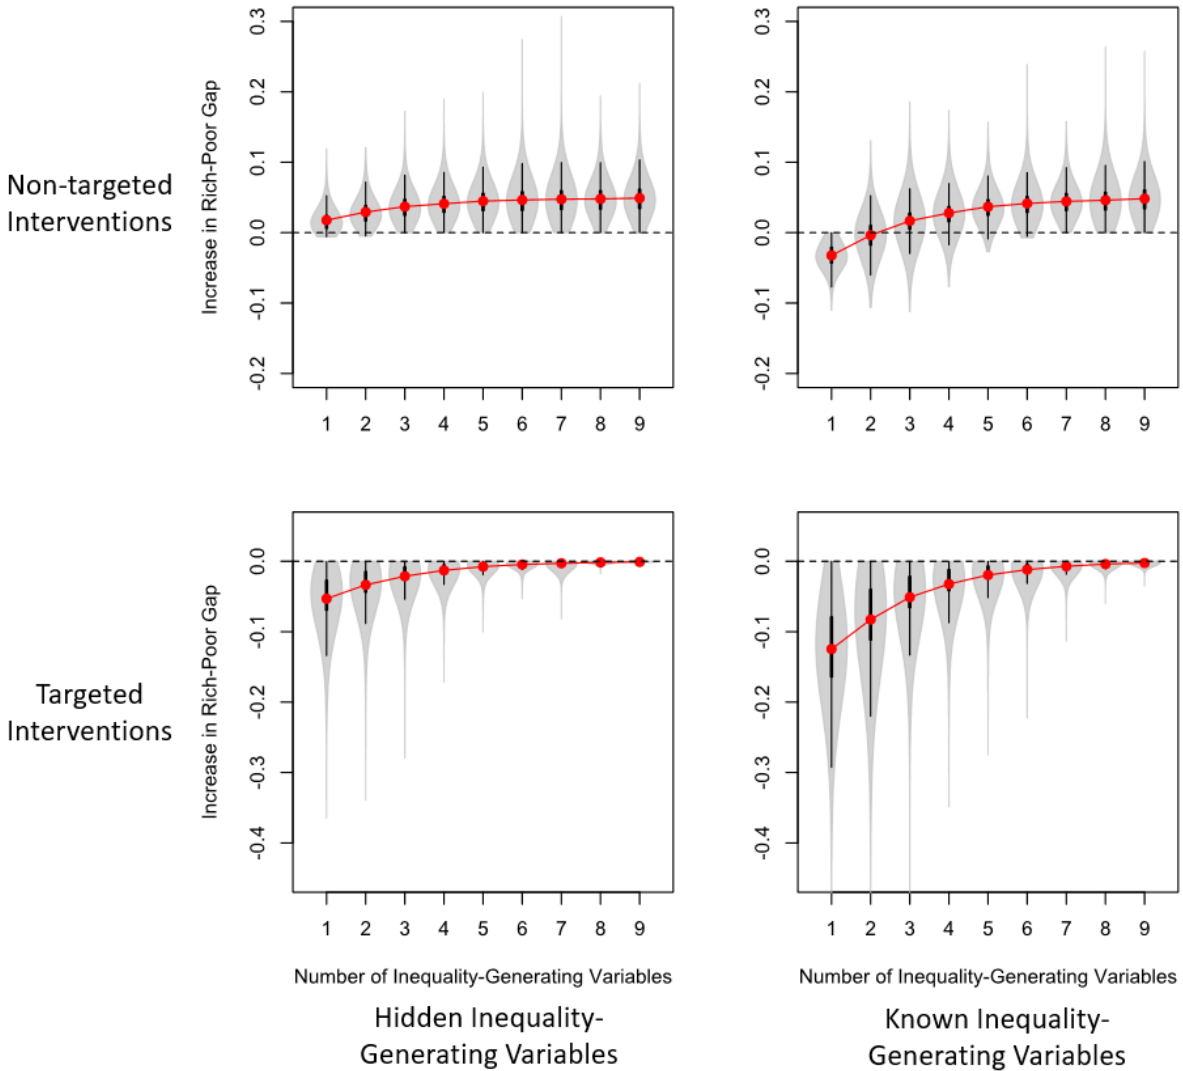


Figure 30: Effects of interventions on the average well-being gap between Group High and Group Low. Each gray distribution represents 1000 simulations. Positive y-values correspond to interventions that increase inequality. The two graphs on the top correspond to the case where interventions are given to all individuals, while the graphs on the bottom represent the case where interventions are targeted only at the disadvantaged Group Low. The left graphs correspond to where the variables that cause inequality are hidden, so interventions focus on a single known variable where both groups are similar. In the right graphs, interventions focus on one of the dimensions where Group High is better off. Inequality becomes more larger and more systemic as the number of inequality-generating variables increases.

The top row of Figure 30 shows that non-targeted interventions tend to increase pre-existing gaps. The clear exception to this case is where there is only one non-hidden inequality-generating variable. In this case, interventions targeting this single factor have the potential to close gaps due to the diminishing marginal returns experience by Group High.

The bottom row gives the results of targeted interventions received only by the disadvantaged group. Targeted interventions clearly decrease gaps between the advantaged and disadvantaged. However, as the number of causes of inequality increases, intervening in a single factor become less useful overall, as an intervention in one factor becomes less able to counteract the systemic differences between groups.

Overall, the results show that targeted interventions are more effective at reducing inequality than non-targeted interventions. However, as inequality becomes more larger and more systemic, the effects of the same single-factor intervention dwindle.

These results suggest that interventions which do not target disadvantaged populations risk increasing inequality – a result consistent with prior research (Lorenc et al., 2013; Veinot et al., 2018). In addition, there is value to uncovering the “hidden” variables that create differences between groups. Targeting the factors that influence overall inequality, rather than just those that are easy to intervene on, can decrease inequality more effectively.

Simulation Details:

This simulation involved ten variables, h of which are inequality-generating. The simulation was run 1000 times for each value of h , with 10,000 individuals split equally into two groups of 5,000. For each iteration, 10 Cobb-Douglas exponents α_i were randomly chosen from a uniform distribution, and then normalized so that they added to 3. For each member of Group High, the value of factor x_i was drawn from an exponential distribution with mean $2\alpha_i$. For Group Low, the values of the $10 - h$ non-inequality-generating (IG) variables were drawn in the same way as for Group High. However, on the h IG variables, all members of Group Low were assigned a value of 0.1.

Once initial values for all agents were generated, their pre-intervention well-being was calculated using the Cobb-Douglas function $f(\vec{x}) = \sum_i x_i^{\alpha_i}$. A 2×2 set of interventions was then given. In each treatment case, the well-being of individuals pre and post intervention was subtracted, averaged among the group, and then plotted in Figure 30. pre-post intervention difference in average well-being between Group High and Group Low was calculated. One treatment dimension (targeted vs non-targeted) involved who received the intervention: In the non-targeted case, all individuals received a boost $x_i \rightarrow x_i + 0.5$ in the intervention variable. In the targeted case, all resources were allocated toward Group Low, so that members of Group Low received a

double-boost $x_i \rightarrow x_i + 1$, while members of Group High received no intervention. The other treatment dimension (hidden vs not-hidden IG variables) involved changing the variables that received an intervention. In the hidden case, the intervention happened to a randomly-chosen non-IG variable, along which both groups were similar. In the non-hidden case, where interventions ostensibly understand all the causes of inequality, the intervention happened on one of the IG variables. Qualitatively similar patterns were shown as the model parameters were varied.

5.3.4 *Single vs Multifactor Interventions*

I've shown that targeted interventions are more effective for decreasing inequality than those that impact everyone universally. So this subsection focuses *only* on interventions that target disadvantaged groups. This allows us to examine the types of targeted interventions that are most effective, rather than having to focus on gaps between groups. In particular, I compare interventions where resources are allocated to a single factor with those where the same resource are allocated to many factors at once.

I will show that, for an individual or group who is disadvantaged in only one factor, a single factor intervention will generally be the best choice. However, if the individual or group is disadvantaged on *many* factors, as in the case of systemic inequality, then multi-factor interventions will be more effective than single-factor interventions. As before, we start with an intuitive example, and then lead into simulations and mathematics.

Optimal Interventions for Individuals

Consider success that is caused by two factors, knowledge k and wealth w , and a utility function $f(k, w)$ with the previously stated properties of a utility function. Intuitively, we can say that it's better to have more wealth. However, too much wealth isn't as helpful if you don't know what to do with it. A similar argument holds for having knowledge without the resources to make your ideas come to life.

Consider an individual with a given value of (k, w) . What type of intervention is most effective for this person? Should the intervention go completely in to knowledge ($k \rightarrow k + \delta$) ? Or wealth ($w \rightarrow w + \delta$) ? Or should the intervention go into wealth and knowledge simultaneously in some proportion $(k, w) \rightarrow (k + p_1\delta, w + p_2\delta)$, where $p_1 + p_2 = 1$?




Figure 31: Optimal interventions when success is dependent on two factors. The blue line is the equilibrium value. Arrows represent optimal intervention types for points in different regions.

Figure 31 gives a visual explanation of the answer to these questions. The formal proofs of these results for the Cobb-Douglas utility function in M factors are given in the Appendix. An individual who has very little wealth, but a lot of knowledge (orange star) will benefit the most through an intervention that provides only wealth. Similarly, someone with a lot of wealth but low knowledge (green square) will benefit most from a single-factor, knowledge-based intervention. Eventually, however, as diminishing returns come into play, the benefits of a

single-factor intervention decrease until single-factor interventions in both factors are equally effective (black line). The equilibrium⁶ curve is the point where single-factor interventions in both knowledge and wealth are equally useful.

For very small interventions, an individual on the equilibrium line would benefit equally from a knowledge intervention, a wealth intervention, or any multi-factor intervention between the two. However, for larger interventions, the most effective intervention for someone at equilibrium involves multiple factors. To see why, consider the red, blue, and purple arrows. To get to the end of the purple arrow, which by assumption is at a higher level of well-being than the purple point, there are two possibilities. One possibility is a single-factor intervention in wealth (red) and then a single-factor intervention in knowledge (blue). Another possibility is a multi-factor intervention along the equilibrium line (purple). However, since the cost of an intervention is just the distance along the path, the multi-factor intervention gets the person to the same point with less cost. This is because the single-factor interventions suffer diminishing returns as we move away from the equilibrium line. The most effective intervention will be in the proportions given on the equilibrium line. For this purple dot, the most effective intervention of δ dollars is

$$(k, w) \rightarrow \left(k + \frac{1}{3}\delta, w + \frac{2}{3}\delta\right).$$

Note that this model does not account for the time-based effects discussed in Section 5.2. For instance, giving teenagers with no money financial training about how to handle their wealth might quickly be forgotten without an opportunity to use that knowledge. In this case, the *long-term* effectiveness of an intervention providing knowledge would be decreased by the lack of wealth. Alternatively, providing large amounts of wealth to a teenager may provide them access to a financial planner who could teach them long-term strategies.

For the Cobb-Douglas utility function $f(x_1, x_2, \dots, x_M) = x_1^{\alpha_1} x_2^{\alpha_2} \dots x_M^{\alpha_M}$, we can explicitly write out some formulas. In this case the equilibrium curve is the straight line defined by the equations:

$$\frac{x_1}{\alpha_1} = \frac{x_2}{\alpha_2} = \dots = \frac{x_M}{\alpha_M}$$

⁶ What I call the equilibrium curve might also be called the path of steepest ascent or a maximal growth path. For the Cobb-Douglas function, the *rate* of steepest ascent is $\tilde{\alpha} = \sum_i \alpha_i$.

We can explicitly calculate the proportional loss generated by a single-factor intervention for an individual at equilibrium.⁷ The proportional loss is the lost benefit from investing in a single-factor intervention divided by the total benefit of the optimal multi-factor intervention. To first order, the proportional loss from investing δ dollars in factor j alone is:

$$\begin{aligned} \text{proportional loss} &= \frac{(\text{multi-factor benefit} - \text{single-factor benefit})}{\text{multi-factor benefit}} \\ &= \frac{1}{2} \left[\frac{1}{\alpha_j} - \frac{1}{\tilde{\alpha}} \right] \gamma \delta \end{aligned}$$

Where $\tilde{\alpha} = \sum_i \alpha_i$ and $\gamma = \frac{\alpha_1}{x_1} = \dots = \frac{\alpha_M}{x_M}$.

As the size of the intervention, δ , increases, so does the proportional loss. This suggests that larger, more costly interventions should target multiple factors, since they might create the largest waste. The factor γ is larger for more disadvantage individuals (when x_i is small). So interventions that target the worse off (which, by assumption, is all targeted interventions) should focus on multiple factors.

To understand the effect of $\left[\frac{1}{\alpha_j} - \frac{1}{\tilde{\alpha}} \right]$, it will help to think of α_j as a measure of the importance of factor j in an individual's well-being⁸. Similarly, $\tilde{\alpha}$ can be thought of as the total importance of *all* factors in an individual's well-being. So $\left[\frac{1}{\alpha_j} - \frac{1}{\tilde{\alpha}} \right]$ gets larger as factor j is less important, as shown in Figure 32. If α_j is reasonably large compared to $\tilde{\alpha}$, so that well-being is mostly caused by a single factor, say wealth, then single-factor interventions in wealth are not as inefficient. However, as we move into the systemic inequality regime, where many factors contribute to inequality (and $\alpha_j \ll \tilde{\alpha}$), then single-factor interventions become very inefficient.

⁷ A detailed derivation of this is in the appendix.

⁸ One way to see this is to look at a point on the equilibrium line. On that line, the ratio of two factors is equal to the ratio of their exponents $\frac{x_i}{x_j} = \frac{\alpha_i}{\alpha_j}$. So, if for instance, $\alpha_{knowledge}$ is twice as big as α_{wealth} , then at equilibrium an individual's knowledge would be twice their wealth.




Figure 32: Relationship between α_j and the factor $\left[\frac{1}{\alpha_j} - \frac{1}{\tilde{\alpha}}\right]$ in the proportional loss equation. This graph assumes that $\tilde{\alpha} = 1$. A similar shape holds for other values of $\tilde{\alpha}$.

Optimal Interventions for Groups

The previous discussion focused on individuals. Groups, in contrast, often have significant heterogeneity which should influence the optimal choice of intervention. We’ll see that, when all members of a group are disadvantaged for a single reason, it makes sense to invest in that factor. However, when group members are disadvantaged in many dimensions, then a multi-factor intervention is again the best choice.

Consider the two groups shown in Figure 33. The black line is the line of equilibrium, where single-factor investments in knowledge and wealth give the same return. Group A has a mean knowledge of 15 and a mean wealth of 10. Group B has a mean knowledge of 10 and a mean wealth of 15. In a real-world scenario, we might notice many people in Group A dropping out of school because they don’t have enough money to pay for it. In the case of Group B, which is more in the systemic inequality regime, students would likely have very different explanations for dropping out.⁹

⁹ This is the case for many young adults who leave school without graduating (J. Johnson & Rockkind, 2009).

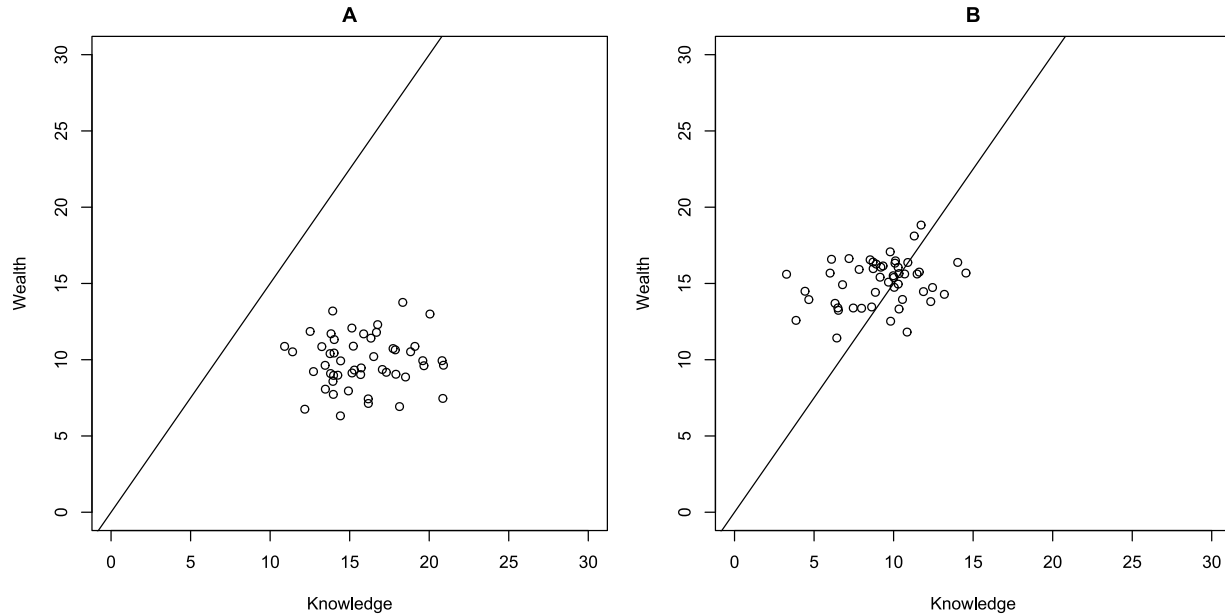


Figure 33: Two heterogeneous groups of people. Group A has knowledge, but less wealth. Group B is more balanced in their relative amounts of wealth and knowledge.

Our previous discussion shows that the best intervention for every member of Group A is a single-factor intervention focused on wealth. However, each side of the equilibrium line contains members of Group B. Very finely-tuned individual-level interventions might give some portion of Group B wealth, while others would get knowledge. However, in practice, such finely-tuned interventions are not possible due to measurement error and administrative challenges. This is especially true in the case with many dimensions. If we have to choose a single intervention for all members of Group B, we would expect that a multi-factor intervention would increase *average* well-being the most, since it helps those who are disadvantaged in wealth, disadvantaged in knowledge, and those who are balanced between the two factors.

An important caveat is necessary: Very large or very small interventions do not follow our intuition here. For instance, a very large wealth intervention for Group A (for instance by adding 30 to wealth) would lead to diminishing returns, where Group A needs more knowledge to manage their wealth. Similarly, a very small multi-factor intervention for Group B would likely not be much of an improvement over a single-factor intervention for the reasons outlined earlier. A more robust simulation also fits our intuition. Consider a group of people whose well-being is determined by ten factors, but the group is systematically disadvantaged on one of those factors

by being below equilibrium. Figure 34 examines the relative effectiveness of an intervention as we make that group more or less disadvantaged in the single factor.

Return on Investment of Intervention as a Function of Systemicity of Inequality

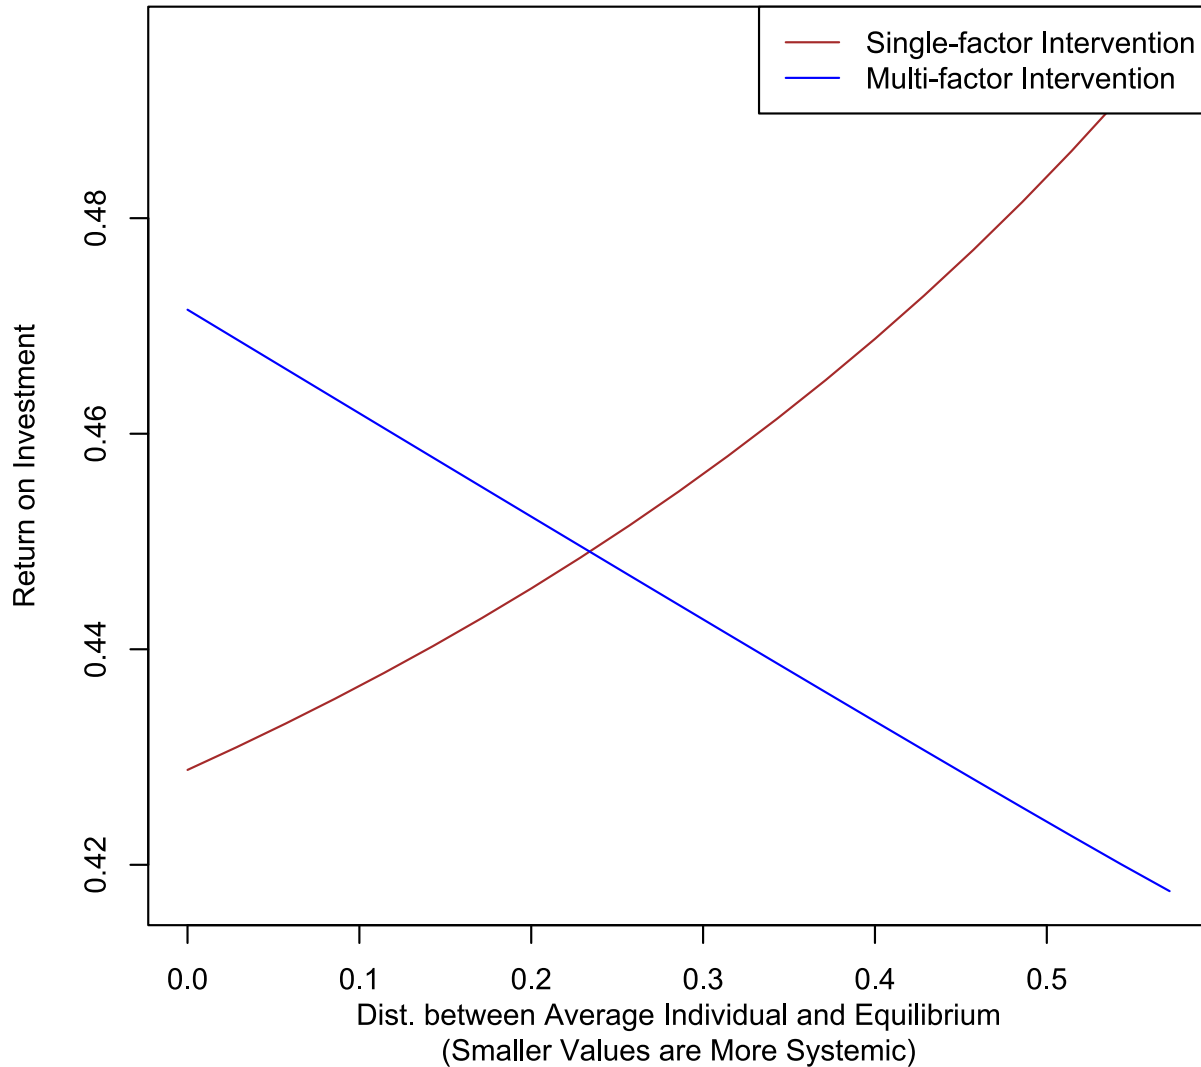


Figure 34: Comparative return to single-factor vs multi-factor interventions as a group moves closer to equilibrium. The x-axis describes the distance between the centroid of the data and the equilibrium line. The y-axis represents the increase in well-being divided by the size of the investment.

When the group is significantly lower in a single factor (the right side of Figure 34), then single-factor interventions are more effective. However, as improvements get made which move the average of the group toward the equilibrium line, the causes of well-being become more various

and idiosyncratic (left side of Figure 34). In this case, which represents more systemic inequality, multi-factor interventions become more effective.

These results support the idea that, when inequality is systemic, interventionists should focus on improving many factors at once. Intervening in multiple factors at once has been shown a powerful approach in education (Levin & García, 2018; Miller & Weiss, 2022) and aging (Shaposhnikov et al., 2022). Furthermore, multi-pronged approaches are represented in the recommendations of many professional bodies (Lichtenstein et al., 2006; National Research Council, 2001; WHO, 2022).

Simulation Details: Figure 34 was created using a simulated experiment with a 10-factor Cobb-Douglas utility function. Cobb-Douglas exponents were chosen randomly from a uniform distribution, sorted so that $\alpha_i < \alpha_{i+1}$, and then normalized so they summed to two. A single population was created, where each individual k 's initial values of x_i (for $i \neq 5$) were drawn from a normal distribution which depended on the value of α_i .

$$\begin{aligned} x_{i,k} &\in \mathcal{N}(\mu_i, \sigma_i) \\ \mu_i &\in \begin{cases} \text{Unif}(8\alpha_i, 10\alpha_i) & \text{if } i \neq 5 \\ \text{Unif}(6\alpha_i, 7\alpha_i) & \text{if } i = 5 \end{cases} \\ \sigma_i &\in \text{Unif}(\alpha_i, 3\alpha_i) \end{aligned}$$

The middle value of x_5 was chosen to be different since it had an average effect on well-being compared with the other variables. Very small values of $x_{i,k}$ were rounded up to 0.01 to avoid singularities in the model. This rounding up involved less than 0.01% of values. The initial population values were then roughly on the equilibrium line $\frac{x_1}{\alpha_1} = \dots = \frac{x_M}{\alpha_M}$, except for x_5 which was below what would be expected at equilibrium. Initial population values were then varied continuously as function of t along a line until the centroid of the population was $\langle 10\alpha_1, 10\alpha_2, \dots, 10\alpha_M \rangle$. This mean the points varied from being off the equilibrium line in the direction of small x_5 to being on the equilibrium. For each t , both a single-factor and multi-factor intervention were given equivalent to $\frac{10\alpha_5 - \mu_5}{2}$, which was 0.29 in this simulation. The benefit of each intervention on average well-being was calculated, and then divided by 0.29 to get the return on investment.

5.4 Conclusion

In this chapter, I first provided a framework for thinking of systemic inequality as arising from a large number of causally interrelated factors, many of which are hard to understand or influence at scale. The research literature supports the idea that some of the stickiest types of inequality do arise from complex webs of causes that propagate through time, embedded in individuals, families, institutions, and societal structures. This framework was broad, in that it can take into account many different analytical approaches. I then gave two specific modeling approaches for understanding systemic inequality. Because these approaches rely on abstract models, they provide generalizable results for the reader to apply to a broad array of areas. The first approach involved the accumulation model from Chapter Four, which explores how multifactor cumulative advantage or disadvantage processes change over time. Section 5.2 showed how the indicators of success need not be the causes, and that, without a focus on equity, opportunities for individuals to more effectively use their skills could increase inequality. It also examined how the addition of a factor influencing success, such as a new technology, could either (a) decrease inequality if it provided a relatively independent way for individuals to be successful, or (b) increase inequality if the new factor had a large enough impact on overall success. The second approach involved a utility function with diminishing marginal returns to examine the interventions that are most effective at dealing with systemic inequality. In Section 5.3 I showed that, when inequality is systemic, interventions targeted at the most disadvantaged groups and at many factors simultaneously will tend to reduce inequality most effectively. However, in the non-systemic regime, when there is only one known factor which generates the inequality, interventionists would do well to focus on that single factor.

Given the modern day availability of data, a reader might ask why this modeling approach is valuable. There are a few reasons. First, regardless of a researcher's care in explaining the difference between causality and correlation, the choice of variables to study directs a reader's mind to certain causal relationships and hides others. For instance, consider the fact that the number of successful friends a person has is predictive of their economic mobility (Chetty et al., 2022). Many people may consider this as advice about how to move up the economic ladder. However, those readers may be ignoring the causal effects of, say, growing up in a high-income neighborhood. To their credit, Chetty, Jackson, and colleagues explore some of these questions. However, not all researchers are able to explore these questions. And even with sophisticated

causal inference methods, any conclusions about causal effects are at most bounded. So it is worth highlighting one of the key results of this chapter – that the strongest indicators of success need not have the strongest causal effect.

Second, the abundance of causal variables and practical measurement challenges mean that data-focused approaches will always be ignoring some of the systemic inequality story. Using models with arbitrarily large numbers of variables provides a different set of insights. Of course, the method taken in this chapter also has its limitations. However, the results are broadly generalizable to domains which are data-scarce. Leaders and policy makers often need to make decisions without the benefit of multi-pronged randomized controlled trials. These results give insight into how to proceed when empirical causal results are limited. Even more, this paper makes an argument that randomized controlled trials are limited in their ability to understand systemic inequality. Using the scientific method of changing one variable at a time ignores potentially synergistic effects that arise from many interacting factors.

5.5 Appendix to Chapter 5

This appendix supplements Chapter Five. It has results about the optimal intervention for an individual with M factors and a Cobb-Douglas utility function determining their well-being.

Theorem A1: Consider the case of a small, finite investment δ which can be invested into M factors $\{x_1, x_2, \dots, x_M\}$ according to some allocation vector \vec{p} (with $\sum_i p_i = 1$). Let an individual's well-being be given by a Cobb-Douglas utility function $U = f(\vec{x}) = x_1^{\alpha_1} x_2^{\alpha_2} \dots x_M^{\alpha_M}$ with $0 < \alpha_i < 1$, and set $\tilde{\alpha} = \sum_i \alpha_i$. Assume that the variables $\{x_i\}$ are scaled so that an investment of δ dollars increases every variable x_i to $x_i + \delta$, and that, pre-intervention, the values are $\vec{x}_0 = \langle x_{1,0}, x_{2,0}, \dots, x_{M,0} \rangle$. Furthermore, and without loss of generality, assume that

$$\frac{\alpha_1}{x_{1,0}} \geq \frac{\alpha_2}{x_{2,0}} \geq \dots \geq \frac{\alpha_M}{x_{M,0}}.$$

(A) If $\frac{\alpha_1}{x_{1,0}} > \frac{\alpha_2}{x_{2,0}}$ at \vec{x}_0 , then the optimal intervention is to invest completely in x_1 . (i.e. $\vec{p} = \langle 1, 0, 0, \dots, 0 \rangle$)

(B) If $\frac{\alpha_1}{x_{1,0}} = \frac{\alpha_2}{x_{2,0}} > \frac{\alpha_3}{x_{3,0}}$ at \vec{x}_0 , then the optimal intervention is $\vec{p} = \frac{1}{\alpha_1 + \alpha_2} \langle \alpha_1, \alpha_2, 0, 0, \dots, 0 \rangle$

(C) If $\frac{\alpha_1}{x_{1,0}} = \frac{\alpha_2}{x_{2,0}} = \dots = \frac{\alpha_k}{x_{k,0}} > \frac{\alpha_{k+1}}{x_{k+1,0}}$ at \vec{x}_0 , then the optimal intervention is $\vec{p} =$

$$\frac{1}{\sum_{i=1}^k \alpha_i} \langle \alpha_1, \alpha_2, \dots, \alpha_k, 0, \dots, 0 \rangle$$

(D) If $\frac{\alpha_1}{x_{1,0}} = \frac{\alpha_2}{x_{2,0}} = \dots = \frac{\alpha_M}{x_{M,0}}$ at \vec{x}_0 , then the optimal intervention is $\vec{p} = \frac{1}{\bar{\alpha}} \langle \alpha_1, \alpha_2, \dots, \alpha_M \rangle$.

Corollary: The optimal *path* as investments increase is given by the process in the Theorem. In other words, the optimal intervention for an individual becomes more and more multi-factor as the total amount of investment increases.

In other words, as individuals get to “higher states of well-being”, they need more comprehensive interventions. For someone in a poor country, providing food and basic medical care may increase well-being significantly. However, for someone in a richer country where food and basic medicine are available (and who is therefore better off than the person in the poor country), improving well-being may require nutrition education, opportunities to exercise, and sophisticated forms of health insurance.

The result above is only about optimal interventions for individuals, not groups. Groups are more complicated, since any intervention needs to account for the heterogeneity in the group. Chapter Five examines this case.

Theorem A2: In cases (B)-(D) of Theorem A1, the proportional loss of utility from investing in a single factor, rather than the optimal multi-factor intervention, is given by:

$$\begin{aligned} \text{proportional loss} &= \frac{(\text{multi-factor benefit} - \text{single-factor benefit})}{\text{multi-factor benefit}} \\ &= \frac{1}{2} \left[\frac{1}{\alpha_j} - \frac{1}{\sum_{i=1}^k \alpha_i} \right] \gamma \delta \end{aligned}$$

Where $\gamma = \frac{\alpha_1}{x_1} = \dots = \frac{\alpha_M}{x_M}$.

Proof of Theorem A1: Statements for (B)-(D) are identical, if we just vary k from 2 to M . So it suffices to show (A) and then (C). The proof will use a Taylor expansion of f .

First, we observe that $\frac{\partial f}{\partial x_i} = \frac{\alpha_i}{x_i} f(\vec{x})$.

For a finite investment δ allocated according to \vec{p} , the first order Taylor approximation of f around gives the change in utility:

$$\begin{aligned}
\Delta U &= \nabla f(\vec{x}_0) \cdot \nabla \vec{x} \\
&= \left\langle \frac{\alpha_1}{x_1} f(\vec{x}_0), \frac{\alpha_2}{x_2} f(\vec{x}_0), \dots, \frac{\alpha_M}{x_M} f(\vec{x}_0) \right\rangle \cdot \langle \delta p_1, \delta p_2, \dots, \delta p_M \rangle \\
&= \delta f(\vec{x}_0) \sum_{i=1}^M \frac{\alpha_i}{x_{i,0}} p_i
\end{aligned}$$

Our goal is to choose \vec{p} that maximizes this express. The first part, $\delta f(\vec{x}_0)$ is independent of \vec{p} . So we need only maximize the dot product:

$$\left\langle \frac{\alpha_1}{x_{1,0}}, \frac{\alpha_2}{x_{2,0}}, \dots, \frac{\alpha_M}{x_{M,0}} \right\rangle \cdot \vec{p} = \sum_{i=1}^M \frac{\alpha_i}{x_{i,0}} p_i$$

Subject to the linear constraint $\sum_i p_i = 1$. This is a linear function in each p_i , subject to a linear constraint. So it hits its maximum on the boundary of the M -simplex given by $\sum_i p_i = 1$. Since $\frac{\alpha_1}{x_{1,0}}$ is strictly larger than any of the other coefficients $\frac{\alpha_i}{x_{i,0}}$, the maximum can be found by setting $p_1 = 1$ and all the other values to 0. In other words, invest completely in factor x_1 . This shows (A).

To show (C), we first note that if $\frac{\alpha_1}{x_{1,0}} = \frac{\alpha_2}{x_{2,0}} = \dots = \frac{\alpha_k}{x_{k,0}} > \frac{\alpha_{k+1}}{x_{k+1,0}}$, then the first k terms of the first-order Taylor expansion are equivalent:

$$\begin{aligned}
\Delta U &\sim \delta f(\vec{x}_0) \left[\frac{\alpha_1}{x_{1,0}} p_1 + \frac{\alpha_2}{x_{2,0}} p_2 + \dots + \frac{\alpha_M}{x_{M,0}} p_M \right] \\
&= \delta f(\vec{x}_0) \left[\frac{\alpha_1}{x_{1,0}} p_1 + \frac{\alpha_1}{x_{1,0}} p_2 + \dots + \frac{\alpha_1}{x_{1,0}} p_k + \frac{\alpha_{k+1}}{x_{k+1,0}} p_{k+1} + \dots + \frac{\alpha_M}{x_{M,0}} p_M \right] \\
&= \delta f(\vec{x}_0) \left[\frac{\alpha_1}{x_{1,0}} (p_1 + p_2 + \dots + p_k) + \frac{\alpha_{k+1}}{x_{k+1,0}} p_{k+1} + \dots + \frac{\alpha_M}{x_{M,0}} p_M \right]
\end{aligned}$$

The coefficient of the part in blue is larger than the other coefficients by assumption. A similar argument to the one above shows that any investment restricted to the first k factors is optimal. As long as $p_1 + p_2 + \dots + p_k = 1$ and $p_j = 0$ for $j > k$, this expression is optimized.

So in the case of (C), there are an infinite number of allocations that optimize the first-order approximation of utility. This makes sense. The slope in each of the first k directions is the same at \vec{x}_0 . So the tangent hyperplane at this point will not give us sufficient information. We need to examine the second-order Taylor series to figure out how to allocate our resources optimally.

$$\Delta U = \delta \vec{p} \cdot \nabla f(\vec{x}_0) + \frac{1}{2} \delta^2 \vec{p}^\top \left[\frac{\partial^2 f}{\partial \vec{x}^2} \right]_{\vec{x}=\vec{x}_0} \vec{p}$$

Here $\frac{\partial^2 f}{\partial \vec{x}^2}$ is the Hessian matrix of f which has an ij entry equivalent to the second derivative $\frac{\partial^2 f}{\partial x_i \partial x_j}$. These second derivatives are given by:

$$\frac{\partial^2 f}{\partial x_i^2} = \frac{\alpha_i(\alpha_i - 1)}{x_i^2} f(\vec{x})$$

$$\frac{\partial^2 f}{\partial x_i \partial x_j} = \frac{\alpha_i \alpha_j}{x_i x_j} f(\vec{x})$$

When $i \neq j$.

Since the first-order term is equivalent for all allocations in (only) the first k variables, we will assume from here out that $p_j = 0$ for $j > k$ and focus on the second-order term.

$$\begin{aligned} \frac{1}{2} \delta^2 \vec{p}^\top \frac{\partial^2 f}{\partial \vec{x}^2} \vec{p} &= \frac{1}{2} \delta^2 \sum_{i=1}^k \sum_{j=1}^k p_i p_j \frac{\partial^2 f}{\partial x_i \partial x_j} \\ &= \frac{1}{2} \delta^2 f(\vec{x}) \sum_{i=1}^k p_i \left[\frac{\alpha_i(\alpha_i - 1)}{x_i^2} p_i + \sum_{j \neq i}^k \frac{\alpha_i \alpha_j}{x_i x_j} p_j \right] \end{aligned}$$

I will be omitting the little 0's from here out. The reader should remember that everything happens at a specific point. We want to choose the p_i that optimize this expression subject to $\sum_i p_i = 1$. This is again a linear optimization problem. In this case, we will find the expression on the interior of the simplex. This becomes a Lagrange multiplier problem with a Lagrangian.

$$\mathcal{L} = \frac{1}{2} \delta^2 \sum_{i=1}^k \sum_{j=1}^k p_i p_j \frac{\partial^2 f}{\partial x_i \partial x_j} + \lambda \left(\sum_i p_i - 1 \right)$$

We can then use the method of Lagrange multipliers to find \vec{p} .

$$\begin{aligned} 0 = \frac{\partial \mathcal{L}}{\partial p_n} &= \frac{1}{2} \delta^2 \left[\frac{\partial}{\partial p_n} \left(p_n^2 \frac{\partial^2 f}{\partial x_n^2} \right) + 2 \frac{\partial}{\partial p_n} \sum_{j \neq n}^k p_n p_j \frac{\partial^2 f}{\partial x_n \partial x_j} \right] + \lambda \\ &= \frac{1}{2} \delta^2 \left[2 p_n \frac{\partial^2 f}{\partial x_n^2} + 2 \sum_{j \neq n}^k p_j \frac{\partial^2 f}{\partial x_n \partial x_j} \right] + \lambda \end{aligned}$$

$$\begin{aligned}
&= \delta^2 \left[p_n \frac{\alpha_n(\alpha_n - 1)}{x_n^2} f(\vec{x}) + \sum_{j \neq n}^k p_j \frac{\alpha_n \alpha_j}{x_n x_j} f(\vec{x}) \right] + \lambda \\
&= \delta^2 f(\vec{x}) \left[p_n \frac{\alpha_n^2}{x_n^2} - p_n \frac{\alpha_n}{x_n^2} + \sum_{j \neq n}^k p_j \frac{\alpha_n \alpha_j}{x_n x_j} \right] + \lambda
\end{aligned}$$

The first term is equivalent to the expression inside the sum for $j = n$. So we can combine them.

$$= \delta^2 f(\vec{x}) \left[-p_n \frac{\alpha_n}{x_n^2} + \sum_{j=1}^k p_j \frac{\alpha_n \alpha_j}{x_n x_j} \right] + \lambda$$

Remember that $\frac{\alpha_j}{x_j} = \frac{\alpha_n}{x_n} = \frac{\alpha_1}{x_1}$ as long as $j, n \leq k$. So we can rewrite the part inside the sum, and

reduce using the restriction that $\sum_{i=1}^k p_i = 1$:

$$\begin{aligned}
\sum_{j=1}^k p_j \frac{\alpha_n \alpha_j}{x_n x_j} &= \sum_{j=1}^k p_j \frac{\alpha_n \alpha_n}{x_n x_n} \\
&= \frac{\alpha_n^2}{x_n^2} \sum_{j=1}^k p_j \\
&= \frac{\alpha_n^2}{x_n^2} (1) = \frac{\alpha_n^2}{x_n^2} = \frac{\alpha_1^2}{x_1^2}
\end{aligned}$$

So our Lagrange multiplier problem reduces to the set of equations (one for each n):

$$0 = \delta^2 f(\vec{x}) \left[-p_n \frac{\alpha_n}{x_n^2} + \frac{\alpha_1^2}{x_1^2} \right] + \lambda$$

Simplifying, we get:

$$0 = \delta^2 f(\vec{x}) \left[-\frac{p_n \alpha_n}{x_n x_n} + \frac{\alpha_1^2}{x_1^2} \right] + \lambda$$

$$0 = \delta^2 f(\vec{x}) \left[-\frac{p_n \alpha_1}{x_n x_1} + \frac{\alpha_1^2}{x_1^2} \right] + \lambda$$

$$0 = \delta^2 f(\vec{x}) \frac{\alpha_1}{x_1} \left[-\frac{p_n}{x_n} + \frac{\alpha_1}{x_1} \right] + \lambda$$

$$-\frac{\lambda}{\delta^2 f(\vec{x}) \alpha_1} \frac{x_1}{\alpha_1} = -\frac{p_n}{x_n} + \frac{\alpha_1}{x_1}$$

$$\frac{p_n}{x_n} = \frac{\alpha_1}{x_1} + \frac{\lambda}{\delta^2 f(\vec{x}) \alpha_1} \frac{x_1}{\alpha_1}$$

The expression on the left, while complicated, is independent of n . And the expression above must hold for all $n \leq k$. So we conclude that the optimal solution arises when:

$$\frac{p_1}{x_1} = \frac{p_2}{x_2} = \dots = \frac{p_k}{x_k}$$

This is very similar to our assumption that:

$$\frac{\alpha_1}{x_{1,0}} = \frac{\alpha_2}{x_{2,0}} = \dots = \frac{\alpha_k}{x_{k,0}}$$

Any vector \vec{p} which is a linear multiple of $\langle \alpha_1, \alpha_2, \dots, \alpha_k, 0, \dots, 0 \rangle$ will optimize ΔU . Since $\sum_i p_i = 1$, our optimal investment is then:

$$\vec{p} = \frac{1}{\sum_{i=1}^k \alpha_i} \langle \alpha_1, \alpha_2, \dots, \alpha_k, 0, \dots, 0 \rangle$$

This proves (C), as well as (B) and (D).

Proof of Theorem A2: We can calculate the second-order Taylor approximations of the improved utility from the single-factor intervention $\vec{p}_1 = \langle 1, 0, 0, \dots, \rangle$. (Again, omitting the subscript 0's.)

$$\begin{aligned} \Delta U_1 &= \delta \vec{p}_1 \cdot \nabla f + \frac{1}{2} \delta^2 \vec{p}_1^\top \frac{\partial^2 f}{\partial \vec{x}^2} \vec{p}_1 \\ &= \delta \frac{\partial f}{\partial x_1} + \frac{1}{2} \delta^2 \frac{\partial^2 f}{\partial x_1^2} \end{aligned}$$

Substituting the explicit expression for the derivatives gives us:

$$\begin{aligned} \Delta U_1 &= \left[\delta \frac{\alpha_1}{x_1} + \frac{1}{2} \delta^2 \frac{\alpha_1(\alpha_1 - 1)}{x_1^2} \right] f(\vec{x}) \\ &= \left[\delta \frac{\alpha_1}{x_1} + \frac{1}{2} \delta^2 \frac{\alpha_1^2}{x_1^2} - \frac{1}{2} \delta^2 \frac{\alpha_1}{x_1^2} \right] f(\vec{x}) \end{aligned}$$

We can do the same for the optimal multi-factor intervention $\vec{p}_m = \frac{1}{\bar{\alpha}} \langle \alpha_1, \alpha_2, \dots, \alpha_k, 0, \dots, 0 \rangle$,

where we set $\bar{\alpha} = \sum_{i=1}^k \alpha_i$:

$$\begin{aligned} \Delta U_m &= \delta \vec{p}_m \cdot \nabla f + \frac{1}{2} \delta^2 \vec{p}_m^\top \frac{\partial^2 f}{\partial \vec{x}^2} \vec{p}_m \\ &= \frac{\delta}{\bar{\alpha}} \sum_{i=1}^k \alpha_i \frac{\partial f}{\partial x_i} + \frac{\delta^2}{2\bar{\alpha}^2} \sum_{i=1}^k \sum_{j=1}^k \alpha_i \frac{\partial^2 f}{\partial x_i \partial x_j} \alpha_j \end{aligned}$$

$$= \frac{\delta}{\bar{\alpha}} f(\vec{x}) \sum_{i=1}^k \alpha_i \frac{\alpha_i}{x_i} + \frac{\delta^2 f(\vec{x})}{2\bar{\alpha}^2} \left[\sum_{i=1}^k \sum_{j=1}^k \alpha_i \frac{\alpha_i}{x_i} \frac{\alpha_j}{x_j} \alpha_j - \sum_{i=1}^k \alpha_i^2 \frac{\alpha_i}{x_i^2} \right]$$

The extra sum inside the brackets accounts for the “-1” when we differentiate with respect to the same variable twice: $\frac{\partial^2 f}{\partial x_i^2} = \frac{\alpha_i(\alpha_i-1)}{x_i^2} f(\vec{x})$. We can, again, substitute $\frac{\alpha_i}{x_i} = \frac{\alpha_1}{x_1}$ to get:

$$\begin{aligned} \Delta U_m &= \frac{\delta f(\vec{x})}{\bar{\alpha}} \sum_{i=1}^k \alpha_i \frac{\alpha_1}{x_1} + \frac{\delta^2 f(\vec{x})}{2\bar{\alpha}^2} \left[\sum_{i=1}^k \sum_{j=1}^k \alpha_i \left(\frac{\alpha_1}{x_1} \right)^2 \alpha_j - \sum_{i=1}^k \alpha_i \frac{\alpha_1^2}{x_1^2} \right] \\ &= \frac{\delta f(\vec{x})}{\bar{\alpha}} \frac{\alpha_1}{x_1} \sum_{i=1}^k \alpha_i + \frac{\delta^2 f(\vec{x})}{2\bar{\alpha}^2} \left(\frac{\alpha_1}{x_1} \right)^2 \left[\sum_{i=1}^k \sum_{j=1}^k \alpha_i \alpha_j - \sum_{i=1}^k \alpha_i \right] \end{aligned}$$

Note that $\sum_{i=1}^k \alpha_i = \bar{\alpha}$. So we can remove the sums.

$$\begin{aligned} &= \frac{\delta f(\vec{x})}{\bar{\alpha}} \frac{\alpha_1}{x_1} \bar{\alpha} + \frac{\delta^2 f(\vec{x})}{2\bar{\alpha}^2} \left(\frac{\alpha_1}{x_1} \right)^2 [\bar{\alpha}^2 - \bar{\alpha}] \\ \Delta U_m &= \left[\delta \frac{\alpha_1}{x_1} + \frac{1}{2} \delta^2 \left(\frac{\alpha_1}{x_1} \right)^2 - \frac{1}{2} \delta^2 \frac{\alpha_1}{x_1^2} \left(\frac{\alpha_1}{\bar{\alpha}} \right) \right] f(\vec{x}) \end{aligned}$$

Combining expressions to find the absolute loss:

$$\begin{aligned} \Delta U_m - \Delta U_1 &= \left[\delta \frac{\alpha_1}{x_1} + \frac{1}{2} \delta^2 \left(\frac{\alpha_1}{x_1} \right)^2 - \frac{1}{2} \delta^2 \frac{\alpha_1}{x_1^2} \left(\frac{\alpha_1}{\bar{\alpha}} \right) \right] f(\vec{x}) - \left[\delta \frac{\alpha_1}{x_1} + \frac{1}{2} \delta^2 \frac{\alpha_1^2}{x_1^2} - \frac{1}{2} \delta^2 \frac{\alpha_1}{x_1^2} \right] f(\vec{x}) \\ &= \left[+\frac{1}{2} \delta^2 \frac{\alpha_1}{x_1^2} - \frac{1}{2} \delta^2 \frac{\alpha_1}{x_1^2} \left(\frac{\alpha_1}{\bar{\alpha}} \right) \right] f(\vec{x}) \\ &= \frac{1}{2} \delta^2 f(\vec{x}) \frac{\alpha_1}{x_1^2} \left(1 - \frac{\alpha_1}{\bar{\alpha}} \right) \end{aligned}$$

Pulling out a α_1 gives us:

$$\Delta U_m - \Delta U_1 = \frac{1}{2} \delta^2 f(\vec{x}) \frac{\alpha_1^2}{x_1^2} \left(\frac{1}{\alpha_1} - \frac{1}{\bar{\alpha}} \right)$$

This is the absolute loss. The relative loss gives us a unitless measure of inefficiency: how much we lost as a percentage of how much we could have gained. Since δ is assumed to be small, we use the first order approximation for the denominator, which gives the expected formula.

$$\frac{\Delta U_m - \Delta U_1}{\Delta U_m} \sim \frac{\frac{1}{2} \delta^2 f(\vec{x}) \frac{\alpha_1^2}{x_1^2} \left(\frac{1}{\alpha_1} - \frac{1}{\bar{\alpha}} \right)}{\delta \frac{\alpha_1}{x_1} f(\vec{x})}$$

$$= \frac{1}{2} \delta \frac{\alpha_1}{x_1} \left[\frac{1}{\alpha_1} - \frac{1}{\bar{\alpha}} \right]$$

Chapter 6

General Discussion

This dissertation examined how many causally related, and often hidden, factors can influence inequality. It explored how our simple explanations for systemic inequality may be counterproductive, and the role that technology can play in magnifying those simple stories or in influencing growth and inequality.

In Chapter Two, we showed that students' ability to be successful in school can be measured by viewing success from their point of view. We developed original statistical methods and used administrative data to show universal trends in community college student capital. We showed that student ability to be successful had the same distribution across 140 cohorts of students. This suggests that student capital behaves like a limited resource.

In Chapter Three, we examined how the term *white privilege* could increase polarization, decrease civility, and erode support for progressively racial policies. Individuals who would otherwise be supportive of those policies were more likely to remove their support or avoid the conversation altogether, due to different ways the language was received.

A major contribution of this dissertation is its transdisciplinarity. An examination of literature across disciplines points out connections between the different factors behind success and well-being. Chapter Four drew on this connectedness to develop a model for multi-factor cumulative advantage using a causal network. I showed that a number of concepts fall out of such a model, including the systemic interest multiplier a , as well as vectors describing both the relative causal strength and long-term growth ratios of various factors. In addition, I showed that distributions of outcome variables generated through this process were eventually stable and tended to be qualitatively similar to distribution in student capital, income (Tao et al., 2019), and depression (Tomitaka et al., 2015).

In Chapter Five, I presented a discipline-independent framework for systemic inequality and connected it with the research literature. I used models to show how the causes and indicators of

success need not be the same, that opportunities to use skills and resources more effectively can increase both inequality and growth, and that adding a new causal factor can either increase or decrease inequality. I then explored the effects of interventions on systemic inequality, showing that the most effective interventions (a) target disadvantaged individuals and (b) aim to improve multiple factors simultaneously.

6.1 Implications for Policymakers and Interventionists

This work was motivated by the hard work of people trying to improve well-being for others. The primary implications are therefore targeted towards creating more effective policies and interventions. Systemic inequality is generated by a large number of interacting and often hidden factors. Practitioners can recognize systemic inequality when (a) individuals collectively refer to a large number of problems causing their disadvantage or (b) repeated, single-factor interventions have not recently made a significant impact.

- **Start with data-gathering.** Change-makers should first try to understand the variety of factors influencing success. All humans have blind spots, so there should be a purposeful and inclusive effort to understand the problem from multiple angles. Approaches might include qualitative surveys of both advantaged and disadvantaged populations, expert opinions from practitioners working on the problem, established research, and large scale administrative data. Rather than rejecting any particular factor(s), the process should be inclusive of potential causes of inequality, while also recognizing that culturally-generated narratives may make some factors seem more impactful than they really are.
- **Target disadvantaged individuals.** Since advantaged individuals have more resources, they are often better positioned to take advantage of opportunities. This can both increase inequality and increase average growth, as the people at the top benefit more than those at the bottom. So, when possible, interventions and policies should target those people/groups that need the most help. Even for interventions that cannot or should not be restricted to disadvantaged groups, special attention can be paid to the least well-off. This can be done through targeted outreach, financial support, or other mechanisms. This special attention is particularly reasonable for profitable technological innovations, where some economic growth can be diverted toward equity.

- **Recognize that advantage/disadvantage is usually a continuum, not a set of discrete groups.** Since systemic inequality is characterized by many causes, it typically does not split people cleanly into groups. It may be practically useful for interventions or policies to target discrete groups. However, these groups will typically contain significant heterogeneity.
- **Focus on building up resources and capabilities within disadvantaged individuals and communities.** Interventions that aim to reduce barriers might help people progress through those particular sticking points. However, it will not make it easier for them to bypass the next barrier that arises. Instead, interventionists should focus on giving people the skills, traits, and resources that they need to effectively surpass a wide variety of barriers on their own.
- **Target many factors simultaneously.** If inequality has a single cause, it makes sense to target that cause. However, if inequality is more systemic, then multi-factor interventions will typically be more effective.
- **Attempt, where possible, to target factors that have been previously been ignored or misunderstood.** If a problem has been around for a long time, previous attempts at improvement have likely targeted – and improved – the obvious, easy-to-impact causes of inequality. So effective interventions might focus on the less obvious factors. If these idiosyncratic factors are hard to identify or influence, then policies and interventions might focus on providing resources and skills (such as money, social networks, or communication skills) that individuals can use to shore up the areas that *they* decide need improvement.

6.2 Implications for Researchers

This dissertation also suggests a variety of directions for future research. Since inequality tends to become more systemic over time (Petrunyk & Pfeifer, 2016; Sacerdote, 2005; Scholte et al., 2015), a deeper examination of these forces is warranted. One approach, following work on the fundamental causes of health inequalities literature (Phelan & Link, 2005), would be to more thoroughly explore how well-being and increase inequality over time. One potential mechanism comes from interventions: It is quite likely that the interventions that get implemented attempt to improve factors which are (i) easier to understand and implement, (ii) more broadly effective,

and (iii) more isolatable. This leaves those factors which are (i) harder to understand and influence, (ii) affect only a small number of people, and (iii) more interrelated with other factors. This process by itself could explain the rise of systemic inequality. Another possible mechanism is the multivariate Matthew effect outlined in Chapter Four: Individuals with resources can often use those resources to improve their well-being in other ways, leading to the “multivariately rich” getting richer in many ways. This process by itself could also explain increasing systemic inequality. Of course, these processes could be occurring simultaneously at multiple levels. Communities with more resources may have made more significant investments in the well-being of their citizens at the same time as those wealthier citizens are reinvesting their own money into various forms of well-being. The data supports this, since as the GDP of a city¹⁰ increases, so do (i) all quintiles of its citizens income and (ii) inequality in the city (Heinrich Mora et al., 2021). This differential arises because the richest citizens’ incomes grow much more quickly than the poorest (Heinrich Mora et al., 2021). Future studies could tease out the relative effects of these different drivers of systemicity in inequality.

The concept of *student capital* as a coarse-grained combination of many traits, resources, and skills has not been thoroughly studied in the research literature. Our method of estimating student capital was specific to community colleges, since it relied on the rather large dropout rates of these students. Dropout rates throughout most of K-12 and in selective colleges tend to be much lower, which hides the information our method used. Another method that might work for K-12 or selective colleges can be found in (A. B. Mitnitski et al., 2001). They operationalized frailty by adding up the number of health “deficits” that an individual had. A similar approach could be used to measure student capital at a given age. For instance, student capital in second grade might be the sum of an inventory of multiple, observed binary outcomes: whether a child could add two digit numbers, communicate respectfully with classmates, write a complete sentence with correct punctuation, sit still for a certain length of time, etc. This metric would emphasize breadth, rather than the depth that is sought for in standard learning assessments. Alternatively, a deficit approach could be used, where each item corresponds to an area where

¹⁰ More precisely, the population of a city scales with inequality and income quintiles. However, the population of a city is strongly related to its GDP (Bettencourt, 2013). So population size is a measure of the economic strength of a city.

the child is falling short of grade-level expectations. I hypothesize that such a measure could be a strong predictor of long-term student outcomes.

A different study, inspired by Chapter Three, could examine the effects of racial salience on long-term behavior. Our study found a short-term intersectional effect, where whites showed different responses to the term *white privilege*. Research on social media has found that exposure to views or people who are different from us can reinforce our pre-existing views (Bail et al., 2018) or make us more open (Pettigrew & Tropp, 2006). The difference in effects is due to a variety of factors (Cottrell & Neuberg, 2005; Munger, 2017). It would be interesting to examine the heterogeneity in the long-term effects on different groups of whites of repeated exposure to terms like *white privilege*.

In addition, more work could be done to better adapt the methods in this dissertation to data. The analytical model in Chapter Two was able to estimate the “average student capital” in a cohort of students. This is a point estimate of a population parameter which could be used much like graduation or transfer rates in college administration. However, to create more robust analyses, it would be useful to know the standard error and bias in generating this point estimate. One promising approach is to use Fisher Information to give a measure of confidence in the estimate of average student capital. In addition, we model required waiting five years after a student’s initial enrollment, which gave students the opportunity to complete college. This long wait is a problem for colleges (as it is with waiting for graduation rates), because colleges need to make decisions on a quicker time scale. Future work could find a method to estimate the average student capital in a cohort after only a few terms, along with providing time-dependent standard error measurements. These analyses would allow institutional researchers to know whether the average student capital was a practically useful metric of student success.

The model in Chapter Four could be more effectively adapted to data. The model itself starts with the assumption that things like income, student capital, and frailty arise from an accumulation process. So an important goal of that work would be to fit the model to, say, income distributions and infer information about the systemic interest rate a or other components of the model. Currently, it is possible to fit the model to a dataset representing an accumulated distribution, as long as one knows the *cumulants* of the data distribution. However, estimating the cumulants of a population using sample data is notoriously difficult (Chan et al., 2020). Cumulants give information about the tail of the distribution, which is often where things like

sampling bias seep in. Further work could develop a method using some variant of maximum likelihood estimation to fit the model to data. Such a method could give insight into how inequality becomes systemic and leads to the stable distributions that we see in society. Another study could examine socioeconomic status (SES) in the context of the accumulation model in Chapter Four. The model in Chapter Four led to long-term trends that were one-dimensional in nature. The matrix A contained the information about how individuals reinvest their skills, traits, and resources based on their ability and goals. The causal network then led to a set of weights, where each factor's weight corresponded to its size relative to the other factors in the long-term. This created a latent variable which explained a larger and larger proportion of the variance in the data as time progressed. SES is just such a variable. Researchers have used dimension reduction techniques to examine SES as a latent variable explaining various forms of well-being (Cowan et al., 2012). I hypothesize that the principal components found through SES represents the set of indicator weights in my model. A variety of tests could be used to test this hypothesis. One approach could examine whether SES explains an increasing amount of variance as individuals age, as my model predicts. Another approach could examine whether different demographic groups have different first components which represent different values and opportunity to reinvest resources. Those components might then correspond to existing literature about skills required or valued in certain environments (Anderson, 2000; Uskul et al., 2019; Yosso, 2005).

6.3 Conclusion

Millions of people have made contributions toward improving societal well-being. This has led to tremendous progress in a multitude of dimensions. Future progress needs to account for the fact that well-being is multifaceted. When generations have already invested considerable effort in making the world a better place, the old ideas may become less effective. Simple stories that effectively described the challenges people faced may not apply in our interconnected world. Technology can be, and has been, a boon. But it can also create opportunities for the well-off to benefit more, or reinforce simple stories about complex problems thereby exacerbating inequality. This dissertation has taken an interdisciplinary approach to shed light on how many intertwined and often hidden factors influence inequality, and explored some ways to analyze and influence systemic inequality.

References

- Abrams, D., & Hogg, M. (2010). Social identity and self-categorization. In J. F. Dovidio, M. Hewstone, P. Glick, & V. M. Esses (Eds.), *The SAGE Handbook of Prejudice, Stereotyping and Discrimination* (pp. 179–193). SAGE Publications.
- Adams, S. [ScottAdamsSays]. (2020, June 25). *I've decided to reclaim some of the free speech I have been losing lately. Here's my free speech opinion for today [Tweet]*.
<https://twitter.com/scottadamssays/status/1276130129970753538>
- Alegría, M., NeMoyer, A., Falgàs Bagué, I., Wang, Y., & Alvarez, K. (2018). Social Determinants of Mental Health: Where We Are and Where We Need to Go. *Current Psychiatry Reports*, 20(11), 95. <https://doi.org/10.1007/s11920-018-0969-9>
- Algan, Y., Beasley, E., Côté, S., Park, J., Tremblay, R. E., & Vitaro, F. (2022). The Impact of Childhood Social Skills and Self-Control Training on Economic and Noneconomic Outcomes: Evidence from a Randomized Experiment Using Administrative Data. *American Economic Review*, 112(8), 2553–2579. <https://doi.org/10.1257/aer.20200224>
- Allison, P. D., Long, J. S., & Krauze, T. A. D. K. (2018). *Cumulative Advantage and Inequality in Science Author (s): Paul D. Allison, J. Scott Long and Tad K. Krauze Published by: American Sociological Association Stable URL: http://www.jstor.org/stable/2095162 American Sociological Association is collab.* 47(5), 615–625.
- Amis, J. M., Mair, J., & Munir, K. A. (2020). The Organizational Reproduction of Inequality. *Academy of Management Annals*, 14(1), 195–230. <https://doi.org/10.5465/annals.2017.0033>
- Anderson, E. (2000). *Code of the Street: Decency, Violence, and the Moral Life of the Inner City*. W.W. Norton & Company, Inc.
- Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. B. F., Lee, J., Mann, M., Merhout, F., & Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Science*, 115(37), 9216–9221. <https://doi.org/https://doi.org/10.1073/pnas.1804840115>

- Bandourian, R., McDonald, J. B., & Turley, R. S. (2002). A Comparison of Parametric Models of Income Distribution Across Countries and Over Time. In *SSRN Electronic Journal* (No. 305; Luxembourg Income Study Working Paper). <https://doi.org/10.2139/ssrn.324900>
- Banton, M. (2015). Racism. In *The Wiley Blackwell Encyclopedia of Race, Ethnicity, and Nationalism* (pp. 1–8). John Wiley & Sons, Ltd.
<https://doi.org/https://doi.org/10.1002/9781118663202.wberen539>
- Bazarova, N. N., Choi, Y. H., Sosik, V. S., Cosley, D., & Whitlock, J. (2015). Social sharing of emotions on Facebook: Channel differences, satisfaction, and replies. *Proceedings of the 2015 ACM International Conference on Computer-Supported Cooperative Work and Social Computing*, 154–164. <https://doi.org/10.1145/2675133.2675297>
- Becker, G. S. (1994). *Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education* (3rd ed.). The University of Chicago Press.
<https://www.nber.org/books/beck94-1>
- Bettencourt, L. M. A. (2013). The Origins of Scaling in Cities. *Science*, *340*, 1438–1441.
<https://doi.org/10.1126/science.167.3924.1461>
- Bin Naeem, S., & Kamel Boulos, M. N. (2021). COVID-19 Misinformation Online and Health Literacy: A Brief Overview. *International Journal of Environmental Research and Public Health*, *18*(15), 8091. <https://doi.org/10.3390/ijerph18158091>
- Blau, P. M., & Duncan, O. D. (1967). *The American Occupational Structure*. Wiley.
- Bloome, D. (2015). Income Inequality and Intergenerational Income Mobility in the United States. *Social Forces*, *93*(3), 1047–1080. <https://doi.org/10.1016/j.jhsa.2009.09.008>. Validity
- Boaler, J. (1998). Open and Closed Mathematics : Student Experiences and Understandings. *Journal for Research in Mathematics Education*, *29*(1), 41–62.
<https://doi.org/10.2307/749717>
- Bogg, T., & Roberts, B. W. (2004). Conscientiousness and Health-Related Behaviors: A Meta-Analysis of the Leading Behavioral Contributors to Mortality. *Psychological Bulletin*, *130*(6), 887–919. <https://doi.org/10.1037/0033-2909.130.6.887>
- Böhm, S. (2013). Behavior and expectations of mobile job seekers. *Proceedings of the 2013 Annual Conference on Computers and People Research*, 105–110.
<https://doi.org/10.1145/2487294.2487316>
- Boschloo, R. D. (1970). Raised conditional level of significance for the 2×2 -table when testing

- the equality of two probabilities. *Statistica Neerlandica*, 24(1), 1–9.
<https://doi.org/10.1111/j.1467-9574.1970.tb00104.x>
- Boshara, R., Emmons, W. R., & Noeth, B. J. (2015). *The Demographics of Wealth: Education and wealth* (Issue 2). <https://www.stlouisfed.org/household-financial-stability/the-demographics-of-wealth/essay-2-the-role-of-education>
- Bowles, S., & Gintis, H. (2002). The Inheritance of Inequality. *Journal of Economic Perspectives*, 16(3), 3–30. <https://doi.org/10.1257/089533002760278686>
- Boyce, W. T. (2019). *The Orchid and the Dandelion*. Alfred A. Knopf.
- Boydstun, A. E., Card, D., Gross, J. H., Resnik, P., & Smith, N. A. (2014). Tracking the development of media frames within and across policy issues. In *Working paper*.
<http://www.cs.cmu.edu/~nasmith/papers/boydstun+card+gross+resnik+smith.apsa14.pdf>
- Boyle, E. A., Li, Y. I., & Pritchard, J. K. (2017). An Expanded View of Complex Traits: From Polygenic to Omnigenic. *Cell*, 169(7), 1177–1186.
<https://doi.org/10.1016/j.cell.2017.05.038>
- Brady, W. J., McLoughlin, K., Doan, T. N., & Crockett, M. J. (2021). How social learning amplifies moral outrage expression in online social networks. *Science Advances*, 7(33), 1–15. <https://doi.org/10.1126/sciadv.abe5641>
- Branscombe, N. R., Schmitt, M. T., & Schiffhauer, K. (2007). Racial attitudes in response to thoughts of white privilege. *European Journal of Social Psychology*, 37(2), 203–215.
<https://doi.org/10.1002/ejsp.348>
- Braveman, P. A., Arkin, E., Proctor, D., Kauh, T., & Holm, N. (2022). Systemic And Structural Racism: Definitions, Examples, Health Damages, And Approaches To Dismantling. *Health Affairs*, 41(2), 171–178. <https://doi.org/10.1377/hlthaff.2021.01394>
- Breivik, H., Collett, B., Ventafridda, V., Cohen, R., & Gallacher, D. (2006). Survey of chronic pain in Europe: Prevalence, impact on daily life, and treatment. *European Journal of Pain*, 10(4), 287–287. <https://doi.org/10.1016/j.ejpain.2005.06.009>
- Brimi, H. M. (2011). Reliability of grading high school work in English. *Practical Assessment, Research and Evaluation*, 16(17), 1–12. <https://doi.org/10.1086/435971>
- Brown, M. (2020, August). How to explain white privilege in term simple enough for a child. *Parents*. <https://www.parents.com/kids/responsibility/racism/how-to-explain-white-privilege-in-term-simple-enough-for-a-child/>

- Brown, R. P., Wohl, M. J. A., & Exline, J. J. (2008). Taking up offenses: Secondhand forgiveness and group identification. *Personality and Social Psychology Bulletin*, *34*(10), 1406–1419. <https://doi.org/10.1177/0146167208321538>
- Bruch, E. E., & Newman, M. E. J. (2019). Structure of online dating markets in U.S. cities. *Sociological Science*, *6*, 219–234. <https://doi.org/10.15195/V6.A9>
- Budak, C., Garrett, R. K., Resnick, P., & Kamin, J. (2017). Threading is Sticky: How Threaded Conversations Promote Comment System User Retention. *Proc. ACM Hum.-Comput. Interact. Article*, *1*(20). <https://doi.org/10.1145/3134662>
- Bunker, C. J., & Varnum, M. E. W. (2021). How strong is the association between social media use and false consensus? *Computers in Human Behavior*, *125*(December), 106947. <https://doi.org/10.1016/j.chb.2021.106947>
- Butler, E. A., Egloff, B., Wilhelm, F. H., Smith, N. C., Erickson, E. A., & Gross, J. J. (2003). The Social Consequences of Expressive Suppression. *Emotion*, *3*(1), 48–67. <https://doi.org/10.1037/1528-3542.3.1.48>
- Calhoun, P. (2019). *Exact: Unconditional Exact Test* (R package version 2.0). <https://cran.r-project.org/package=Exact>
- Campbell, T. A., & Campbell, D. E. (1997). Faculty/student mentor program: Effects on academic performance and retention. *Research in Higher Education*, *38*(6), 727–742. <https://doi.org/https://doi.org/10.1023/A:1024911904627>
- Case, K. A., & Rios, D. (2017). Educational interventions to raise awareness of white privilege. *Journal on Excellence in College Teaching*, *28*(1), 137–156.
- Castex, G., & Dechter, E. K. (2014). The changing roles of education and ability in wage determination. *Journal of Labor Economics*, *32*(4), 685–710. <https://doi.org/10.1086/676018>
- Cech, E. A., & Waidzunas, T. J. (2021). Systemic inequalities for LGBTQ professionals in STEM. *Science Advances*, *7*(3). <https://doi.org/10.1126/sciadv.abe0933>
- Chan, L. H., Chen, K., Li, C., Wong, C. W., & Yau, C. Y. (2020). On higher-order moment and cumulant estimation. *Journal of Statistical Computation and Simulation*, *90*(4), 747–771. <https://doi.org/10.1080/00949655.2019.1700987>
- Chater, N., & Loewenstein, G. (2016). The under-appreciated drive for sense-making. *Journal of Economic Behavior & Organization*, *126*, 137–154.

<https://doi.org/10.1016/j.jebo.2015.10.016>

- Chetty, R., Jackson, M. O., Kuchler, T., Stroebe, J., Hendren, N., Fluegge, R. B., Gong, S., Gonzalez, F., Grondin, A., Jacob, M., Johnston, D., Koenen, M., Laguna-Muggenburg, E., Mudekereza, F., Rutter, T., Thor, N., Townsend, W., Zhang, R., Bailey, M., ... Wernerfelt, N. (2022). Social capital I: measurement and associations with economic mobility. *Nature*, *608*(7921), 108–121. <https://doi.org/10.1038/s41586-022-04996-4>
- Clotfelter, C. T., Ladd, H. F., Muschkin, C. G., & Vigdor, J. L. (2013). Success in community college: Do institutions differ? *Research in Higher Education*, *54*, 805–824. <http://link.springer.com/article/10.1007/s11162-013-9295-6>
- Coleman, J. S. (1988). Social Capital in the Creation of Human Capital. *American Journal of Sociology*, *94*(Supplement), S95–S120.
- Cook, B. L., Trinh, N.-H., Li, Z., Hou, S. S.-Y., & Progovac, A. M. (2017). Trends in Racial-Ethnic Disparities in Access to Mental Health Care, 2004–2012. *Psychiatric Services*, *68*(1), 9–16. <https://doi.org/10.1176/appi.ps.201500453>
- Coppock, A. (2019). Generalizing from Survey Experiments Conducted on Mechanical Turk: A Replication Approach. *Political Science Research and Methods*, *7*(3), 613–628. <https://doi.org/10.1017/psrm.2018.10>
- Cortright, J. (2006). Making sense of clusters: regional competitiveness and economic development. In *The Brookings Institution*. https://www.brookings.edu/wp-content/uploads/2016/06/20060313_Clusters.pdf
- Cottrell, C. A., & Neuberg, S. L. (2005). Different emotional reactions to different groups: A sociofunctional threat-based approach to “prejudice.” *Journal of Personality and Social Psychology*, *88*(5), 770–789. <https://doi.org/10.1037/0022-3514.88.5.770>
- Cowan, C. D., Hauser, R. M., Kominski, R. A., Levin, H. M., Lucas, S. R., Morgan, S. L., Spencer, M. B., & Chaptman, C. (2012). *Improving the Measurement of Socioeconomic Status for the NAEP: A Theoretical Foundation*. <https://linkinghub.elsevier.com/retrieve/pii/B9780407729032500101>
- Cox, R. D. (2015). “You’ve Got to Learn the Rules”: A Classroom-Level Look at Low Pass Rates in Developmental Math. *Community College Review*, *43*(3), 264–286. <https://doi.org/10.1177/0091552115576566>
- Cuban, L. (2013). *Inside the Black Box of Classroom Practice: Change Without Reform in*

- American Education*. Harvard Education Press.
- Dagum, C. (1977). A new model of personal income distribution: Specification and estimation. *Economie Appliquee*, 30, 413–437.
- Danescu-Niculescu-Mizil, C., West, R., Jurafsky, D., & Potts, C. (2013). No Country for Old Members : User Lifecycle and Linguistic Change in Online Communities. *Proceedings of the 22nd International Conference on World Wide Web*, 307–317.
<https://doi.org/10.1145/2488388.2488416>
- Dannefer, D. (2003). Cumulative Advantage/Disadvantage and the Life Course: Cross-Fertilizing Age and Social Science Theory. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 58(6), S327–S337.
<https://doi.org/10.1093/geronb/58.6.S327>
- Davies, J. B., & Shorrocks, A. F. (2000). The distribution of wealth. In *Handbook of Income Distribution* (Vol. 1, pp. 605–675). [https://doi.org/https://doi.org/10.1016/S1574-0056\(00\)80014-7](https://doi.org/https://doi.org/10.1016/S1574-0056(00)80014-7)
- Davis, J., Wolff, H.-G., Forret, M. L., & Sullivan, S. E. (2020). Networking via LinkedIn: An examination of usage and career benefits. *Journal of Vocational Behavior*, 118(January), 103396. <https://doi.org/10.1016/j.jvb.2020.103396>
- De Martino, A., & De Martino, D. (2018). An introduction to the maximum entropy approach and its application to inference problems in biology. *Heliyon*, 4.
<https://doi.org/10.1016/j.heliyon.2018.e00596>
- Deming, D. J. (2017). The growing importance of social skills in the labor market. *Quarterly Journal of Economics*, 132(4), 1593–1640.
<https://doi.org/https://doi.org/10.1093/qje/qjx022>
- Diehl, T., Weeks, B. E., & Gil de Zúñiga, H. (2016). Political persuasion on social media: Tracing direct and indirect effects of news use and social interaction. *New Media and Society*, 18(9), 1875–1895. <https://doi.org/10.1177/1461444815616224>
- Dillon, E. W., & Smith, J. A. (2017). Determinants of the Match between Student Ability and College Quality. *Journal of Labor Economics*, 35(1), 45–66. <https://doi.org/10.1086/687523>
- DiPrete, T. A., & Eirich, G. M. (2006). Cumulative Advantage as a Mechanism for Inequality: A Review of Theoretical and Empirical Developments. *Annual Review of Sociology*, 32(1), 271–297. <https://doi.org/10.1146/annurev.soc.32.061604.123127>

- Doane, W. (2003). Rethinking Whiteness Studies. In A. W. Doane & E. Bonilla-Silva (Eds.), *White Out: The Continuing Significance of Racism* (pp. 3–18). Routledge.
- Doosje, B., Branscombe, N. R., Spears, R., & Manstead, A. S. R. (1998). Guilty by association: When one's group has a negative history. *Journal of Personality and Social Psychology*, 75(4), 872–886. <https://doi.org/10.1037/0022-3514.75.4.872>
- Drăgulescu, A., & Yakovenko, V. M. (2000). Statistical mechanics of money. *European Physical Journal B*, 17, 723–729.
- Drăgulescu, A., & Yakovenko, V. M. (2001). Exponential and power-law probability distributions of wealth and income in the United Kingdom and the United States. *Physica A: Statistical Mechanics and Its Applications*, 299(1–2), 213–221. [https://doi.org/10.1016/S0378-4371\(01\)00298-9](https://doi.org/10.1016/S0378-4371(01)00298-9)
- Duncombe, C. (2019). The Politics of Twitter: Emotions and the Power of Social Media. *International Political Sociology*, 13(4), 409–429. <https://doi.org/10.1093/ips/olz013>
- Dunning, T. (2008). Improving Causal Inference. *Political Research Quarterly*, 61(2), 282–293. <https://doi.org/10.1177/1065912907306470>
- Easley, D., & Kleinberg, J. (2010). *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press.
- Edin, P.-A., Fredriksson, P., Nybom, M., & Öckert, B. (2022). The Rising Return to Noncognitive Skill. *American Economic Journal: Applied Economics*, 14(2), 78–100. <https://doi.org/10.1257/app.20190199>
- El-Sayed, A. M., & Galea, S. (Eds.). (2017). *Systems Science and Population Health*. Oxford University Press.
- Farrell, S. G., Mitnitski, A. B., Theou, O., Rockwood, K., & Rutenberg, A. D. (2018). Probing the network structure of health deficits in human aging. *Physical Review E*, 98(3), 032302. <https://doi.org/10.1103/PhysRevE.98.032302>
- Farrington, C. A., Roderick, M., & Allensworth, E. (2012). *Teaching Adolescents to Become Learners: The Role of Noncognitive Factors in Shaping School Performance--A Critical Literature Review*. University of Chicago Consortium on Chicago School Research. [https://consortium.uchicago.edu/sites/default/files/publications/Noncognitive Report.pdf](https://consortium.uchicago.edu/sites/default/files/publications/Noncognitive%20Report.pdf)
- Fe, H., & Sanfelice, V. (2022). How bad is crime for business? Evidence from consumer behavior. *Journal of Urban Economics*, 129(February 2021), 103448.

<https://doi.org/10.1016/j.jue.2022.103448>

- Fowler, F. J. J., & Cosenza, C. (2008). Writing effective questions. In E. D. de Leeuw, J. J. Jox, & D. A. Dillman (Eds.), *International Handbook of Survey Methodology* (pp. 136–160). Psychology Press.
- Fox, J., & Holt, L. F. (2018). Fear of Isolation and Perceived Affordances: The Spiral of Silence on Social Networking Sites Regarding Police Discrimination. *Mass Communication and Society*, 21(5), 533–554. <https://doi.org/10.1080/15205436.2018.1442480>
- Freedman, D. H. (2010). Why scientific studies are so often wrong: The streetlight effect. *Discover*. <https://www.discovermagazine.com/the-sciences/why-scientific-studies-are-so-often-wrong-the-streetlight-effect>
- Fried, L. P., Hadley, E. C., Walston, J. D., Newman, A. B., Guralnik, J. M., Studenski, S., Harris, T. B., Ershler, W. B., & Ferrucci, L. (2005). From Bedside to Bench: Research Agenda for Frailty. *Science of Aging Knowledge Environment*, 2005(31). <https://doi.org/10.1126/sageke.2005.31.pe24>
- Fristedt, B., & Gray, L. (1997). *A Modern Approach to Probability Theory*. Springer.
- Fritzell, J., & Henz, U. (2021). Household income dynamics: mobility out of and into low income over the life-course. In J. Jonsson & C. Mills (Eds.), *Cradle to Grave: Life-Course Change in Modern Sweden* (pp. 184–210). Sociology Press. <https://doi.org/10.4324/9781315074610-9>
- Fry, L., Santos, Y. E., & Zhang, Y. (2015). Health information use: Preliminary results from a systematic review. *Proceedings of the Association for Information Science and Technology*, 52(1), 1–3. <https://doi.org/10.1002/pra2.2015.1450520100112>
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic Decision Making. *Annual Review of Psychology*, 62(1), 451–482. <https://doi.org/10.1146/annurev-psych-120709-145346>
- Gillespie, T. (2018). *Custodians of the internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press.
- Goff, P. A., Steele, C. M., & Davies, P. G. (2008). The Space Between Us: Stereotype Threat and Distance in Interracial Contexts. *Journal of Personality and Social Psychology*, 94(1), 91–107. <https://doi.org/10.1037/0022-3514.94.1.91>
- Goldin, C., & Katz, L. F. (2008). *The Race between Education and Technology*. Belknap Press.
- Goldrick-Rab, S. (2018). Addressing Community College Completion Rates by Securing

- Students' Basic Needs. *New Directions for Community Colleges*, 2018(184), 7–16.
<https://doi.org/10.1002/cc.20323>
- Goldrick-Rab, S., Richardson, J., & Hernandez, A. (2017). *Hungry and Homeless in College: Results From a National Study of Basic Needs Insecurity in Higher Education*.
<http://wihopelab.com/publications/Hungry-and-Homeless-in-College-Report.pdf>
- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and Conservatives Rely on Different Sets of Moral Foundations. *Journal of Personality and Social Psychology*, 96(5), 1029–1046.
<https://doi.org/10.1037/a0015141>
- Granovetter, M. (1983). The Strength of Weak Ties: A Network Theory Revisited. *Sociological Theory*, 1, 201–233. <https://doi.org/10.2307/202051>
- Grubb, W. N. (2001). *From black box to pandora's box: Evaluating remedial/developmental education*.
- Hamermesh, D. S. (2011). *Beauty Pays*. Princeton University Press.
- Hargittai, E. (2018). Potential Biases in Big Data: Omitted Voices on Social Media. *Social Science Computer Review*, 1–15. <https://doi.org/10.1177/0894439318788322>
- Harte, J. (2011). *Maximum Entropy and Ecology: A Theory of Abundance, Distribution, and Energetics*. Oxford University Press.
- Hawley, P. H. (2002). Social dominance and prosocial and coercive strategies of resource control in preschoolers. *International Journal of Behavioral Development*, 26(2), 167–176.
<https://doi.org/10.1080/01650250042000726>
- Heckman, J. J., Stixrud, J., & Urzua, S. (2006). The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior. *Journal of Labor Economics*, 24(3), 411–482. <https://doi.org/10.1086/504455>
- Heikamp, T., Phalet, K., Van Laar, C., & Verschueren, K. (2020). To belong or not to belong: Protecting minority engagement in the face of discrimination. *International Journal of Psychology*, 55(5), 779–788. <https://doi.org/10.1002/ijop.12706>
- Heinrich Mora, E., Heine, C., Jackson, J. J., West, G. B., Yang, V. C., & Kempes, C. P. (2021). Scaling of urban income inequality in the USA. *Journal of The Royal Society Interface*, 18(181). <https://doi.org/10.1098/rsif.2021.0223>
- Henle, J. M., Horton, N. J., & Jakus, S. J. (2008). Modelling Inequality with a Single Parameter. In D. Chotikapanich (Ed.), *Modeling Income Distributions and Lorenz Curves* (pp. 255–

- 269). Springer. https://doi.org/10.1007/978-0-387-72796-7_14
- Herrnstein, R. J., & Murray, C. (1994). *The Bell Curve: Intelligence and class structure in American life*. The Free Press.
- Hidalgo, C. A. (2015). *Why Information Grows*. Basic Books.
- Hidalgo, C. A. (2021). Economic complexity theory and applications. *Nature Reviews Physics*, 3(2), 92–113. <https://doi.org/10.1038/s42254-020-00275-1>
- Hout, M. (2015). A Summary of What We Know about Social Mobility. *Annals of the American Academy of Political and Social Science*, 657(1), 27–36. <https://doi.org/10.1177/0002716214547174>
- Hovmand, P. S. (2014). *Community Based System Dynamics: Lessons from The Field*.
- Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N., & Westwood, S. J. (2019). The Origins and Consequences of Affective Polarization in the United States. *Annual Review of Political Science*, 22(1), 129–146. <https://doi.org/10.1146/annurev-polisci-051117-073034>
- Jacoby, W. G. (2000). Issue Framing and Public Opinion on Government Spending. *American Journal of Political Science*, 44(4), 750–767. <https://www.jstor.org/stable/2669279>
- Jenkins, D., Ellwein, T., Wachen, J., Kerrigan, M. R., & Cho, S.-W. (2009). *Achieving the Dream Colleges in Pennsylvania and Washington State: Early Progress toward Building a Culture of Evidence* (Issue March). <https://academiccommons.columbia.edu/doi/10.7916/D8445JKS>
- Joensen, J. S., & Nielsen, H. S. (2018). Spillovers in education choice. *Journal of Public Economics*, 157(November 2015), 158–183. <https://doi.org/10.1016/j.jpubeco.2017.10.006>
- John, N. A., & Dvir-Gvirzman, S. (2015). I Don't Like You Any More: Facebook Unfriending by Israelis During the Israel-Gaza Conflict of 2014. *Journal of Communication*, 65(6), 953–974. <https://doi.org/10.1111/jcom.12188>
- Johnson, C. R., & Tarazaga, P. (2004). On Matrices with Perron–Frobenius Properties and Some Negative Entries. *Positivity*, 8(4), 327–338. <https://doi.org/10.1007/s11117-003-3881-3>
- Johnson, J., & Rochkind, J. (2009). *With their whole lives ahead of them: Myths and realities about why so many students fail to finish college*. <http://www.publicagenda.org/files/theirwholelivesaheadofthem.pdf>
- Jones, L., & Tertilt, M. (2008). An Economic History of Fertility in the U.S.: 1826-1960. In *Frontiers of Family Economics* (pp. 165–230). <https://doi.org/10.3386/w12796>

- Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Keith, V. M., & Herring, C. (1991). Skin tone and stratification in the black community. *American Journal of Sociology*, 97(3), 760–778.
- Killewald, A., Pfeffer, F. T., & Schachner, J. N. (2017). Wealth Inequality and Accumulation. *Annual Review of Sociology*, 43(1), 379–404. <https://doi.org/10.1146/annurev-soc-060116-053331>
- King, D. L., & Delfabbro, P. H. (2020). Video game addiction. In *Adolescent Addiction* (Second Edi, pp. 185–213). Elsevier. <https://doi.org/10.1016/B978-0-12-818626-8.00007-4>
- Kirilenko, A. P., & Stepchenkova, S. (2016). Inter-coder agreement in one-to-many classification: Fuzzy kappa. *PLoS ONE*, 11(3), 1–14. <https://doi.org/10.1371/journal.pone.0149787>
- Koch, P. (2021). Economic complexity and growth: Can value-added exports better explain the link? *Economics Letters*, 198, 109682. <https://doi.org/10.1016/j.econlet.2020.109682>
- Kolenovic, Z., Linderman, D., & Karp, M. M. (2013). Improving Student Outcomes via Comprehensive Supports : Three-Year Outcomes From CUNY’s Accelerated Study in Associate Programs (ASAP). *Community College Review*, 41(3), 271–291. <https://doi.org/10.1177/0091552113503709>
- Kraus, M. W., & Park, J. W. (2017). The structural dynamics of social class. *Current Opinion in Psychology*, 18, 55–60. <https://doi.org/10.1016/j.copsyc.2017.07.029>
- Krüger, S., & Degel, A. (2022). The fall of the East German trickster: A comparative cultural analysis of East and West German jokes before and after reunification. *Psychoanalysis, Culture & Society*, 27(2–3), 181–200. <https://doi.org/10.1057/s41282-022-00280-6>
- Kullback, S., & Liebler, R. A. (1951). On information and sufficiency. *Annals of Mathematical Statistics*, 22(1), 79–86. <https://doi.org/doi:10.1214/aoms/1177729694>
- Lakner, C., & Milanovic, B. (2013). Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession. In *Policy Research Working Paper* (No. 6719). <https://doi.org/10.1596/29118>
- Lam, D. (1997). Chapter 18 Demographic variables and income inequality. In *Handbook of Population and Family Economics* (Vol. 1, Issue PART B, pp. 1015–1059). [https://doi.org/10.1016/S1574-003X\(97\)80010-4](https://doi.org/10.1016/S1574-003X(97)80010-4)
- Landale, N. S., Oropesa, R. S., & Noah, A. J. (2017). Experiencing discrimination in Los

- Angeles: Latinos at the intersection of legal status and socioeconomic status. *Social Science Research*, 67, 34–48. <https://doi.org/10.1016/j.ssresearch.2017.05.003>
- Lareau, A. (2011). *Unequal Childhoods* (2nd ed.). University of California Press.
- Lashitew, A. A., Branzei, O., & van Tulder, R. (2023). Community Inclusion under Systemic Inequality: How For-Profit Businesses Pursue Social Purpose. *Journal of Management Studies*. <https://doi.org/10.1111/joms.12907>
- Lee, N. Y., & Kim, Y. (2014). The spiral of silence and journalists' outspokenness on Twitter. *Asian Journal of Communication*, 24(3), 262–278. <https://doi.org/10.1080/01292986.2014.885536>
- Lemon, J. (2006). Plotrix: a package in the red light district of R. *R-News*, 6(4), 8–12.
- Lerman, K., Yan, X., & Wu, X.-Z. (2016). The “majority illusion” in social networks. *PLoS ONE*, 11(2). <https://doi.org/10.1371/journal.pone.0147617>
- Levin, H. M., & Garcia, E. (2013). *Benefit-cost analysis of accelerated study in associate programs (ASAP) of the City University of New York (CUNY)* (Issue May). https://www1.nyc.gov/assets/opportunity/pdf/Levin_ASAP_Benefit_Cost_Report_FINAL_05212013.pdf
- Levin, H. M., & García, E. (2018). Accelerating Community College Graduation Rates: A Benefit–Cost Analysis. *The Journal of Higher Education*, 89(1), 1–27. <https://doi.org/10.1080/00221546.2017.1313087>
- Levy, B. L., Phillips, N. E., & Sampson, R. J. (2020). Triple Disadvantage: Neighborhood Networks of Everyday Urban Mobility and Violence in U.S. Cities. *American Sociological Review*, 85(6), 925–956. <https://doi.org/10.1177/0003122420972323>
- Lichtenstein, A. H., Appel, L. J., Brands, M., Carnethon, M., Daniels, S., Franch, H. A., Franklin, B., Kris-Etherton, P., Harris, W. S., Howard, B., Karanja, N., Lefevre, M., Rudel, L., Sacks, F., Van Horn, L., Winston, M., & Wylie-Rosett, J. (2006). Summary of American Heart Association Diet and Lifestyle Recommendations Revision 2006. *Arteriosclerosis, Thrombosis, and Vascular Biology*, 26(10), 2186–2191. <https://doi.org/10.1161/01.ATV.0000238352.25222.5e>
- Lieberman, M. D., Hariri, A., Jarcho, J. M., Eisenberger, N. I., & Bookheimer, S. Y. (2005). An fMRI investigation of race-related amygdala activity in African-American and Caucasian-American individuals. *Nature Neuroscience*, 8(6), 720–722. <https://doi.org/10.1038/nn1465>

- Long, M. C., Conger, D., & Iatarola, P. (2012). Effects of High School Course-Taking on Secondary and Postsecondary Success. *American Educational Research Journal*, 49(2), 285–322. <https://doi.org/10.3102/0002831211431952>
- Lorenc, T., Petticrew, M., Welch, V., & Tugwell, P. (2013). What types of interventions generate inequalities? Evidence from systematic reviews. *Journal of Epidemiology and Community Health*, 67(2), 190–193. <https://doi.org/10.1136/jech-2012-201257>
- Lowery, B. S., Knowles, E. D., & Unzueta, M. M. (2007). Framing Inequity Safely: Whites’ Motivated Perceptions of Racial Privilege. *Personality and Social Psychology Bulletin*, 33(9), 1237–1250. <https://doi.org/10.1177/0146167207303016>
- Lowery, B. S., & Wout, D. A. (2010). When inequality matters: The effect of inequality frames on academic engagement. *Journal of Personality and Social Psychology*, 98(6), 956–966. <https://doi.org/10.1037/a0017926>
- Lukacs, E. (1970). *Characteristic Functions* (2nd ed.). Hafner.
- Mackie, D. M., Devos, T., & Smith, E. R. (2000). Intergroup emotions: Explaining offensive action tendencies in an intergroup context. *Journal of Personality and Social Psychology*, 79(4), 602–616. <https://doi.org/10.1037/0022-3514.79.4.602>
- Mackie, D. M., & Smith, E. R. (2015). Intergroup emotions. In M. Mikulincer & P. R. Shaver (Eds.), *APA Handbook of Personality and Social Psychology* (Vol. 2, pp. 263–293). American Psychological Association. <https://doi.org/10.1037/14342-010>
- Malik, K. (2020, June 14). “White privilege” is a distraction, leaving racism and power untouched. *The Guardian*. <https://www.theguardian.com/commentisfree/2020/jun/14/white-privilege-is-a-lazy-distraction-leaving-racism-and-power-untouched>
- Manduca, R. A. (2019). The Contribution of National Income Inequality to Regional Economic Divergence. *Social Forces*, 1–27. <https://doi.org/10.1093/sf/soz013>
- Markus, H. R., & Hamedani, M. G. (2019). People are Culturally-Shaped Shapers: The Psychological Science of Culture and Culture Change. In D. Cohen & S. Kitayama (Eds.), *The Handbook of Cultural Psychology* (2nd ed., pp. 11–52). Guilford.
- Matthes, J., Knoll, J., & von Sikorski, C. (2018). The “Spiral of Silence” Revisited: A Meta-Analysis on the Relationship Between Perceptions of Opinion Support and Political Opinion Expression. *Communication Research*, 45(1), 3–33. <https://doi.org/10.1177/0093650217745429>

- Matysiak, A., & Steinmetz, S. (2008). Finding Their Way? Female Employment Patterns in West Germany, East Germany, and Poland. *European Sociological Review*, 24(3), 331–345. <https://doi.org/10.1093/esr/jcn007>
- Mayhew, M. J., Rockenbach, A. N., Bowman, N. A., Seifert, T. A., Wolniak, G. C., Pascarella, E. T., & Terenzini, P. T. (2016). *How College Affects Students, Volume 3: 21st Century evidence that higher education works*. Jossey-Bass.
- McFarland, J., Cui, J., Rathbun, A., & Holmes, J. (2018). *Trends in High School Dropout and Completion Rates in the United States: 2018*.
- McGregor, S. C. (2019). Social media as public opinion: How journalists use social media to represent public opinion. *Journalism*, 20(8), 1070–1086. <https://doi.org/10.1177/1464884919845458>
- McGregor, S. C. (2020). “Taking the Temperature of the Room”: How Political Campaigns Use Social Media to Understand and Represent Public Opinion. *Public Opinion Quarterly*, 84(S1), 236–256. <https://doi.org/10.1093/poq/nfaa012>
- McIntosh, P. (1990). White privilege: Unpacking the invisible knapsack. *Independent School, Winter*, 31–36.
- McLanahan, S., & Jacobsen, W. (2015). Diverging destinies revisited. In P. R. Amato, A. Booth, S. M. McHale, & J. Van Hook (Eds.), *Families in an Era of Increasing Inequality: National Symposium on Family Issues* (pp. 3–23). Springer. <https://doi.org/10.1007/978-3-319-08308-7>
- McLanahan, S., & Percheski, C. (2008). Family Structure and the Reproduction of Inequalities. *Annual Review of Sociology*, 34(1), 257–276. <https://doi.org/10.1146/annurev.soc.34.040507.134549>
- Meier, H. E., & Mutz, M. (2016). Sport-Related National Pride in East and West Germany, 1992-2008. *SAGE Open*, 6(3), 215824401666589. <https://doi.org/10.1177/2158244016665893>
- Milanovic, B., Lindert, P. H., & Williamson, J. G. (2011). Pre-industrial inequality. *The Economic Journal*, 121(551), 255–272. <https://doi.org/10.1111/j.1468-0297.2010.02403.x>.
- Miller, C., & Weiss, M. J. (2022). Increasing Community College Graduation Rates: A Synthesis of Findings on the ASAP Model From Six Colleges Across Two States. *Educational Evaluation and Policy Analysis*, 44(2), 210–233.

<https://doi.org/10.3102/01623737211036726>

Mitnitski, A. B., Mogilner, A. J., & Rockwood, K. (2001). Accumulation of Deficits as a Proxy Measure of Aging. *The Scientific World JOURNAL*, 1, 323–336.

<https://doi.org/10.1100/tsw.2001.58>

Mitnitski, A., Bao, L., & Rockwood, K. (2006). Going from bad to worse: A stochastic model of transitions in deficit accumulation, in relation to mortality. *Mechanisms of Ageing and Development*, 127(5), 490–493. <https://doi.org/10.1016/j.mad.2006.01.007>

Mittos, A., Zannettou, S., Blackburn, J., & De Cristofaro, E. (2020). “And we will fight for our race!” A measurement study of genetic testing conversations on Reddit and 4chan. *Proceedings of the 14th International AAAI Conference on Web and Social Media, ICWSM 2020*, 14(1), 452–463. <https://ojs.aaai.org/index.php/ICWSM/article/view/7314>

Montroll, E. W. (1981). On the entropy function in sociotechnical systems. *Proceedings of the National Academy of Sciences of the United States of America*, 78(12), 7839–7843.

<https://doi.org/10.1073/pnas.78.12.7839>

Moore, W. S., Lardner, E., Malnarich, G., & Davis, M. (2013). *Rethinking Pre-college Math: Pedagogy, professional responsibility, and student success* (Issue February).

<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.666.5369&rep=rep1&type=pdf>

Morris, A. S., Silk, J. S., Steinberg, L., Myers, S. S., & Robinson, L. R. (2007). The role of the family context in the development of emotion regulation. *Social Development*, 16(2), 361–388. <https://doi.org/10.1111/j.1467-9507.2007.00389.x>

Muchnik, L., Aral, S., & Taylor, S. J. (2013). Social influence bias: a randomized experiment. *Science (New York, N.Y.)*, 341(6146), 647–651. <https://doi.org/10.1126/science.1240466>

Mulder, M. B., Bowles, S., Hertz, T., Bell, A., Beise, J., Clark, G., Fazzio, L., Gurven, M., Hill, K., Hooper, P. L., Irons, W., Kaplan, H., Leonetti, D., Low, B., Marlowe, F., McElreath, R., Naidu, S., Nolin, D., Piraino, P., ... Wiessner, P. (2009). Supplemental Materials: Intergenerational wealth transmission and the dynamics of inequality in small-scale societies. *Science*, 326(5953), 682–688. <https://doi.org/10.1126/science.1178336>

Munger, K. (2017). Tweetment Effects on the Tweeted: Experimentally Reducing Racist Harassment. *Political Behavior*, 39(3), 629–649. [https://doi.org/10.1007/s11109-016-9373-](https://doi.org/10.1007/s11109-016-9373-5)

5

Myers, S. A., Sharma, A., Gupta, P., & Lin, J. (2014). Information Network or Social Network ?

- The Structure of the Twitter Follow Graph. *Proceedings of the 23rd International Conference on the World Wide Web*, 493–498.
- National Research Council. (2001). *Adding It Up: Helping Children Learn Mathematics* (J. Kilpatrick, J. Swafford, & B. Findell (Eds.)). National Academies Press.
<https://doi.org/10.17226/9822>
- Neubaum, G., & Krämer, N. C. (2017). Monitoring the Opinion of the Crowd: Psychological Mechanisms Underlying Public Opinion Perceptions on Social Media. *Media Psychology*, 20(3), 502–531. <https://doi.org/10.1080/15213269.2016.1211539>
- Newman, M. (2018). *Networks* (2nd ed.). Oxford University Press.
- Nirei, M., & Souma, W. (2007). A two factor model of income distribution dynamics. *Review of Income and Wealth*, 53(3), 440–459. <https://doi.org/doi.org/10.1111/j.1475-4991.2007.00242.x>
- Noelle-Neumann, E. (1974). The spiral of silence: A theory of public opinion. *Journal of Communication, Spring*, 43–51.
- Nolan, J. P. (2020). *Univariate Stable Distributions: Models for Heavy Tailed Data*. Springer.
- Olfson, M., Druss, B. G., & Marcus, S. C. (2015). Trends in Mental Health Care among Children and Adolescents. *New England Journal of Medicine*, 372(21), 2029–2038.
<https://doi.org/10.1056/NEJMsa1413512>
- Page, S. E., & Zelner, J. (2020). Population Health as a Complex Adaptive System of Systems. In *Complex Systems and Population Health* (pp. 33–44). Oxford University Press.
<https://doi.org/10.1093/oso/9780190880743.003.0003>
- Pareto, V. (1897). *Cours d'Economie Politique*.
- Pascarella, E. T., & Terenzini, P. T. (2005). *How College Affects Students, Volume 2: A third decade of research*. Jossey-Bass.
https://edocs.uis.edu/Departments/LIS/Course_Pages/LIS301/papers/How_college_effects_students_534-545.pdf
- Petrunyk, I., & Pfeifer, C. (2016). Life Satisfaction in Germany After Reunification: Additional Insights on the Pattern of Convergence. *Jahrbücher Für Nationalökonomie Und Statistik*, 236(2), 217–239. <https://doi.org/10.1515/jbnst-2015-1010>
- Pettigrew, T. F., & Tropp, L. R. (2006). A meta-analytic test of intergroup contact theory. *Journal of Personality and Social Psychology*, 90(5), 751–783.

<https://doi.org/10.1037/0022-3514.90.5.751>

- Pew Research Center. (2019a). *In a Politically Polarized Era, Sharp Divides in Both Partisan Coalitions* (Issue December).
- Pew Research Center. (2019b). *Race in America 2019* (Issue April).
- Pfeffer, F. T., & Killewald, A. (2018). Generations of Advantage. Multigenerational Correlations in Family Wealth. *Social Forces*, 96(4), 1411–1442. <https://doi.org/10.1093/sf/sox086>
- Phelan, J. C., & Link, B. G. (2005). Controlling Disease and Creating Disparities: A Fundamental Cause Perspective. *The Journals of Gerontology: Series B*, 60(Special_Issue_2), S27–S33. https://doi.org/10.1093/geronb/60.Special_Issue_2.S27
- Phillips, L. T., & Lowery, B. S. (2018). Herd Invisibility: The Psychology of Racial Privilege. *Current Directions in Psychological Science*, 27(3), 156–162. <https://doi.org/10.1177/0963721417753600>
- Pinckard, K., Baskin, K. K., & Stanford, K. I. (2019). Effects of Exercise to Improve Cardiovascular Health. *Frontiers in Cardiovascular Medicine*, 6(June), 1–12. <https://doi.org/10.3389/fcvm.2019.00069>
- Plaut, V. C., Garnett, F. G., Buffardi, L. E., & Sanchez-Burks, J. (2011). “What about me?” Perceptions of exclusion and Whites’ reactions to multiculturalism. *Journal of Personality and Social Psychology*, 101(2), 337–353. <https://doi.org/10.1037/a0022832>
- Porchea, S. F., Allen, J., Robbins, S., & Phelps, R. P. (2010). Predictors of long-term enrollment and degree outcomes for community college students: Integrating academic, psychosocial, socio-demographic, and situational. *The Journal of Higher Education*, 81(6), 680–708.
- Powell, A. A., Branscombe, N. R., & Schmitt, M. T. (2005). Inequality as ingroup privilege or outgroup disadvantage: The impact of group focus on collective guilt and interracial attitudes. *Personality and Social Psychology Bulletin*, 31(4), 508–521. <https://doi.org/10.1177/0146167204271713>
- Price, D. de S. (1976). A General Theory of Bibliometric and Other Cumulative Advantage Processes. *Journal of the American Society for Information Science*, 27(5), 292–306. <https://doi.org/https://doi.org/10.1002/asi.4630270505>
- Prichard, J., Watters, P., Krone, T., Spiranovic, C., & Cockburn, H. (2015). Social Media Sentiment Analysis: A New Empirical Tool for Assessing Public Opinion on Crime? *Current Issues in Criminal Justice*, 27(2), 217–236.

<https://doi.org/10.1080/10345329.2015.12036042>

- Putnam, R. (2001). Social capital: Measurement and consequences. *Isuma: Canadian Journal of Policy Research*, 2(Spring 2001), 41–51.
- Quarles, C. L., Budak, C., & Resnick, P. (2018). Understanding Human Capital Using Distributions of Student Outcomes. *MIDAS Symposium*.
- Quarles, C. L., Budak, C., & Resnick, P. (2020). The shape of educational inequality. *Science Advances*, 6(29), 1–9. <https://doi.org/10.1126/sciadv.aaz5954>
- Quarles, C. L., & Davis, M. (2017). Is Learning in Developmental Math Associated With Community College Outcomes? *Community College Review*, 45(1), 33–51. <https://doi.org/10.1177/0091552116673711>
- R Core Team. (2018). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. <http://www.r-project.org/>
- Rajadesingan, A., Resnick, P., & Budak, C. (2020). Quick, community-specific learning: How distinctive toxicity norms are maintained in political subreddits. *Proceedings of the International AAAI Conference on Web and Social Media*, 14(1), 557–568.
- Reardon, S. (2011). The widening academic achievement gap between the rich and the poor: New evidence and possible explanations. In R. Murnane & G. Duncan (Eds.), *Whither Opportunity? Rising Inequality and the Uncertain Life Chances of Low-Income Children*. Russell Sage Foundation Press.
- Reardon, S. F., & Bischoff, K. (2011). Income inequality and income segregation. *American Journal of Sociology*, 116(4), 1092–1153. <https://doi.org/10.1086/657114>
- Riedl, D., & Schüßler, G. (2017). The Influence of Doctor-Patient Communication on Health Outcomes: A Systematic Review. *Zeitschrift Für Psychosomatische Medizin Und Psychotherapie*, 63(2), 131–150. <https://doi.org/10.13109/zptm.2017.63.2.131>
- Rockwood, K., & Howlett, S. E. (2019). Age-related deficit accumulation and the diseases of ageing. *Mechanisms of Ageing and Development*, 180(February), 107–116. <https://doi.org/10.1016/j.mad.2019.04.005>
- Rockwood, K., Mogilner, A., & Mitnitski, A. (2004). Changes with age in the distribution of a frailty index. *Mechanisms of Ageing and Development*, 125(7), 517–519. <https://doi.org/10.1016/j.mad.2004.05.003>
- Saad, L. (2020). *U.S. Perceptions of White-Black Relations Sink to New Low*.

- <https://news.gallup.com/poll/318851/perceptions-white-black-relations-sink-new-low.aspx>
- Sacerdote, B. (2005). Slavery and the Intergenerational Transmission of Human Capital. *Review of Economics and Statistics*, 87(2), 217–234. <https://doi.org/10.1162/0034653053970230>
- Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science*, 311(5762), 854–856. <https://doi.org/10.1126/science.1121066>
- Salganik, M. J., Lundberg, I., Kindel, A. T., Ahearn, C. E., Al-Ghoneim, K., Almaatouq, A., Altschul, D. M., Brand, J. E., Carnegie, N. B., Compton, R. J., Datta, D., Davidson, T., Filippova, A., Gilroy, C., Goode, B. J., Jahani, E., Kashyap, R., Kirchner, A., McKay, S., ... McLanahan, S. (2020). Measuring the predictability of life outcomes with a scientific mass collaboration. *Proceedings of the National Academy of Sciences*, 117(15), 8398–8403. <https://doi.org/10.1073/pnas.1915006117>
- Salminen, J., Sengün, S., Corporan, J., Jung, S., & Jansen, B. J. (2020). Topic-driven toxicity: Exploring the relationship between online toxicity and news topics. *PLoS ONE*, 15(2), 1–24. <https://doi.org/https://doi.org/10.1371/journal.pone.0228723>
- Sarabia, J. M., Jordá, V., & Trueba, C. (2017). The Lamé class of Lorenz curves. *Communications in Statistics - Theory and Methods*, 46(11), 5311–5326. <https://doi.org/10.1080/03610926.2013.775306>
- Saunders, M. G., & Voth, G. A. (2013). Coarse-Graining Methods for Computational Biology. *Annual Review of Biophysics*, 42(1), 73–93. <https://doi.org/10.1146/annurev-biophys-083012-130348>
- Sawhill, I. V., & Reeves, R. V. (2016). Modeling Equal Opportunity. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 2(2), 60–97. <https://doi.org/10.7758/rsf.2016.2.2.03>
- Sawhill, I. V., & Karpilow, Q. (2015). *How Much Could We Improve Children's Life Chances by Intervening Early and Often*.
- Sbardella, A., Pugliese, E., & Pietronero, L. (2017). Economic development and wage inequality: A complex system analysis. *PLOS ONE*, 12(9), e0182774. <https://doi.org/10.1371/journal.pone.0182774>
- Scholte, R. S., van den Berg, G. J., & Lindeboom, M. (2015). Long-run effects of gestation during the Dutch Hunger Winter famine on labor market and hospitalization outcomes.

- Journal of Health Economics*, 39, 17–30. <https://doi.org/10.1016/j.jhealeco.2014.10.002>
- Scott-Clayton, J., Crosta, P. M., & Belfield, C. R. (2014). Improving the Targeting of Treatment: Evidence From College Remediation. *Educational Evaluation and Policy Analysis*, 36(3), 371–393. <https://doi.org/10.3102/0162373713517935>
- Seip, E. C., Van Dijk, W. W., & Rotteveel, M. (2014). Anger motivates costly punishment of unfair behavior. *Motivation and Emotion*, 38, 578–588. <https://doi.org/10.1007/s11031-014-9395-4>
- Semega, J., Kollar, M., Shrider, E. A., & Creamer, J. F. (2021). Income and Poverty in the United States: 2019. In *Current Population Reports*. <https://www.census.gov/library/publications/2020/demo/p60-270.html>
- Sewell, W. H., Haller, A. O., & Portes, A. (1969). The Educational and Early Occupational Attainment Process. *American Sociological Review*, 34(1), 82. <https://doi.org/10.2307/2092789>
- Shalley, C. E. (2012). Writing good theory: Issues to consider. *Organizational Psychology Review*, 2(3), 258–264. <https://doi.org/10.1177/2041386611436029>
- Shapiro, D., Dundar, A., Huie, F., Wakhungu, P. K., Bhimdiwala, A., & Wilson, S. E. (2018). *Completing College: A National View of Student Completion Rates - Fall 2012 Cohort (Signature Report No. 16)*. <https://nscresearchcenter.org/wp-content/uploads/SignatureReport16.pdf>
- Shaposhnikov, M. V., Guvatova, Z. G., Zemskaya, N. V., Koval, L. A., Schegoleva, E. V., Gorbunova, A. A., Golubev, D. A., Pakshina, N. R., Ulyasheva, N. S., Solovev, I. A., Bobrovskikh, M. A., Gruntenko, N. E., Menshanov, P. N., Krasnov, G. S., Kudryavseva, A. V., & Moskalev, A. A. (2022). Molecular mechanisms of exceptional lifespan increase of *Drosophila melanogaster* with different genotypes after combinations of pro-longevity interventions. *Communications Biology*, 5(1), 566. <https://doi.org/10.1038/s42003-022-03524-4>
- Shi, J., Song, X., Yu, P., Tang, Z., Mitnitski, A., Fang, X., & Rockwood, K. (2011). Analysis of frailty and survival from late middle age in the Beijing Longitudinal Study of Aging. *BMC Geriatrics*, 11(1), 17. <https://doi.org/10.1186/1471-2318-11-17>
- Singer, M., Bulled, N., & Ostrach, B. (2020). Whither syndemics?: Trends in syndemics research, a review 2015–2019. *Global Public Health*, 15(7), 943–955.

- <https://doi.org/10.1080/17441692.2020.1724317>
- Sirin, S. R. (2005). Socioeconomic Status and Academic Achievement : A Meta-Analytic Review of Research. *Review of Educational Research*, 75(3), 417–453.
<https://doi.org/https://doi.org/10.3102%2F00346543075003417>
- Sleeper, M., Balebako, R., Das, S., McConahy, A. L., Wiese, J., & Cranor, L. F. (2013). The post that wasn't: Exploring self-censorship on Facebook. *Proceedings of the 2013 Conference on Computer Supported Cooperative Work - CSCW '13*, 793–802.
<https://doi.org/10.1145/2441776.2441865>
- Sparks, B., Zidenberg, A. M., & Olver, M. E. (2022). *An Exploratory Study of Incels' Dating App Experiences, Mental Health, and Relational Well-Being*. June.
<https://doi.org/10.13140/RG.2.2.29838.23362>
- Starch, D., & Elliott, E. C. (1912). Reliability of the grading of high-school work in English. *The School Review*, 20(7), 442–457.
- Starch, D., & Elliott, E. C. (1913). Reliability of grading work in mathematics. *The School Review*, 21(4), 254–259. <https://doi.org/10.1086/436086>
- Steck, L. W., Heckert, D. M., & Heckert, D. A. (2003). The salience of racial identity among African-American and white students. *Race and Society*, 6(1), 57–73.
<https://doi.org/10.1016/j.racsoc.2004.09.005>
- Steele, C. M. (2010). *Whistling Vivaldi*. W.W. Norton & Company, Inc.
- Stewart, T. L., Latu, I. M., Branscombe, N. R., Phillips, N. L., & Ted Denney, H. (2012). White Privilege Awareness and Efficacy to Reduce Racial Inequality Improve White Americans' Attitudes Toward African Americans. *Journal of Social Issues*, 68(1), 11–27.
<https://doi.org/10.1111/j.1540-4560.2012.01733.x>
- Stone, C., Trisi, D., Sherman, A., & Beltrán, J. (2020). *A Guide to Statistics on Historical Trends in Income Inequality*. <https://www.cbpp.org/research/poverty-and-inequality/a-guide-to-statistics-on-historical-trends-in-income-inequality>
- Sun, N., Rau, P. P.-L., & Ma, L. (2014). Understanding lurkers in online communities: A literature review. *Computers in Human Behavior*, 38, 110–117.
<https://doi.org/10.1016/j.chb.2014.05.022>
- Taber, C. S., Cann, D., & Kucsova, S. (2009). The motivated processing of political arguments. *Political Behavior*, 31(2), 137–155. <https://doi.org/10.1007/s11109-008-9075-8>

- Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In W. G. Austin & S. Worchel (Eds.), *The Social Psychology of Intergroup Relations*. Brooks-Cole.
- Talaska, C. A., Fiske, S. T., & Chaiken, S. (2008). Legitimizing racial discrimination: Emotions, not beliefs, best predict discrimination in a meta-analysis. *Social Justice Research, 21*, 263–296. <https://doi.org/10.1007/s11211-008-0071-2>
- Tan, J. J. X., Kraus, M. W., Carpenter, N. C., & Adler, N. E. (2020). The association between objective and subjective socioeconomic status and subjective well-being: A meta-analytic review. *Psychological Bulletin, 146*(11), 970–1020. <https://doi.org/10.1037/bul0000258>
- Tao, Y., Wu, X., Zhou, T., Yan, W., Huang, Y., Yu, H., Mondal, B., & Yakovenko, V. M. (2019). Exponential structure of income inequality: evidence from 67 countries. *Journal of Economic Interaction and Coordination, 14*(2), 345–376. <https://doi.org/10.1007/s11403-017-0211-6>
- Tazhitdinova, A. (2022). Increasing Hours Worked: Moonlighting Responses to a Large Tax Reform. *American Economic Journal: Economic Policy, 14*(1), 473–500. <https://doi.org/10.1257/pol.20190786>
- Thill, C. (2015). Listening for policy change: how the voices of disabled people shaped Australia’s National Disability Insurance Scheme. *Disability and Society, 30*(1), 15–28. <https://doi.org/10.1080/09687599.2014.987220>
- Thorson, K., Cotter, K., Medeiros, M., & Pak, C. (2021). Algorithmic inference, political interest, and exposure to news and politics on Facebook. *Information, Communication & Society, 24*(2), 183–200. <https://doi.org/10.1080/1369118X.2019.1642934>
- Tomitaka, S., Kawasaki, Y., & Furukawa, T. (2015). Right tail of the distribution of depressive symptoms is stable and follows an exponential curve during middle adulthood. *PLoS ONE, 10*(1), 1–8. <https://doi.org/10.1371/journal.pone.0114624>
- Tomitaka, S., Kawasaki, Y., Ide, K., Akutagawa, M., Yamada, H., Ono, Y., & Furukawa, T. A. (2018). Distributional patterns of item responses and total scores on the PHQ-9 in the general population: data from the National Health and Nutrition Examination Survey. *BMC Psychiatry, 18*(1), 108. <https://doi.org/10.1186/s12888-018-1696-9>
- Torche, F. (2018). Prenatal Exposure to an Acute Stressor and Children’s Cognitive Outcomes. *Demography, 55*(5), 1611–1639. <https://doi.org/10.1007/s13524-018-0700-9>
- Turchin, P. (2007). *War and Peace and War: The Rise and Fall of Empires*. Pi Press.

- Tyree, A., Freiberg, J. W., Ong, K., Raczynski, D., Shosid, N., & Steeg, D. Ver. (1971). The Dickensian Occupational Structure. *Sociological Inquiry*, 41(1), 95–105.
<https://doi.org/10.1111/j.1475-682X.1971.tb01206.x>
- U.S. Department of Education. (2017). *Digest of Education Statistics*. Digest of Education Statistics. https://nces.ed.gov/programs/digest/d17/tables/dt17_219.57.asp
- UN General Assembly. (1948). *The Universal Declaration of Human Rights*.
- Uskul, A. K., Cross, S. E., Gunsoy, C., & Gul, P. (2019). Cultures of Honor. In S. Kitayama & D. Cohen (Eds.), *Handbook of Cultural Psychology* (pp. 793–821). The Guilford Press.
- Vandenbroeck, P., Goossens, J., & Clemens, M. (2007). Tackling Obesities: Future Choices – Building the Obesity System Map. In *Foresight*. www.foresight.gov.uk
- Veinot, T. C., Mitchell, H., & Ancker, J. S. (2018). Good intentions are not enough: how informatics interventions can worsen inequality. *Journal of the American Medical Informatics Association*, 25(8), 1080–1088. <https://doi.org/10.1093/jamia/ocy052>
- Wang, X., Wickersham, K., Lee, Y., & Chan, H. Y. (2018). Exploring sources and influences of social capital on community college students' first-year success: Does age make a difference? *Teachers College Record*, 120.
- WHO. (2022). *What works to prevent violence against children online?*
<https://www.who.int/publications/i/item/9789240062061>
- Wight, V., Kaushal, N., Waldfogel, J., & Garfinkel, I. (2014). Understanding the link between poverty and food insecurity among children: Does the definition of poverty matter? *Journal of Children and Poverty*, 20(1), 1–20. <https://doi.org/10.1080/10796126.2014.891973>
- Wilson, A. G. (1970). The Use of the Concept of Entropy in System Modelling. *Journal of the Operational Research Society*, 21(2), 247–265.
<https://doi.org/https://doi.org/10.1057/jors.1970.48>
- Wilson, A. S. P., & Urick, A. (2022). An intersectional examination of the opportunity gap in science: A critical quantitative approach to latent class analysis. *Social Science Research*, 102(September 2021), 102645. <https://doi.org/10.1016/j.ssresearch.2021.102645>
- Wood, D. M., Auhl, G., & McCarthy, S. (2019). Accreditation and quality in higher education curriculum design: does the tail wag the dog? *5th International Conference on Higher Education Advances (HEAd'19)*, 783–791. <https://doi.org/10.4995/HEAD19.2019.9365>
- Woods, H. C., & Scott, H. (2016). #Sleepyteens: Social media use in adolescence is associated

- with poor sleep quality, anxiety, depression and low self-esteem. *Journal of Adolescence*, 51, 41–49. <https://doi.org/10.1016/j.adolescence.2016.05.008>
- Yakovenko, V. M., & Rosser, J. B. (2009). Colloquium: Statistical mechanics of money, wealth, and income. *Reviews of Modern Physics*, 81(4), 1703–1725. <https://doi.org/10.1103/RevModPhys.81.1703>
- Yang, H., Liao, Q., & Huang, X. (2008). Minorities remember more: The effect of social identity salience on group-referent memory. *Memory*, 16(8), 910–917. <https://doi.org/10.1080/09658210802360629>
- Yeager, D., & Walton, G. (2011). Social-Psychological Interventions in Education: They're Not Magic. *Review of Education Research*, 81(2), 267–301. <http://www.jstor.org/stable/23014370>
- Yee, T. W. (2018). *VGAM: Vector Generalized Linear and Additive Models* (R package version 1.0-6). <https://cran.r-project.org/package=VGAM>
- Yosso, T. J. (2005). Whose culture has capital? A critical race theory discussion of community cultural wealth. *Race Ethnicity and Education*, 8(1), 69–91. <https://doi.org/10.1080/1361332052000341006>
- Zaniboni, S., Kmicinska, M., Truxillo, D. M., Kahn, K., Paladino, M. P., & Fraccaroli, F. (2019). Will you still hire me when I am over 50? The effects of implicit and explicit age stereotyping on resume evaluations. *European Journal of Work and Organizational Psychology*, 28(4), 453–467. <https://doi.org/10.1080/1359432X.2019.1600506>
- Zhao, F., Li, S., Li, T., & Yu, G. (2019). Does Stereotype Threat Deteriorate Academic Performance of High School Students With Learning Disabilities? The Buffering Role of Psychological Disengagement. *Journal of Learning Disabilities*, 52(4), 312–323. <https://doi.org/10.1177/0022219419849107>
- Zhu, Q., Skoric, M., & Shen, F. (2017). I Shield Myself From Thee: Selective Avoidance on Social Media During Political Protests. *Political Communication*, 34(1), 112–131. <https://doi.org/10.1080/10584609.2016.1222471>
- Zimmerman, F. J., & Anderson, N. W. (2019). Trends in Health Equity in the United States by Race/Ethnicity, Sex, and Income, 1993-2017. *JAMA Network Open*, 2(6), e196386. <https://doi.org/10.1001/jamanetworkopen.2019.6386>