ESSAYS ON LABOR MARKETS, EDUCATION AND MIGRATION IN DEVELOPING COUNTRIES

by

Todd M. Pugatch

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2011

Doctoral Committee:

Professor Brian P. McCall, Chair
Professor David A. Lam
Associate Professor Dean Yang
Professor James A. Levinsohn, Yale University

To Mom, for always putting me first.
And to Dad, who is smiling somewhere.

ACKNOWLEDGEMENTS

First and foremost, I thank my committee – Brian McCall, David Lam, James Levinsohn, and Dean Yang – for invaluable guidance and support. It was particularly a pleasure to co-author essays of this dissertation with James Levinsohn (Chapter 3) and Dean Yang (Chapter 4). The chapter-specific acknowledgements below come equally from them as from me. Kevin Stange deserves special thanks for his patient guidance, particularly on the Chapter 2. Chapters 2 and 3 use the Cape Area Panel Study (CAPS), for which I am grateful to Cally Ardington, Nicola Branson, Jessica Goldberg, and David Lam for oriented me to the data and graciously sharing their programs for formatting it. I acknowledge financial support from the Rackham Pre-Doctoral Fellowship to fund my final year of studies. Most importantly, I thank my mother for unwavering support, my siblings David and Marianne for always encouraging me, and my girlfriend Tina for helping me focus on what's really important.

Acknowledgements specific to each chapter follow.

Chapter 2. I thank Cally Ardington, Emily Beam, Elodie Djemai, Rob Garlick, Jessica Goldberg, Osborne Jackson, Erik Johnson, David Lam, Murray Leibbrandt, James Levinsohn, Brian McCall, Zoë McLaren, Peter Orazem, Francesc Ortega, Christina Plerhoples, David Ratner, Matthew Rutledge, Kevin Stange, Elias Walsh, Martin Wittenberg, Dean Yang and seminar participants at DIWDC, IMPAQ International, Instituto Tecnológico Autónomo de México, International Food Policy

Research Institute, Oregon State University, Population Association of America annual meetings, Seattle University, Stanford University, University of Cape Town, University of Michigan, Vanderbilt University, and the World Bank for helpful comments.

Chapter 3. We thank Justin McCrary for his essential input as a co-author on previous versions of this paper. Brian McCall provided invaluable guidance on the structural estimation techniques used in this paper. David Lam, Murray Leibbrandt, Zoë McLaren, Jeffrey Smith, Kevin Stange, Martin Wittenberg, and seminar participants at the University of Michigan, the University of Cape Town, Centre for European Economic Research (ZEW), NEUDC, 3rd Joint IZA/IFAU Conference on Labor Market Policy Evaluation, and 5th IZA/World Bank Conference on Employment and Development provided useful comments.

Chapter 4. We thank John Bound, Charlie Brown, John DiNardo, Lutz Kilian, Francesc Ortega, Yoonseok Lee, Jeff Smith, Mel Stephens, Steve Trejo, and seminar participants at the University of Michigan for useful comments. Justin Ladner provided excellent research assistance. We also acknowledge financial support from the University of Michigan National Poverty Center.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER I

Introduction

This dissertation consists of three separate essays in the fields of labor and development economics. The first two essays, which examine the school to work transitions of South African youth, are closely related. Both consider how South African youth make school enrollment and labor market decisions in an environment of high grade failure and high unemployment, and how these choices affect their economic circumstances. The third essay analyzes the effect of Mexican immigration on labor market outcomes in the United States, using variation in the Mexican immigrant presence across U.S. states generated by weather shocks in Mexico. A focus on individual decision-making in difficult and uncertain economic environments-whether the young South African choosing to enroll for another year in school or search for a job, or a Mexican peasant choosing to migrate to the United States-binds together this research.

The first essay, "Bumpy Rides: School to Work Transitions in South Africa," examines the schooling choices of South African youth. These youth face extraordinarily high unemployment (47 percent of 15-24 year olds were unemployed in South Africa in 2007) and high rates of grade failure and repetition. Their frequent re-enrollment in school after spells of withdrawal is difficult to reconcile with standard human capital investment models, but becomes sensible in a dynamic context in which they are uncertain about future outcomes and the returns associated with their choices. Specifically, I develop and estimate a dynamic model of school advancement and job search that allows for uncertain outcomes and updating about academic and labor market returns during adolescence. In contrast to prevailing dynamic schooling models that estimate an optimal stopping rule, my model allows for re-enrollment in school after periods of dropout, a common feature of the South African youth experience that has been largely ignored in the literature. As youth learn the outcomes of their schooling and labor market choices, they update their expectations about the relative returns to enrollment versus labor market participation, leading some who dropped out of school to re-enroll.

Estimates of the model's structural parameters confirm the hypothesis that enrollment choices reflect dynamic updating of the relative returns to schooling versus labor market participation. In addition to explaining observed patterns of school enrollment and labor market participation, I quantify the importance of re-enrollment by conducting simulations in which the option to re-enroll before completing high school is restricted. I find that the re-enrollment restriction increases the proportion completing 12 years of schooling by 8 percentage points, as youth reconsider the consequences of early dropout. The results suggest that the option to re-enroll is an important component of the incentives South African youth face when making schooling decisions. This research is the first dynamic, sequential schooling model applied to South African data, and the first in the wider dynamic sequential schooling literature to focus on the importance of re-enrollment after dropout.

In the accompanying second essay, "The Role of Reservation Wages in Youth Unemployment in South Africa: A Structural Approach" (co-authored with James

Levinsohn), we examine the post-schooling labor market outcomes of South African youth in greater detail, using a job search model to quantify the importance of reservation wages in a setting where job offer arrivals and offered wages are uncertain. We find that inclusion of survey data on reservation wages, as opposed to alternative methods more commonly used in the literature, implies a labor market in which job offers are relatively frequent but at wages that tend to be too low to be accepted. Using a novel procedure, we estimate the full distribution of job search costs in the sample, and show that these estimates confirm the model's predictions that those with lower search costs are more likely to remain unemployed and hold out for higher wages. We also present evidence that an employer wage subsidy for hiring unemployed youth would increase accepted wages and decrease the probability of lengthy unemployment spells, even after accounting for how the subsidy would increase reservation wages. To our knowledge, this is the first attempt to apply survey data on reservation wages to a structurally estimated search model for a developing country, and contributes to the debate on the role of reservation wages in South African unemployment.

The final essay, "The Impact of Mexican Immigration on U.S. Natives: Evidence from Migrant Flows Driven by Rainfall Shocks" (co-authored with Dean Yang), examines the effect of immigration driven by another source of uncertainty: weather shocks. We report a new estimate of the effect of immigration on U.S. labor market outcomes that exploits rainfall shocks in migration source areas, a novel source of exogenous variation in migrant inflows. Spatial and temporal variation in rainfall within Mexico generates variation in immigration to the U.S. from specific Mexican locations. Crucially for our identification strategy, there is historical persistence in the U.S. destinations of migrants from particular source regions in Mexico. We predict changes in Mexican immigration in a given U.S. state on the basis of rainfall shocks occurring in Mexican locations that historically sent migrants to that U.S. state. Rainfall shocks at the Mexican state level, operating through established migration channels, are strongly correlated with changes in the Mexican labor force share in an annual panel of U.S. states. Using our instrument to estimate the effect of the Mexican labor force share on native labor market outcomes, we find that a higher Mexican share of the labor force leads to lower wages and higher unemployment for non-Mexicans. Our point estimates are substantially larger in magnitude than previous estimates of the wage impact of immigration.

CHAPTER II

Bumpy Rides: School to Work Transitions in South Africa

2.1 Introduction

More than a decade after the fall of apartheid in South Africa, economic outcomes for previously disenfranchised groups remain bleak, particularly for youth: 52% of 20-24 year olds were unemployed in 2005 (Banerjee, Galiani, Levinsohn, McLaren and Woolard 2008). South Africa's youth unemployment is severe even by the standards of its region, with the lowest employment/population ratio among its neighbors and other large African economies shown in Figure 2.1. Behind South Africa's poor youth labor market outcomes lies an especially slow and bumpy transition from school to work: a remarkable 22% of 20 year-olds were enrolled in grades K-10 in 2001, reflecting frequent grade repetition and re-enrollment; Figure 2.2 shows that only Brazil has a higher rate among a group of comparison countries. Numerous studies have examined the mechanisms generating youth economic outcomes in South Africa, but most rely on static theories that assume students have perfect foresight about future job opportunities when making schooling decisions. Even the small number of dynamic sequential schooling models in the wider literature dismiss re-enrollment in school after spells of withdrawal as a "rare event" (Keane and Wolpin 1997, pp. 487) or ignore it entirely. Yet in South Africa, where lack of labor market opportunities and high rates of school failure make enrollment decisions difficult for many youth, re-enrollment is a key feature of the school to work transition: in my sample, one third re-enroll at some point in their schooling careers, including 20% before completing high school.[1] While static models are incapable of explaining this irregular enrollment behavior, re-enrollment becomes sensible in a dynamic context in which youth are uncertain about future outcomes and the returns associated with their choices.

This paper aims to fill this gap in the literature by quantifying the importance of the option to re-enroll in the South African school to work transition. Specifically, I develop and estimate a dynamic model of school advancement and job search that allows for uncertain outcomes and re-enrollment following dropout during the schooling career. As youth learn the results of their enrollment and job search choices, they update their expectations about the relative returns to enrollment versus labor market participation, leading some who dropped out of school to re-enroll. In addition to explaining observed patterns of school enrollment, completion, and labor market participation, I quantify the importance of re-enrollment by conducting simulations in which the option to re-enroll before completing high school is restricted. I find that the re-enrollment restriction increases the proportion completing 12 years of schooling by 8 percentage points, as youth who would have dropped out under unrestricted re-enrollment reconsider the long-term consequences of doing so. To my knowledge, this is the first dynamic, sequential schooling model applied to South African data, and the first in the wider development literature to focus on the importance of school re-enrollment after periods of dropout.

The difficult school to work transition of South African youth in the post-apartheid era has been well documented. Racial disparities in school quality and student out-

[1]Comparable rates in the NLSY79 for the United States are 20% and 5%, respectively (Keane and Wolpin 1997).

comes found under apartheid (Case and Deaton 1999, Case and Yogo 1999) persisted in the post-apartheid era (Yamauchi 2005). Fiske and Ladd (2004) and Lam, Ardington and Leibbrandt (2011) find that high rates of grade repetition lead to lengthy school careers for black and coloured youth, while Lam, Leibbrandt and Mlatsheni (2009) find slow transitions from school to employment, with only 37% of black males aged 21-22 reporting ever obtaining employment. When facing such poor labor market outcomes, forward-looking agents may find it optimal to remain in school despite relatively low probabilities of advancement. Others may drop out and later re-enroll if unsuccessful in the labor market. Static human capital investment models, however, will not capture this behavior if agents adapt their expectations and alter their choices due to acquisition of new information, making a dynamic model appropriate. Eckstein and Wolpin (1999), Arcidiacono (2004), Stange (2009) and Joensen (2008), among others, estimate dynamic schooling models in which agents face uncertainty over academic advancement and labor market opportunities.

Like each of those studies, the model in this paper will allow agents to update their expectations about the relative returns to school enrollment and labor market participation in a dynamic discrete choice framework. I extend the previous work by allowing for an option to re-enroll in school after a spell of labor market participation, an important feature of school to work transitions in South Africa that prevailing dynamic schooling models have overlooked.[2] The pervasiveness of adverse shocks to South African youth – grade repetition while in school, high and lengthy unemployment when in the labor market – lead to the frequently rocky transitions from school to work observed in the data. Dropouts who update their expectations

[2]There are important exceptions. Keane and Wolpin (1997) do not restrict re-enrollment in their model despite referring to it as a "rare event" (pp. 16), while Eckstein and Wolpin (1999) allow re-enrollment within 5 years of entering high school. Belzil and Hansen (2002) estimate a dynamic discrete choice model of school enrollment that allows for re-enrollment after periods of "interruption," but they assume that such interruptions occur exogenously. Light (1995) estimates a hazard model of re-enrollment. Each of these papers uses U.S. data.

about future labor market opportunities during labor market spells may choose to re-enroll as a result. In this paper, I model these transitions and estimate their underlying structural parameters.

Estimates of the model's structural parameters confirm the hypothesis that enrollment choices reflect dynamic updating of the relative returns to schooling versus labor market participation. The model replicates basic patterns of enrollment, grade advancement and employment observed in the data, according to characteristics such as completed schooling and recent and cumulative school failures. I use my estimates of the underlying structural parameters relevant to enrollment to conduct policy simulations. To quantify the importance of re-enrollment in enrollment choices and schooling outcomes, I simulate the model under removal of the re-enrollment option. Enrollment rates rise sharply throughout the schooling distribution after restricting re-enrollment before high school completion, relative to both the data and results from the unrestricted model. I find that the re-enrollment restriction increases the proportion completing 12 years of schooling by 8 percentage points: by increasing the opportunity cost of dropout, restricted re-enrollment prolongs enrollment spells. Absent sources of market failure, however, such a policy would reduce social welfare, and is therefore intended as a useful thought experiment about the incentives to re-enroll rather than a specific policy proposal. Additional policy simulations, such as enforcement and extension of compulsory schooling and increasing school pass rates, lead to qualitatively similar results as the restricted re-enrollment simulation. The results suggest that the option to re-enroll is an important component of the incentives South African youth face when making schooling decisions.

I organize the paper as follows. The next section explains the motivation for the paper in greater depth. Section 2.3 develops the formal model and explains the estimation method. Section 2.4 describes the data. Section 2.5 presents results, model fit, and robustness checks. Section 2.6 presents results of policy simulations. Section 2.7 concludes.

2.2 Motivation for a dynamic schooling model with re-enrollment

2.2.1 Weaknesses in existing static and dynamic human capital models

In this section, I develop intuition for the dynamic model of school to work transitions that I formulate and estimate in this paper. I argue that both the traditional, static human capital model and prevailing dynamic schooling models fail to capture salient features of the South African educational and labor market context. A model that allows agents to update their expectations about future outcomes and alter their decisions accordingly is therefore appropriate. To illuminate these features, I focus on the choice to re-enroll in school after a spell in the labor market, a common behavior among the youth in my sample but one that is largely ignored in both static and dynamic human capital models in the literature. With a simple stylized model, I argue that dynamic updating of expectations based on past outcomes is essential to understand the choice to re-enroll in school.

First, consider a simple version of the classic Becker (1994) model of human capital investment, as discussed by Card (1999). The agent maximizes an indirect utility function based on the present discount value of (log) earnings net of schooling costs:

$$U(y(s)) = \ln y(s) - c(s) \tag{2.1}$$

where y is the annual earnings function, assumed concave in schooling s; and c is the cost of schooling, assumed convex in s. The agent chooses s to satisfy the first-order condition $\frac{y'(s)}{y(s)} = c'(s)$, where the marginal benefit of an additional year of schooling equals its marginal cost. This static model assumes the agent faces no uncertainty over schooling outcomes or the shapes of the earnings and cost functions, and therefore need not update her choice of s as new information about the returns to schooling (due to say, imperfect knowledge of ability or labor market opportunities) arrives. Yet such dynamic considerations are likely to be quite important to schooling choices, particularly for young adults. As Card (1999, pp. 1810-1812) acknowledges, "In fact the transition from school to work is often a bumpy one, as young adults move back and forth between full-time or part-time enrollment and part-time or full-time work...individuals do not necessarily know the parameters of their earnings functions when they make their schooling choices."

To address these concerns, a number of studies have developed and estimated dynamic human capital models, in which agents may update their schooling decisions based on arrival of new information about academic ability (Stange 2009, Eckstein and Wolpin 1999) and labor market opportunities (Keane and Wolpin 1997, Joensen 2008). However, even in most dynamic models (such as Heckman and Navarro 2007, Stange 2009, Joensen 2008), exiting school to enter the labor market is a terminal action. While this may be appropriate for developed country settings where re-enrollment in school after a spell in the labor market is rare, South African youth face high levels of grade repetition and unemployment that increase the uncertainty associated with their enrollment choices and make them especially vulnerable to adverse shocks as they transition from school to work.[3] The corresponding prevalence of re-enrollment in South Africa (in my sample, nearly one third of youth re-enroll at some point before age 24) make re-enrollment essential to incorporate in a dynamic human capital model.

[3]Figure 2.2, showing K-10 enrollment among 20 year-olds in selected countries, suggests that grade repetition and re-enrollment in South Africa are high in an international context.

2.2.2 Why agents re-enroll

In examining the importance of the option to re-enroll, first note that agents with this choice are unambiguously better off *ex ante* than those without, by virtue of the option value of re-enrollment.[4]

Under what circumstances will an agent with the option to re-enroll choose to do so? Consider a stylized dynamic setting in which labor market outcomes and academic success are stochastic. The agent is risk neutral and maximizes expected lifetime utility. For simplicity, assume no discounting of future payoffs. The model has three periods: in period one, the agent enters the labor market with zero skill units and finds employment with probability p. If she finds employment, she earns wage w_0, otherwise her payoff is 0. In period two, she may choose to remain in the labor market or re-enroll in school. If she remains in the labor market, she earns w_0 with certainty if she worked in period one, or else she must search again with probability p of success (hence expected payoff pw_0). If she re-enrolls in school, she earns a flow payoff of 0 and passes the grade with probability q, earning one skill unit. In period three, she must enter the labor market, but this time finds employment with certainty regardless of her employment status in previous periods, and earns a wage based on her skills.[5] If she never re-enrolled or re-enrolled and failed, she earns w_0. If she re-enrolled and passed, she earns $w_1 > w_0$. Figure 2.3 summarizes the

[4] To see how the re-enrollment option increases welfare, let $V^j(s)$ be the value function for an agent currently in the labor market, who receives flow payoff $w(s)$ from working (as a function of schooling s, the state variable) and discounts the future by discount factor β. Let $V^e(s)$ be the value function for enrollment, and s' the future level of schooling. Consider the Bellman equations for two young people in the labor market who differ only in that the second enjoys the option to re-enroll in school (in order to increase his or her stock of human capital and earn higher future wages) while the first does not:

$$V^j(s) = w(s) + \beta\mathbb{E}(V^j(s)) \tag{2.2}$$

$$V^j(s) = w(s) + \beta\mathbb{E}(\max\{V^j(s), V^e(s')\}) \tag{2.3}$$

The continuation values show that the agent with the re-enrollment option is better off, because $\mathbb{E}(\max\{V^j(s), V^e(s')\}) \geq \mathbb{E}(V^j(s))$.

[5] Assuming certain employment in period three simplifies the exposition without affecting my conclusions.

model; Appendix 2.A.1 extends the model to consider the initial choice to drop out of school.

Now consider her decision rule for re-enrollment. Since her only choice occurs in period two, we need to consider only her expected payoffs from period two forward, conditional on the period one outcome. If she found a job in period one, then she earns w_0 in each of periods two and three if she remains in the labor market, whereas she earns an expected wage of $\mathbb{E}(w) = qw_1 + (1-q)w_0$ if she re-enrolls. Hence an agent who is successful in period one job search will re-enroll if:

$$\begin{aligned}
&\mathbb{E}(U(\text{re-enroll})) \geq \mathbb{E}(U(\text{work})) \Leftrightarrow \\
&qw_1 + (1-q)w_0 \geq 2w_0 \Leftrightarrow \\
&\qquad\qquad q \geq \frac{w_0}{w_1 - w_0} \equiv \underline{q}_j
\end{aligned} \tag{2.4}$$

where the subscript j denotes that the agent obtained a job in period one. Re-enrollment thus follows a threshold rule, where the agent re-enrolls if the grade advancement probability exceeds $\underline{q}$, and remains in the labor market otherwise.

An agent who did not find work in period one and who remains in the labor market in period two will earn expected payoffs pw_0 in period two and w_0 in period three. If she re-enrolls, she earns the expected wage defined above. Therefore the re-enrollment rule for an agent who did not work in period one is:

$$\begin{aligned}
&\mathbb{E}(U(\text{re-enroll})) \geq \mathbb{E}(U(\text{work})) \Leftrightarrow \\
&qw_1 + (1-q)w_0 \geq pw_0 + w_0 \Leftrightarrow \\
&\qquad\qquad q \geq \frac{pw_0}{w_1 - w_0} \equiv \underline{q}_u
\end{aligned} \tag{2.5}$$

where the subscript u denotes that the agent was unemployed in period one. Note that because $\underline{q}_j \equiv \frac{w_0}{w_1-w_0} > \frac{pw_0}{w_1-w_0} \equiv \underline{q}_u$, the threshold probability of passing required to re-enroll is greater for an agent who obtained work in period one than for an agent who was unemployed. This means that all else equal, adverse labor market shocks make an agent more likely to re-enroll. These heterogeneous propensities to (re-)enroll based on information acquired from past outcomes will form the basis for the structural estimation in this paper. The result squares with economic intuition: an agent with bleaker labor market prospects will be more willing to risk failure in order to increase her human capital and earn a higher subsequent wage. The comparative statics of the model are also quite intuitive: the threshold probability to re-enroll $\underline{q}$ is decreasing in the return to schooling $w_1 - w_0$, so that agents are more likely to invest in uncertain human capital acquisition the greater the returns; and $\underline{q}$ is increasing in the probability of finding employment p, so that skill acquisition must be more assured the more likely one is to succeed in job search in the absence of such skills.

Most notable about the result, however, is that agents with identical expected lifetime incomes upon entering the labor market in period one nonetheless may exhibit different subsequent behaviors after their uncertainty about the initial job search outcome is resolved. This dynamic updating in response to labor market outcomes, and the associated behavioral sorting based on past outcomes to which it leads, is difficult to reconcile with standard human capital models assuming perfect foresight.[6]

[6]The notion that dynamic updating could drive re-enrollment accords with Heckman and Navarro (2007). Their sequential schooling model does not allow re-enrollment, but they nonetheless speculate, "In a general model, different persons could drop out and return to school at different times as information sets are revised" (pp. 364).

2.3 A model of the school to work transition

2.3.1 Timing and preferences

The model of the preceding section provided intuition about how one feature of the dynamic environment, the uncertainty of labor market outcomes, may lead to divergent choices and outcomes among otherwise identical agents. Yet this model is too simple to capture the observed behavior of agents making schooling decisions over their youth. In particular, the stylized model assumed homogeneity in the school advancement probability, no financial or psychic benefits or costs of schooling, and no wage return to labor market experience. In this section, I build on the intuition of the previous stylized model to develop a more complete model of enrollment choice in the presence of uncertain schooling and labor market outcomes. I relax each of the assumptions mentioned above, and explain how I estimate the structural parameters of the model using panel data on South African youth.

Consider the discrete-time, dynamic environment of a finitely-lived agent seeking to maximize lifetime utility by making discrete choices about school enrollment and labor market participation. In each period, the agent observes the state vector S_t, which summarizes all known information relevant to her choice at time t. She then chooses whether to enroll in school or enter the labor market; let $d_t = \mathbb{I}(\text{enroll at time } t)$. The agent may enroll in periods $t = t_0, \ldots, T_d$, but in all subsequent periods $t = T_d + 1, \ldots, T$, the agent must participate in the labor market. In each period, the agent receives choice-specific flow payoffs $U_t^d(S_t)$. Thus the agent's optimization problem at any time in the decision period $t \leq T_d$ is:

$$\max_d \mathbb{E}\left[\sum_{t=\tau}^{T_d} \beta^{t-\tau}\left\{d_t U_t^e(S_t) + (1-d_t)U_t^j(S_t)\right\} + \sum_{t=T_d+1}^{T} \beta^{t-\tau} U_t^j(S_t)\right] \tag{2.6}$$

where U^e and U^j are the utility functions for enrollment and labor market participation, respectively; β is the agent's discount factor; and the expectation is taken with respect to the evolution of the state space. The (indirect) utility functions are expressed in monetary units (South African rand per year), and capture the agent's preferences for enrollment and labor market participation as a function of S_t. I assume the agent is risk neutral. I also assume that the utility functions are additively separable in observable (to the econometrician) components X_t and choice-specific shocks observed only by the agent, ϵ_t^d, so that $U_t^d(S_t) = u_t^d(X_t) + \epsilon_t^d$. The agent learns the shocks $\epsilon_t = (\epsilon_t^j, \epsilon_t^e)$ at the beginning of each period, prior to the enrollment decision.

If the agent chooses to enter the labor market, she must search for employment, with probability of success conditional on current information contained in S. Her (flow) payoff in the labor market is her expected wage, i.e., the product of the probability that she will find employment and the wage she would earn if successful, all conditional on information known at the beginning of the period:[7]

$$\begin{aligned} U_t^j(S_t) &= \mathbb{E}[w(X_t)] + \epsilon_t^j \\ &= \Pr(\text{work}|X_t) \times w(X_t) + \epsilon_t^j \end{aligned} \tag{2.7}$$

I model wages as a linear function the observable state variables, i.e., $w(X_t) = X_t\gamma$. The probability of working is estimated as a logit that is linear in the same covariates as the wage equation.[8]

[7] Note that I implicitly normalize the value of unemployment to zero, consistent with the absence of a widely accessible unemployment insurance system in South Africa. Although other household recipients may receive public transfer payments, such as pensions, and share their resources with youth, such transfers do not depend on whether the youth is enrolled in school or in the labor market.

[8] This treatment of job search nests the approach of a companion paper that analyzes youth's first post-school labor market experiences using the same dataset (Levinsohn and Pugatch 2011). That paper estimates $\Pr(\text{work}|X_t) = q(X_t) \times (1 - F_{w|X}(w^*))$, where q is the arrival rate of job offers, F is the wage offer distribution, and w^* is the reservation wage. Whereas that paper focused on disentangling q from the parameters of F under alternative

If the agent chooses to enroll in school, she passes the grade level and accumulates associated human capital with probability that depends on the current state. Her payoffs depend on her net psychic benefits of schooling (b) based on all information known at period t, less the school fee. Thus we have:

$$U_t^e(S_t) = b(X_t) - fee(X_t) + \epsilon_t^e \tag{2.8}$$

The net psychic benefits of schooling include the agent's self-assessment of her ability based on both pre-determined characteristics and her schooling outcomes up to year t, as well as other characteristics that may influence the non-monetary benefits and costs of school.[9] I parameterize b as a linear function of X, i.e., $b(X_t) = X_t\alpha$.

Table 2.1 describes the elements of X, with further details in the Data section and Appendix 2.A.3. The table also notes the exclusion restrictions for the utility functions (2.7) and (2.8) and the transition equations for enrollment and passing (described in the next subsection).[10] The state space X includes race and gender dummies, indicators for ability quartiles based on results from a standardized test administered in Wave 1 of the panel, and a dummy for completion of grade 6 or higher at age 12,[11] when the model begins. In addition to the state variables in Table 2.1, I also assume that expected wages evolve with a quadratic trend in age beginning at the close of the decision horizon ($t = T_d + 1$) until the end of the model at $t = T$.[12] The model assumes that the returns to schooling are linear, but allows

measures of the reservation wage, my focus here is on the relative returns of enrollment versus job search, for which the distinction between the arrival rate and distribution of wage offers is less important. Therefore, I choose to forego estimation of these parameters, which simplifies estimation of the model.

[9] The psychic benefits of schooling may be negative, if an agent derives disutility from schooling.

[10] These exclusion restrictions are theoretically motivated, and not necessary for identification.

[11] Completing grade 6 by age 12 represents "on-time" school completion for students who enter at age 6, the modal school entry age in South Africa.

[12] I restrict the age-expected wage profile to begin at $t = T_d + 1$ for several reasons. First, because no individual in the data is observed after age 26 ($t = T_d + 3$), I prefer to control for actual work experience, rather than age, prior to T_d. Second, estimating age coefficients with wage data on such young agents is likely to be biased, so I use auxiliary data from the 2001 South African Census to estimate the age-expected wage profile. Details may be found in Appendix 2.A.3.

for changes in slope and intercept at high school completion.[13] By including time-varying covariates such as work experience and an indicator for employment in the last relevant period, the model allows agents to update their expectations of labor market outcomes based on the results of past choices (pass or fail for enrollment, work or search for job search).

In the enrollment utility function, I account for (possible) credit constraints through inclusion of indicators for lower quintiles of household incomes (y_q for quintiles $q = 1, 2$) and their interaction with HSG, reflecting the potential difficulties youth from lower-income households face in financing post-secondary education.[14] The completed schooling terms $(s, HSG, (s-12) * HSG)$ capture how the net psychic benefit of schooling evolves over time. The failure terms (f, f_{tot}) capture how the agent dynamically updates her self-assessed academic ability; including separate terms for recent and cumulative school failure allows the agent to place different weight on recent versus past outcomes.[15] Figure 2.4 summarizes timing, choices, and payoffs in the model.

2.3.2 State transitions

Because the choice environment is dynamic, the agent considers not only current payoffs when deciding to enroll in school, but also how the current period's choice

[13]Case and Yogo (1999) also find evidence for a slope change in the returns to schooling at high school completion, although their outcome measure is log wages, whereas my outcome is expected wages.

[14]This method of allowing household income to proxy for credit constraints follows Cameron and Heckman (2001), Carneiro and Heckman (2002) and Lam, Ardington, Branson, Goostrey and Leibbrandt (2010), although the resulting estimates may reflect the accumulation of educational disadvantage in low income households, rather than credit constraints. Formally incorporating credit constraints would require a richer model that incorporates asset accumulation and savings, as in Keane and Wolpin (2001) and Cameron and Taber (2004).

[15]The process of dynamic updating of the net psychic benefits of schooling in the model may be viewed as an approximation to Bayesian updating. Partition X_t into $[\mathcal{X}_0, \mathcal{X}_t]$, representing its time-invariant and time-varying elements, respectively. If both the prior and posterior distributions for b are normal, then the posterior mean $\mathbb{E}[b|X_t]$ is a linear combination of the prior mean $\mathbb{E}[b|\mathcal{X}_0]$ and the mean of $\mathcal{X}_t$:

$$\mathbb{E}[b|X_t] = a_1 \mathbb{E}[b|\mathcal{X}_0] + a_2 \bar{\mathcal{X}}_t$$

where $a_1 = \frac{1/\sigma_b^2}{(1/\sigma_b^2)+t\sigma_X^2}$ and $a_2 = \frac{t\sigma_b^2}{(1/\sigma_b^2)+t\sigma_X^2}$. Here, σ_b^2 and σ_X^2 are the variances of b and $\mathcal{X}_t$, respectively. Note that as t increases, the agent places relatively more weight on $\mathcal{X}_t$ when updating $\mathbb{E}[b|X_t]$, reflecting the importance of new information. Stange (2009) gives a similar interpretation of the agent's updating rule, noting that the process is similar to the normal learning model.

affects expected future payoffs. Following standard practice in dynamic discrete choice models, I assume that transitions of the state variable follow a first-order Markov process, and that transitions of the unobserved shocks are conditionally independent from those of the observed state.[16]

In the present model, an agent choosing to enter the labor market considers the probability of finding employment in the current period, which enters the labor market utility function, as well as the effect that current success or failure in job search will have on future job opportunities, which enters the state transition function. Because work experience variables are the only elements of the state space that evolve during periods in which the agent enters the labor market, the corresponding state transition simplifies to the probability that the agent finds work:

$$\begin{aligned}\Pr(X_{t+1}|X_t, d_t = 0) &= \Pr(x_{t+1} = 1, x_{tot,t+1} = x_{tot,t} + 1|X_t, d_t = 0)\\ &= \Pr(\text{work}|X_t, d_t = 0) \qquad (2.9)\end{aligned}$$

Similarly, for periods in which the agent enrolls in school, the state transition simplifies to the probability that the agent will pass the grade level:[17]

$$\begin{aligned}\Pr(X_{t+1}|X_t, d_t = 1) &= \Pr(s_{t+1} = s_t + 1, f_{t+1} = 0, f_{tot,t+1} = f_{tot,t}|X_t, d_t = 1)\\ &= \Pr(\text{pass}|X_t, d_t = 1) \qquad (2.10)\end{aligned}$$

[16] These assumptions allow me to estimate transitions of next period's observed state X_{t+1} as a function of the current observed state X_t only. This simplification comes at a cost, as it assumes that the transition probability for X_{t+1} depends on unobserved shocks only through their influence on the current state variable. This property would not hold if, for instance, an unobserved negative household shock that does not cause a youth to fail the current grade nonetheless persists and affects the following period's academic performance, leading her to fail the subsequent grade. However, assuming that persistent shocks affect the observed state variable upon impact, thereby allowing the current observed state to serve as a sufficient statistic for the distribution of next period's observed state, seems reasonable given the large gain in computational tractability.

[17] This treatment of the state transition while enrolled follows Lam et al. (2011), who treat school progression as the outcome of a threshold advancement rule with stochastic shocks.

I estimate both state transitions (2.9) and (2.10) as logits that are linear in X_t, using data on labor market and schooling outcomes, respectively.[18] The labor market utility function (2.7) and employment probability (2.9) functions also include the macro environment indicator ζ, but for simplicity agents forecast next period's macro state to be identical to the present. Denote the logit parameters governing the labor market transition (2.9) and enrollment transition (2.10) as ϕ_j and ϕ_e, respectively.

2.3.3 Model Solution and Estimation

Because the model has a finite horizon, it may be solved by backward recursion. Via Bellman's principle of optimality, rewrite the agent's problem as a dynamic programming problem, where the value function V_t is defined as the maximal expected present value of utility at time t, conditional on the state S_t:

$$V_t(S_t) = U_t(S_t) + \beta \mathbb{E}\left[V_{t+1}(S_{t+1})\right] \tag{2.11}$$

Because the value function assumes optimizing behavior by the agent from period t forward, it may also be expressed as the maximum over alternative-specific value functions:

$$\begin{aligned} V_t(S_t) &= \max_{d}\{V_t^d\} \\ &= \max_{d_t}\left\{U_t^d(S_t) + \beta \mathbb{E}\left[\max_{d_{t+1}}\{V_{t+1}^d(S_{t+1})\}|S_t, d_t\right]\right\} \end{aligned} \tag{2.12}$$

The assumptions of additive separability and conditional independence of the state space allow me to treat the $\mathbb{E}[\max]$ term in (2.12) as a double integral over the marginal distributions of X and ϵ. I discussed the transitions of X in the previous

[18] The state transitions may also be estimated nonparametrically. Since X is discrete, a simple bin estimator would make nonparametric estimation particularly straightforward, for example. However, because many covariate cells would have few observations, consistent estimation of the state transitions using nonparametric methods would be difficult, so I opt instead for parametric estimation using logits.

section. However, to make the problem tractable for estimation, I must assume a parametric form for the distribution of ϵ. I assume that the unobserved shocks are independently and identically distributed type I extreme value, i.e., $F(\epsilon_t^d) = \exp\left(-\exp\left(-\epsilon_t^d\right)\right)$. The i.i.d. assumption on ϵ is admittedly restrictive, as it implies that unobserved shocks affecting enrollment utility occur independently from those affecting labor market utility. Although it is easy to think of unmodeled shocks that could affect enrollment and labor force participation simultaneously (illness, household job loss, etc.), many types of shocks, such as the transfer of a talented teacher or the unexpected destruction of a job due to demand-side factors, will affect the utility of just one choice. Assuming that ϵ is drawn from an i.i.d. Type I EV distribution allows me to write $\mathbb{E}[\max]$ in closed form, a considerable computational savings, though I may relax this assumption in future work.

The assumptions of additively separable utility and conditionally independent and i.i.d. Type I EV shocks greatly simplify the form of the $\mathbb{E}[\max]$ terms of the Bellman equation, which we may rewrite as:

$$\begin{aligned}\mathbb{E}\left[V_{t+1}(S_{t+1})|S_t, d_t\right] &= \mathbb{E}\left[\max_d\{V_{t+1}^d(S_{t+1})\}|S_t, d_t\right] \\ &= \mathbb{E}\left[\ln\left(\sum_d \exp\left[V_{t+1}^d(X_{t+1})\right]\right)|X_t, d_t\right] + \nu \end{aligned} \tag{2.13}$$

where $\nu \cong .577$ is Euler's constant. Substituting the above closed-form expression for $\mathbb{E}[\max]$ into (2.12) serves to calculate the agent's expectations over the unobserved state variables ϵ. Combined with my assumptions on the transitions of the discretized observable state variables X, we have:

$$\mathbb{E}\left[V_{t+1}(S_{t+1})|S_t,d_t\right]=\sum_X\left(\left[\ln\left(\sum_d\exp\left[V^d_{t+1}(X_{t+1})\right]\right)|X_t,d_t\right]+\nu\right)\Pr(X_{t+1}|X_t) \tag{2.14}$$

The model's structural parameters are $\theta=(\phi_j,\phi_e,\gamma,\alpha)$, where ϕ_j and ϕ_e are the labor market and enrollment transition parameters; γ is the parameter vector for the wage equation; and α is the parameter vector describing enrollment preferences. For any value of θ, we may solve the model recursively by using the terminal condition that following the close of the decision horizon at T_d, the agent must enter the labor market. Therefore, I set $U^e_t(X_t)=U^j_t(X_t)=\mathbb{E}(w|X_t)+\epsilon^j_t$ for all $t>T_d$, which allows me to solve for the value functions in periods $t=T_d+1,\ldots,T$ by backward induction. The Type I extreme value assumption on the distribution of utility shocks and their conditional independence over choice alternatives allow me to express the agent's conditional enrollment probability as:

$$\Pr(d_t=1|X_t)=\frac{\exp\left(V^e_t(X_t)\right)}{\exp\left(V^e_t(X_t)\right)+\exp\left(V^j_t(X_t)\right)} \tag{2.15}$$

With data on the sequence of enrollment choices $\{d_t\}_{t=1}^{T_d}$, I decompose an individual's contribution to the likelihood function as:

$$l_e(\phi_e)=\prod_{t=1}^{T_d}\Pr(\text{pass}_t|X_t,d_t=1)^{d_t} \tag{2.16}$$

$$l_j(\phi_j)=\prod_{t=1}^{T_d}\Pr(\text{work}_t|X_t,d_t=0)^{1-d_t} \tag{2.17}$$

$$l_w(\gamma)=\prod_{t=1}^{T_d}f\left(w(X_t)^{\mathbb{I}(\text{work}_t)}\right) \tag{2.18}$$

$$l_d(\alpha)=\prod_{t=1}^{T_d}\Pr(d_t=1|X_t)^{d_t}\Pr(d_t=0|X_t)^{1-d_t} \tag{2.19}$$

where the components are the grade transition (2.16), the labor market transition (2.17), the wage equation (2.18)[19] and the enrollment choice equation (2.19). The individual's likelihood contribution is then:

$$l(\theta|X) = \prod_{t=1}^{T_d} l_e \cdot l_j \cdot l_w \cdot l_d \tag{2.20}$$

With panel data $\{d_{it}, X_{it}\}_{t=1}^{T_d}$ for $i = 1, \ldots, n$ individuals, the likelihood function becomes:

$$L(\theta|X) = \sum_{i=1}^{n} \ln l_i(\theta|X_i) \tag{2.21}$$

Estimation of θ may proceed by a nested fixed point (NFXP) algorithm, as in Rust (1987): in the "inner loop," the current guess of θ is used to solve the model recursively, as described above, while in the "outer loop," each guess of θ is updated through a numerical optimization procedure. The process repeats until convergence.

Because the likelihood function is additively separable in the contributions of enrollment ($d = 1$) and labor market ($d = 0$) components, I am able to estimate these parts sequentially. First, I use data on employment and school advancement to estimate (ϕ_e, ϕ_j) using (2.16) and (2.17); employment and accepted wage data are used to estimate γ using (2.18). The resulting estimates $(\hat{\phi}_j, \hat{\phi}_e, \hat{\gamma})$ are then substituted into (2.19) to estimate α. I set the discount factor β to 0.95. The likelihood function converges to identical parameter estimates from several different arbitrary starting values.

2.3.4 Parameter Identification

This subsection discusses identification of the model's structural parameters $\theta = (\phi_j, \phi_e, \gamma, \alpha)$. Employment and grade advancement probabilities conditional on ob-

[19] Here, $f(\cdot)$ is the density of the wage residuals.

servable characteristics identify the transition parameters ϕ_j and ϕ_e, respectively, as in standard logistic regressions for binary outcomes. Variation in state variables across individuals and over time among those in the labor market identifies γ, the parameters of the expected wage function. Although γ is identified from the selected sample that chooses to participate in the labor market and therefore may not be interpreted as causal, the discrete choice model fully specifies the selection process. Differences in enrollment rates by net psychic benefits of schooling (b), both across individuals and over time, identifies the enrollment utility function parameters α. For example, the difference in enrollment rates between those who failed their last grade enrolled and those who passed identifies the coefficient on the failure indicator in (2.8).

In the model, the transition parameters ϕ_e and ϕ_j capture youth expectations about future schooling and labor market outcomes, respectively. The parameters α and γ capture expectations about the utility associated with enrollment and labor force participation, respectively. Precise estimation and sensible signs of coefficients on time-varying characteristics within each structural parameter would provide evidence in favor of my hypothesis that youth dynamically update their expectations about the relative returns to enrollment versus labor market participation. The most notable coefficients are those on school failure and work experience (recent and cumulative), as such coefficients reflect updating based on previously uncertain schooling and labor market outcomes. Evidence of dynamic updating may also be found in the schooling coefficients (schooling, the indicator for high school graduate, and their interaction), because schooling also evolves based on past enrollment choices and their outcomes.

2.4 Data

2.4.1 The Cape Area Panel Study

I estimate the model with data from the Cape Area Panel Study (CAPS), a longitudinal study of youth in metropolitan Cape Town, South Africa (Lam, Ardington, Branson, Case, Leibbrandt, Menendez, Seekings and Sparks 2008). CAPS sampled about 4,800 youth aged 14-22 in Wave 1 (2002) and currently includes four publicly available waves, the most recent conducted in 2006. In Wave 1, retrospective life histories were collected for each year stretching back to birth, and include information on school enrollment and advancement, job search, and employment. I update this retrospective life history data with information from Waves 2-4 to construct the panel used in this paper. I make several sample restrictions. I keep only those youth observed until at least age 18. Those who report advancing two or more grades in a year, or without continuous information on enrollment, are dropped from the sample. I drop those who report entering school prior to age 4 or exiting school after age 24 (which effectively sets $T_d = 24$ as the decision horizon).[20] I also drop those whose educational histories, by my definition of school advancement (described below), place them with more than 16 years of completed schooling. This leaves $N = 3,806$ individuals in the sample, comprising 68,898 person-years from ages 4-26. Table 2.A.1 shows the panel balance at selected ages. The sample size drops sharply at later ages due to both the young ages of entry into the panel (i.e., right-censoring) and attrition in later waves. Although whites are more likely than blacks or coloureds to be out of the sample after age 18, panel imbalance by race is not severe.

At each age, the data contain information on the youth's enrollment status. If

[20]Reporting school entry prior to age 4 is more likely to reflect measurement error than childhood precociousness, in my view. Few observations are available in the data after age 24 (less than 20% of the sample is observed beyond 24, due to starting ages in Wave 1 and attrition), raising concern that estimating the model with these enrollment choices will result in severe finite-sample bias.

enrolled, the data report the school outcome (pass or fail); if in the labor market, the data report whether employed.[21] Enrollment in school supersedes labor market participation when these are reported to occur simultaneously at a given age, in order to keep these choices consistent with their treatment as mutually exclusive in the model. Work or search while enrolled never exceeds 3% of the sample at any grade level, and never exceeds 2% during grades 1-12, making me confident that modeling enrollment as distinct from labor force participation is reasonable. For post-secondary education (grades 13-16), I expand the definition of grade advancement to include a response of "no grade/continuing" to the survey question on school result, as this is consistent with continuation in a multi-year degree program, and is the modal response for these grades.[22] Table 2.A.2 shows, using monthly calendar data collected during CAPS Waves 1-4, that school enrollment and labor force spells last at least one year on average, regardless of whether I allow a spell to include other concurrent activities or restrict the definition to full-time spells (although removal of censored spells reduces mean duration somewhat, because the longest spells tend to be censored). Moreover, because school enrollment spells tend to follow the January-December academic year (as shown in Figure 2.5), my assumption in the model that time periods last one year seems reasonable.

Data on accepted wages and school fees are available only at each panel interview rather than at each age. To overcome this restriction, I predict wages and school fees by regressing observed values on the state variables in Table 2.1.[23] Appendix 2.A.2

[21] Although search behavior is also reported in the sample, all non-enrolled youth are assumed to be participating in the labor market. In principle, the model could be extended to include non-participation in the labor market as a discrete choice. However, a companion paper using CAPS (Levinsohn and Pugatch 2011) finds that 64% of those who permanently exit school and never report searching nonetheless find employment by the end of the sample, suggesting that labor market participation is not synonymous with active search for these youth.

[22] Unfortunately I am unable to distinguish reliably whether those in post-secondary education are making satisfactory progress to degree completion, because few students report academic failure in these grades.

[23] I predict school fees using only race dummies, completed schooling, and a dummy for high school graduate, to be consistent with racial disparities in school expenditure and the discontinuous jump in fees in post-secondary schooling. I am concerned that if additional information on, say, cumulative grades failed were used to predict school fees, results could reflect unmodeled endogenous school choice and therefore be biased. For both expected wages and

contains additional information on the South African education system, while Appendix 2.A.3 contains additional information on the sample and variable definitions.

For the purposes of the model, the first decision period t_0 occurs at age 12, as no youth enter the labor market and obtain work prior to this age, so extending the model to an earlier age would be superfluous. Of course, by age 12 substantial, non-random differences have already arisen among youth due to unobserved characteristics related to family background and motivation. Therefore, I include a dummy for completion of at least 6 years of schooling at age 12 as a state variable (representing "on-time" completion for a student who enters at age 6), which proxies for the effect of such early life differences. The decision period ends at $T_d = 24$. The model ends at $T = 60$, the age of public pension eligibility in South Africa.[24] I top-code grades failed and work experience since age 12 at 3 and 4, respectively, which reduces the dimension of the state space while still accurately capturing more than 99% of the person-year observations in the sample.

2.4.2 Summary statistics and stylized facts

Table 2.2 presents summary statistics for the sample. The racial distribution reflects the unique racial composition of Cape Town, where coloureds (mixed racial heritage) are prominent. Schooling careers range on average from ages 6 to 18, with a mean of 1.2 grades failed.[25] There is wide variation (standard deviation 1.2 years) in completed schooling even by age 12. Only slightly more than half of the sample has worked, though this in part reflects right-censoring of school careers rather than failure in job search. As mentioned in Section 2.1, one third of the sample has re-

school fees, I replace predicted values below the minimum value observed in the sample with the first percentile from the data, in order to avoid non-positive predicted values and extreme outliers.

[24] Although individuals must also pass a means test to receive a public pension, in practice about 80% of elderly blacks and coloureds receive the pension. The pension is quite generous relative to the median incomes of these groups (Case 2004).

[25] "Failure" refers to any type of failure to complete a grade level while enrolled, and includes withdrawal from school as well as academic failure.

enrolled in school (after disenrolling for at least one year) at some point in the sample, including 20% who re-enroll before completing secondary school.[26] As explained in Section 2.2, this re-enrollment behavior is difficult to reconcile with standard human capital theory, but is a prominent feature of the dynamic model I estimate.[27]

Table 2.3 shows the means of selected variables across various demographic and educational categories. Racial disparities in schooling and labor market outcomes are substantial: blacks enter school one year later and last enroll nearly one year earlier than whites, on average, and fail 1.7 grades, compared to 0.5 for whites. Black wages are less than one third those earned by whites. Coloureds occupy an intermediate group between these extremes. Gender differences are not nearly as pronounced, however. Schooling and labor market outcomes follow the expected patterns when splitting the sample by ability or completed schooling. However, re-enrollment is not correlated with schooling or labor market disadvantage in a simple way. For instance, although similar proportions of black and whites re-enroll at some point in the sample (39% and 43%, respectively), 31% of black youth re-enroll before completing grade 12, compared to only 14% of whites. Similarly, high ability (those scoring above the age-adjusted, within-sample median on the standardized literacy and numeracy evaluation) youth re-enroll at higher rates than those with low ability (39% to 28%), but both groups re-enroll at similar rates before grade 12 (18% and 22%). This suggests that the option to re-enroll may be taken for different reasons according to background characteristics and current levels of completed schooling.

[26]Comparable rates in the NLSY79 are 20 percent ever re-enrolled and 5 percent re-enrolled before high school completion (Keane and Wolpin 1997).

[27]Moreover, it is unlikely that measurement error, in the form of recall bias, leads to a gross overstatement of re-enrollment rates. If recall bias were severe, I would expect re-enrollment rates to be greater for observations from ages before youth entered the panel and were subject to frequent interviews. Yet in a regression of the re-enrollment indicator on a dummy for whether the observation overlapped with the time of the panel, the coefficient is positive and significant (coefficient 0.03, t-statistic 2.5), which is the opposite of the result I would expect to find in the presence of recall bias. The regression includes a full set of age and schooling dummies, controls for work experience and school failure (in both the last relevant period and cumulative), and individual fixed effects, with standard errors clustered at the individual level.

The model accounts for such heterogeneity through its demographic controls and by allowing enrollment preferences and labor market returns to change discontinuously following high school completion.

Table 2.4 depicts patterns of enrollment and school failure by level of completed schooling. Enrollment rates are relatively high among those in the primary and early secondary school grades (1-9), though not universal,[28] falling steadily to a trough around 20% at high school completion. Those with some post-secondary education (grades 13 and above) tend to remain enrolled, however. Enrollment and failure rates are negatively correlated, with failure highest between 9-11 years of completed schooling (when students prepare to pass the "matric," or secondary school completion exam). Table 2.A.3, Panel (a) shows that such failure stems largely from academic failures, rather than withdrawal from school or continuation in an educational program.[29,30] Dropout increases in secondary school and peaks upon secondary completion. Re-enrollment follows a similar pattern, with re-enrollment rates highest among those with 9-11 years or more than 12 years of schooling.

The discrete choice model of this paper treats (expected) wages and school fees as the directly measurable economic factors influencing enrollment decisions; Table 2.4 also shows employment, wages and fees as functions of completed schooling. Employment rates, wages and school fees increase substantially upon completion of secondary school. The apparent non-linearity of returns at secondary completion motivates the inclusion of terms (HSG and $HSG * (s - 12)$, described in Table

[28] Recall that the sample begins at age 12, when many youth will already be close to primary school completion. Extending the sample to age 4 makes primary school enrollment nearly universal.

[29] Of course, school withdrawal could be another form of academic failure in which a failing student stops attending rather than face outright failure.

[30] Note that beginning in post-secondary (completed schooling 12 or above), I alter the definition of failure so that "no grade/continued" counts as passing the level, which explains why this response is zero for this group. Relatively high rates of "no grade/continued" prior to secondary school completion may be due to enrollment in vocational training or equivalency programs, although enrollees in National Training Certificate (NTC) programs report similar "no grade/continued" rates as the full sample.

2.1) allowing the intercept and slope of the expected wage equation to change in post-secondary schooling. Yet despite the large increase in fees for higher education, their (mean) reported level remains markedly lower than the corresponding wage throughout the schooling distribution.[31]

2.4.3 Evidence consistent with dynamic updating of expected returns

Although the descriptive evidence presented thus far in this section helps to illustrate the relative costs and benefits of enrollment over the school career in the South African context, relatively straightforward modifications to the standard, static model of human capital investment (incorporating school failure and employment rates to modify expected returns, for instance) could still capture observed behavior quite well.[32] To justify the specification of a dynamic model, we must also observe behavior that could not be accounted for in the static case, such as enrollment (and re-enrollment) decisions reflecting dynamic updating of expected returns following acquisition of new information. I will argue observed enrollment choices in the sample are consistent with such dynamic updating. In the following subsection, I consider alternative explanations for the observed enrollment and dropout patterns.

I have hypothesized that much labor market entry and exit occurs because youth dynamically update their expectations about the relative returns to enrollment and job search based on the outcomes of their past choices. For instance, a student who fails a grade may be more likely to fail in the future, as she learns that she is less able academically than previously thought. Similarly, a youth who succeeds in her job search may be more likely to succeed in the future, as she builds a professional

[31]These values, as with all currency-denominated units in the paper (unless otherwise noted), are in real South African rand per year, base year 2002. The South African rand traded at 10.3 per US dollar in August 2002, when CAPS Wave 1 began.

[32]For instance, Altonji (1993) develops a model of sequential schooling choice with uncertain completion, but ultimately estimates reduced-form regressions for wages and completion probabilities.

network and firm-specific human capital or forms a long-term contract with her employer. Figure 2.6 shows that these patterns exist in the data.[33] If youth know these conditional probabilities of passing and employment, they may change their behavior as they update their expectations based on the outcomes of their past choices.

The transition matrices in Table 2.5, which show how youth in the sample move between enrollment and non-enrollment, suggest that enrollment decisions reflect such dynamic updating about returns. In the full sample, 82% of those enrolled in a given year will remain enrolled in the following year, i.e., the overall dropout rate is 18%. Of those not enrolled in school, 88% will remain not enrolled in the following year, for a gross re-enrollment rate of 12%. Conditioning on past outcomes reveals behavior consistent with dynamic updating of the relative returns to enrollment versus labor market participation. Among those enrollees who passed the grade level, 87% remain enrolled in the following period, while just 48% of those who failed remain enrolled. Similarly, only 5% of those employed choose to re-enroll in school in the following period, compared with 15% of the unemployed. This pattern of re-enrollment is consistent with the stylized model presented in Section 2.2.2, in which agents who fail to find a job are more likely to re-enroll as they reassess their expected returns to labor market participation.

The middle and rightmost panels of Table 2.5 show that transition patterns consistent with dynamic updating are present regardless of secondary completion. For both those with less than or at least 12 years of completed schooling, enrollees who failed are more likely to drop out of school than those who passed, and those who were unemployed while not enrolled are more likely to re-enroll in school. This descriptive

[33] Of course, these patterns may also simply reflect selection, as the less able and employable are more likely to fail and be unemployed, respectively, both now and in the past. The model fully specifies the dynamic selection process, however, and includes a rich set of controls to account for such selection.

data is consistent with agents dynamically reassessing the returns to enrollment as they acquire new information based on past outcomes. The importance of such updating to the decision-making process makes a dynamic human capital investment model appropriate.

2.4.4 Alternative explanations of observed enrollment behavior

Several alternatives to my preferred "dynamic updating" hypothesis could explain the observed enrollment patterns, among them: 1) responses to household shocks; 2) preferences for leisure; 3) credit constraints and the need to accumulate savings; and 4) behavioral explanations. I address each of these in turn.

Responses to household shocks. Youth experiencing household shocks, such as a parent's job loss or the severe illness of a household member, may choose to change their enrollment behavior abruptly as a result. Unfortunately, the available life history data are not rich enough for me to model or otherwise rule out the role of such shocks. The most detailed data on household shocks are available only during the period of the survey waves (2002-2006), leaving an incomplete panel relative to the longer panel I use in estimation. In this limited panel, household shocks are not statistically associated with dropout or re-enrollment (after controlling for grade failures, work experience, age, schooling and individual fixed effects), as shown in Table 2.A.4.[34] Additional data on household characteristics available throughout youth life histories include marriage, pregnancy, and co-residence with parents or grandparents, which are more properly modeled as endogenous choices rather than external shocks, and would therefore require a richer model.[35] However, the model in

[34]This evidence is consistent with Lam et al. (2010), who find using CAPS data that household shocks (measured in Wave 3) do not significantly affect post-secondary enrollment.

[35]Empirically, pregnancy does not play a prominent role in dropout decisions: among new dropouts' self-reported reasons for leaving school, girls cite pregnancy 4% of the time, though the data is hampered by high rates of non-response.

its current form is consistent with the presence of household shocks, either through their effects on observable outcomes (such as grade advancement and employment) or through the unobservable (to the econometrician) state variable ϵ.

Preferences for leisure. Another alternative explanation for the observed dropout and re-enrollment patterns is that youth prefer to take some time off from school occasionally and enjoy a period of leisure. Such a preference for leisure would be particularly likely at or near the completion of secondary school, when youth might wish to enjoy a "gap year" before continuing their studies. Yet the proportion of youth who cite economic factors (rather than other reasons more likely to signal a preference for leisure) as the reason for dropout *increases* to 23% among those with 12 years of schooling (compared to 12% for those with less than 12 years of schooling), precisely when we would expect youth to be taking a "gap year."[36] Moreover, a sizable fraction (20%) of the sample drops out and re-enrolls prior to completing high school (see Table 2.2), before the typical "gap year" ages. This suggests that preferences for leisure, in the form of an extended period of non-enrollment, are not driving dropout and re-enrollment rates.

Credit constraints. If youth are credit constrained, they may need to drop out of school temporarily in order to accumulate savings, and re-enroll when they have saved enough to pay school fees. Increases in re-enrollment rates after high school completion (Table 2.4) and the relatively high rates of youth from the poorest households citing school fees as their reason for dropout (Table 2.A.3, panel [b]) are consistent with credit constraints. Yet credit constraints do not appear to bind for the majority of youth: Table 2.A.3, panel (b) also shows that for the full sample, the proportion of dropouts reporting that they can not afford to attend school *falls* from 7% to 6%

[36] Economic reasons for dropout include work, job search, or not being able to afford school. Non-response is the modal reason for not enrolling in school among dropouts, but the increase in those citing economic factors among high school graduates suggests that preferences for leisure do not drive dropout at that stage.

among high school graduates, the group most likely to face prohibitively high education expenses.[37] Formally incorporating credit constraints would require a richer model that incorporates asset accumulation and savings, as in Keane and Wolpin (2001) and Cameron and Taber (2004). However, in the model I follow the existing literature (such as Cameron and Heckman 2001, Carneiro and Heckman 2002, Lam et al. 2010) and allow household income to proxy for credit constraints. Specifically, I include indicators for lower household income quintiles and their interaction with an indicator for high school graduation; if youth from the poorest households face credit constraints that affect their enrollment choices, the coefficients on these terms will be negative. It must be noted, however, that the resulting estimates can not prove the existence of credit constraints, but may instead reflect the accumulation of educational disadvantage in low income households.

Behavioral explanations. Many behavioral explanations are possible for the observed enrollment patterns, such as myopia or asymmetric information about the relative returns to schooling versus labor market participation. Although I do not test such alternative behavioral hypotheses directly, if the model provides a good fit to the data, it would show that a model based on rational, forward-thinking, and well-informed agents is consistent with observed behavior. I turn to the results and fit of the model in the next section.

2.5 Results

2.5.1 Parameter estimates

Table 2.6 presents estimates of the model's structural parameters $\theta = (\phi_e, \phi_j, \alpha, \gamma)$.[38] The transition parameters for school advancement and employment (ϕ_e and ϕ_j,

[37] The sample used to calculate the statistic is "new" dropouts, i.e., those who have dropped out of school after enrolling in the previous year.

[38] The model converges to identical parameter estimates from arbitrary starting values.

respectively) in columns (1) and (2) are mostly as expected, though with a few surprises. In the school advancement equation of column (1), I find that blacks, coloureds and males are significantly less likely to pass to the next grade level. The probability of passing is increasing in ability and for those who completed at least 6 years of schooling by age 12. The probability of passing declines with the level of schooling, though the separate slope and intercept terms for high school graduates are positive, reflecting the lower failure rates in post-secondary education reported in Table 2.A.3, panel (a). Perhaps surprisingly, those who failed their last period enrolled are more likely to pass, though this effect is not statistically significant. Unsurprisingly, however, the probability of advancement declines with the cumulative number of grades failed since age 12.

Column (2) shows results for the labor market transition equation (2.9), corresponding to the parameter ϕ_j. Most coefficients have the expected sign: blacks and females are significantly less likely to find work than white males; the probability of employment is increasing in schooling, with a higher intercept (though no steeper gradient) for high school graduates. As hypothesized, employment is path dependent, with the coefficient on employment during the last period of non-enrollment quite large and precise. Surprisingly, those with at least 6 years of schooling at age 12 and more total work experience are less likely to be employed. Employment is quite sensitive to macro conditions, with a precisely estimated negative coefficient on the bad macro environment dummy.

Column (3) presents parameter estimates from the enrollment utility function of the dynamic discrete choice model, i.e., the coefficients from (2.8).[39] The coefficient

[39] A note on the interpretation of coefficients from the enrollment equation in Table 2.6: as is well known, logit coefficients such as these are identified only up to scale. Because the units of the labor market utility function (as well as the school fee term in the enrollment utility function) are in South African rand per year (ten thousands, base year 2002), so too are the units of the enrollment utility coefficients, provided the scale parameter of the i.i.d. Type I extreme value shocks is unity, as assumed.

for black shows that blacks are no more likely to enroll in school than whites, conditional on all other covariates. For coloureds, however, the coefficient is negative and significant, consistent with their shorter school careers. Coefficients on schooling at age 12 and the highest ability group are positive, consistent with greater consumption value of schooling for the more academically able. The schooling coefficient is negative, reflecting increases in both academic difficulty and opportunity costs as youth progress through school. Coefficients on recent and cumulative failure are negative, consistent with dropout after learning about one's academic ability through school outcomes. These school failure coefficients are strong evidence in favor of the hypothesis that youth dynamically update their enrollment behavior based on past schooling outcomes.

Since the school fee term in the enrollment utility function (2.8) accounts for the direct costs of schooling, the negative sign on the indicator for high school graduate may reflect any of several impediments to higher education: greater psychic costs, difficulty in obtaining admission, or credit constraints. The coefficients on all household income terms are negative, which may reflect the accumulated disadvantage of poverty, credit constraints in school financing, or both. The argument for credit constraints strengthens when considering the negative interaction terms on high school graduate and indicators for the first two household income quintiles, which means that conditional on high school graduation, youth from poorer households are less likely to enroll in post-secondary education, as also found by Lam et al. (2010). However, neither of these interaction terms is significant at conventional levels, weakening the case for credit constraints. Ultimately, the model can not determine whether lower post-secondary enrollment by poorer youth stems from credit constraints or accumulated disadvantage.

Coefficients for the wage equation (parameter γ) presented in column (4) generally show the expected pattern with respect to demographic, schooling and ability variables, with a large wage return to higher education evident in the high school graduate slope coefficient. A bad macro environment dampens wages, also as expected. The school fee regressions reported in Table 2.A.5 show that fees are less for blacks and coloureds, likely reflecting lower school quality. School fees increase with grade level, with a large and discontinuous jump in post-secondary schooling.

2.5.2 Model fit

Figure 2.7 assesses model fit by comparing observed versus predicted grade advancement (panel [a]) and employment (panel [b]), corresponding to the transition equations (2.10) and (2.9), respectively. The model fits the data well for both types of transitions. Figure 2.8 plots enrollment probabilities for various state variables. The model does well in predicting enrollment by completed schooling (panel [a]), except in the final levels of post-secondary schooling. The model also does well in predicting enrollment as a function of previous school failure, reflecting students' learning about their academic ability: in panel (b), agents in the data and the model are less likely to enroll following a failed grade, which is also true in the case of cumulative failures, as in panel (c). Panels (b) and (c) show that the model captures how youth dynamically update their behavior based on past schooling outcomes.

Another important dimension of fit to consider is whether the model can replicate the patterns of dropout and re-enrollment observed in the data. Figure 2.9 presents observed and predicted dropout and re-enrollment rates by completed schooling. To calculate predicted values, I simulate 50 enrollment histories for each observation, using the state variable at age 12 (the first decision period) as the initial condition. The model predicts the stylized pattern of dropout fairly well, accurately capturing the rise of dropout through the end of secondary school and its subsequent fall in post-secondary. However, the predicted magnitudes are sometimes quite different from the data, overestimating dropout rates throughout the schooling distribution. The model predicts the secular rise in re-enrollment through the secondary school years (despite noise in the early grades due to small sample sizes), but again does not predict magnitudes well. These discrepancies between actual and predicted dropout and re-enrollment rates likely arise because the model does not adequately capture the switching costs (both monetary and psychic) youth face when considering dropout or re-enrollment. Despite the relatively high re-enrollment rates observed in the data, youth exhibit more continuity in their enrollment choices than the model predicts.

2.5.3 Robustness checks

In considering the school to work transition of South African youth throughout their adolescence, the model of this paper treats transitions between high school and post-secondary education and transitions between schooling and labor force participation in the same framework. Although the model allows for enrollment preferences and the returns to schooling to change discontinuously at high school completion, one might still worry that school to work transitions for those with post-secondary schooling are fundamentally different in nature than for those without. To explore this possibility, I re-estimate the model using only observations from youth who do not enroll in post-secondary schooling. Comparing the results from this selected sample to those from the full sample allows one to see the extent to which the inclusion of those pursuing post-secondary schooling drives my results.

Table 2.A.6 shows results from the sample with no post-secondary schooling. Results are qualitatively similar to those from the full sample presented in Table 2.6. In particular, youth who never enroll in post-secondary schooling are more likely to

drop out of school if they have failed their most recent grade and with cumulative grade failures, just as in the full sample. This result suggests that the process of dynamic updating of the expected relative returns to enrollment versus labor force participation is the same for these youth.

One may also be concerned that my assumptions about how forward-looking youth are in the model drives my results. In particular, one might worry that results will change dramatically depending on my choice of the discount factor β or the time horizon T. Table 2.A.7 shows that this is not the case. The table shows estimates of the enrollment utility parameter (γ) under alternative scenarios for β and T.[40] In column (1), β is set to 0.9 (from 0.95 in the main estimates). In columns (2) and (3), T is set to ages 52 and 65 (from 60 in the main estimates), respectively; $T = 52$ corresponds to South African life expectancy in 2006, while $T = 65$ corresponds to the age of public pension eligibility for males prior to a recent change in the law (and therefore may have been the retirement age males had in mind when making enrollment choices). Across all columns, coefficients are qualitatively similar to those from the main results in column (3) of Table 2.6.

2.6 Policy Simulations

In this section, I return to the motivating question of the paper: how important is the opportunity to re-enroll in the school to work transition of South African youth? To answer this question, I simulate a counterfactual scenario in which the option to re-enroll in school following dropout is restricted, and explore how enrollment decisions and outcomes change as a result. I also consider several other policy simulations that potentially alter the incentives for enrollment versus labor market participation:

[40]Estimates of the enrollment utility parameter γ only are shown because this is the only structural parameter that depends explicitly on the discount factor β and the time horizon T in the model. All other structural parameters remain unchanged.

compulsory schooling, increased pass rates, and a subsidy for post-secondary school fees. I conclude the section by comparing results from the restricted re-enrollment simulation with those from the other policies.

2.6.1 Restricted re-enrollment

Given the prevalence of re-enrollment among youth in the sample, it is natural to ask what role the option to re-enroll plays in decisions about schooling and labor market participation. Youth who know that they may re-enroll in school after a labor market spell may make different choices about human capital investment than those for whom labor market participation is effectively a terminal action. The model developed in this paper is well-suited to explore the ramifications of the option to re-enroll. By fully specifying the enrollment choice environment for youth and recovering its structural parameters, the model may be adapted to consider a counterfactual scenario in which the option to re-enroll is restricted.

To quantify the importance of re-enrollment, I modify the model by restricting re-enrollment prior to completion of high school, making labor market participation before completing grade 12 an absorbing state. I still allow for re-enrollment after high school completion, because restricting re-enrollment in the transition between secondary and post-secondary education seems particularly unrealistic. Although this scenario of restricted re-enrollment is admittedly extreme (and possibly unenforceable under the existing educational infrastructure in South Africa), it nonetheless provides a useful thought experiment, allowing me to assess empirically the importance of the re-enrollment option. The exercise is similar in spirit to that of Heckman and Urzua (2008), who simulate the elimination of the General Educational Development certificate (GED), a high school equivalency certificate available in the U.S. that is typically earned by high school dropouts who later re-enroll.

Specifically, I simulate 50 enrollment histories for each observation, using the state variable at age 12 as the initial condition, and eliminating the option to re-enroll before completing grade 12. Comparing simulated enrollment probabilities and completed schooling to analogous simulation results from the unrestricted model provides an assessment of how the option to re-enroll influences youth schooling choices and outcomes. The expected effects of such a policy are ambiguous: although restricting the freedom to enroll in school among a subpopulation (in this case, dropouts) ought to depress enrollment rates, potential dropouts facing no possibility of returning to school may choose to remain in school as a result, causing overall enrollment rates to rise.

Figure 2.10 shows enrollment and dropout probabilities from the data, the unrestricted model, and the restricted re-enrollment simulation. Enrollment rates in the restricted re-enrollment scenario exceed observed and predicted (from the unrestricted model) enrollment rates throughout most of the schooling distribution (panel [a]). The magnitude of these increases are striking: for example, 86% of those with 7 years of schooling enroll in school under the restricted re-enrollment scenario, compared to 76% in the data and 71% in the unrestricted model. Dropout rates are correspondingly lower in the restricted re-enrollment simulation (panel [b]). The intuition for these results is that youth facing restricted re-enrollment will be less willing to experiment in the labor market if the opportunity cost includes not only foregone human capital during a short labor market spell, but also foregone human capital in all future periods.

Table 2.7 shows the distribution of completed schooling at age 20 from the data and the restricted re-enrollment simulation. Column (1) reports the percentage of the sample in each schooling category under simulation of the unrestricted model;

column (2) reports percentages under the restricted re-enrollment simulation; and column (3) reports the difference (standard errors in parentheses). The results show a clear rightward shift in the schooling distribution. The shares of the sample completing high school or attaining some post-secondary education climb by 8.3 percentage points each, with corresponding declines in the lower schooling categories. The results are qualitatively similar to those of Heckman and Urzua (2008), who use a structural schooling model to find that elimination of the GED would increase secondary completion by 2.1 percentage points in the U.S.[41] In both cases, restricting the option to re-enroll reduces dropout. While the magnitudes of the effects I find may be surprisingly large, the large share of my sample re-enrolling in school prior to completing grade 12 (20%) suggests that the restricted re-enrollment policy should indeed have a dramatic effect on the incentives faced by many youth.

A final note is in order on the welfare implications of the restricted re-enrollment simulation considered here. Restricted re-enrollment imposes costs on youth, even if it induces some to remain in school longer and complete more schooling than they otherwise would. In the unrestricted model, the possibility of re-enrollment confers option value on agents who have dropped out. Youth who choose to disenroll and re-enroll do so rationally. As forward-looking agents who compare the present value of enrollment and labor market participation in each period, their dropout and re-enrollment decisions are rational responses to shocks such as school failure and unemployment. Any restricted re-enrollment policy reduces youth welfare by shrinking their choice set (as noted in Section 2.2.2). From a social welfare perspective, restricted re-enrollment is efficient only if the resulting reduction in school dropout generates sufficiently high social returns, which I do not model in this paper. This

[41] Similarly, Heckman, LaFontaine and Rodriguez (2008) find, using a panel of U.S. states, that an increase in the difficulty of passing the GED test reduced high school dropout rates.

same caveat about welfare implications applies to other policy simulations as well, in particular to the compulsory schooling simulation considered below.

2.6.2 Compulsory schooling

Schooling is compulsory in South Africa from ages 7 to 15 (or until completion of grade 9, if this occurs before age 15). Child labor laws also prohibit employment by those under age 15. In the data, however, truancy rates in ages 12-15 reach as high as 8%, with child labor rates as high as 21% among truants, as shown in Table 2.A.8. Therefore, the model of this paper assumes that neither the compulsory schooling nor child labor laws are enforced. I use the model to simulate the enforcement or extension of these laws by restricting youth to be enrolled in school until a certain age; this restriction models enforcement of both the compulsory schooling and child labor laws because youth in the model may only work if they are not enrolled in school. As before, I simulate 50 enrollment histories for each observation under both the unrestricted model and the compulsory schooling policy, using the state variable at age 12 as the initial condition. The results of the simulation, in addition to being interesting in their own right, also help to shed additional light on the importance of the option to re-enroll: making schooling compulsory over a period of time also removes the option to drop out and re-enroll during that period.

Table 2.8 presents results of simulations in which I make schooling compulsory through ages 15 and 16 (columns [2] and [3], respectively). Panel (a) shows the percentage of the sample in each schooling category at age 20, while panel (b) reports differences between the compulsory schooling simulation and the unrestricted model. Compulsory schooling leads to a redistribution of youth from the lower schooling categories (less than 12 years) to the higher schooling categories (12 or more years). Enforcing compulsory schooling until age 15 (the existing law), for instance, leads to a decrease of 2.5 percentage points in those completing less than 9 years of schooling, but a 0.5 percentage point increase in those with exactly 12 years of schooling and a 2.6 percentage point increase in those with more than 12 years. The effects are generally larger in magnitude when extending compulsory schooling to age 16.

The results of the compulsory schooling simulation are qualitatively similar to those of the restricted re-enrollment simulation of the preceding section: in each case, the shares of youth in the two highest schooling categories increase. These results are consistent with each other because each policy increases the cost of early dropout (or completely restricts it, in the case of compulsory schooling), changing the incentives for youth who might otherwise end their schooling careers. For both simulations, however, the restrictions come at the cost of reducing youth's choice sets, and therefore their individual welfare.

2.6.3 Increased pass rates

As noted in section 2.4, South African youth face high rates of school failure and grade repetition. Lam et al. (2011) document how grade failure and repetition lead to lengthy school careers in South Africa. One might therefore argue that passing thresholds are too high in South African schools, and that relaxing standards would lead to improved outcomes for youth. The model of this paper is well suited to explore such a policy, because it makes pass rates a key factor that youth consider when making enrollment choices. Using the same methodology as the previous policy simulations, I consider the effects of a 10 percentage point increase in pass rates at all levels of schooling.[42]

Table 2.9 presents results of the increased pass rate simulation. Not surprisingly, the distribution of completed schooling among 20 year-olds shifts to the right relative

[42]Pass rates are set to 100% when the increase would result in rates exceeding 100%.

to the unrestricted model, with those completing 12 years increasing 4.3 percentage points and those with more than 12 years increasing 14 percentage points. The results are qualitatively similar to those from both the restricted re-enrollment and compulsory schooling simulations, showing that increasing the benefits of enrollment (through an increase in school pass rates) has similar effects as increasing the costs of dropout. The increased pass rate policy has important limitations, however, because increasing pass rates without commensurate increases in (unmodeled) school quality may not translate to gains in labor market outcomes.

2.6.4 Post-secondary schooling subsidy

Given the sharp increase in school fees for post-secondary education (Table 2.A.5), another policy simulation I consider is a subsidy for post-secondary education. Specifically, I simulate a 25% reduction in school fees for post-secondary education. As with the previous policy simulations, I report results from simulating 50 enrollment histories for each observation, using the state variable at age 12 as the initial condition, and modifying the school fee term in the enrollment utility function (2.8) to reflect the subsidy.

The post-secondary fee subsidy has negligible effects on the proportion of the sample enrolling in post-secondary school by age 22. Table 2.10 shows that in the full sample, the subsidy increases post-secondary enrollment by 0.3 percentage points, but the increase is not statistically significant. The subsidy has identical effects on youth from the first two quintiles of household incomes, the group that we would expect to respond most to the subsidy. The results suggest that fees do not explain much about post-secondary enrollment patterns, though I must note that because mine is not a formal model of asset accumulation and borrowing, the simulation does not allow me to conclude anything about the presence or absence of credit constraints.

2.6.5 Summary of policy simulations

In this section, I have considered a variety of policies that alter the enrollment incentives of youth: restricted re-enrollment prior to high school completion, enforcement and extension of compulsory schooling, increased pass rates, and a subsidy for post-secondary schooling. Table 2.A.9 presents a summary of results of these policy simulations (omitting the post-secondary fee subsidy, where the focus was on a different outcome). The results from each policy simulation may be compared directly to the others because I use the same random numbers to simulate the model under each policy. The policies have qualitatively similar effects: each policy increases the proportion of 20 year-olds in the sample who complete 12 or more years of schooling, while decreasing the proportion completing less than 12 years, as shown in panel (b). The restricted re-enrollment simulation has the largest effect among the policies considered on those just below high school completion (9-11 years) and those at exactly 12 years of schooling. Restricted re-enrollment also has effects that are generally larger in magnitude across the schooling distribution than both compulsory schooling simulations, suggesting that the option to re-enroll has more influence on enrollment decisions than enforcement and extension of compulsory schooling. The dramatic effect of the re-enrollment restriction demonstrates the importance of the option to re-enroll in the schooling decisions of South African youth.

2.7 Conclusion

In this paper, I quantify the importance of school re-enrollment in the school to work transition of South African youth. Specifically, I formulate a dynamic discrete choice model of the transition between school and work, and estimate it using a panel of young South Africans. The model accounts for uncertainty in schooling and labor market outcomes, and allows for dynamic updating of expected returns based on the outcomes of past choices. The model also allows students to re-enroll in school after a spell in the labor market, a frequent choice by South African youth, but one that is largely ignored in both the literature on South Africa and in prevailing dynamic human capital models. Each of these features of the model – uncertain outcomes, dynamic updating of expectations, and the option to re-enroll – would be missing from standard static human capital models, but are essential to understand the difficult transitions between school and work faced by South African youth.

Structural parameter estimates confirm the hypothesis that youth dynamically update their expectations about the relative returns to enrollment versus labor market participation based on schooling and labor market outcomes. Evidence of such dynamic updating is a key finding of this paper. The model replicates basic patterns of grade advancement, employment and enrollment observed in the data, according to characteristics such as completed schooling and recent and cumulative school failures. However, the model performs less well in predicting dropout and re-enrollment rates because it does not adequately account for switching costs between enrollment and labor force participation. Nonetheless, the model matches the stylized pattern of dropout and re-enrollment throughout the schooling distribution, making the model an appropriate basis for policy simulation.

I use the estimated structural parameters of the model to conduct several policy simulations that alter the incentives for enrollment versus labor market participation. In the first simulation, I restrict the option to re-enroll in school before completing grade 12, making the labor market an absorbing state for those who drop out prior to completing high school. Enrollment rates rise sharply under this restriction on re-enrollment, relative to both the data and results from the unrestricted model. The re-enrollment restriction increases the proportion of the sample completing 12 years of schooling by 8 percentage points. The results show that the option to re-enroll is an important component of the incentives South African youth face when making schooling decisions.

Simulations of the enforcement (and extension) of compulsory schooling and an increase in school pass rates show qualitatively similar effects as the restricted re-enrollment policy: in all cases, the proportion of the sample completing less than 12 years of schooling falls, while the proportion completing 12 or more years increases. Simulation of a 25% fee subsidy for post-secondary education shows no significant effects on the proportion of the sample enrolling in post-secondary schooling, however. The results suggest increasing the opportunity cost of school dropout (as in the restricted re-enrollment and compulsory schooling simulations) or raising the expected benefits of enrollment (as in the increased pass rate simulation) can have dramatic effects on youth enrollment decisions. Youth are less responsive to changes in the direct costs of schooling, however.

These policy simulations illustrate how enrollment behavior changes in the context of the model, and are therefore useful for quantifying the importance of various features of the school to work transition for young South Africans. It is important to remember, however, that each policy may impose important costs on public budgets (in the case of the post-secondary fee subsidy, for instance) or on youth (limiting choice through restricted re-enrollment or compulsory schooling) that are not considered in this paper. Proper policy guidance would require a more comprehensive cost-benefit analysis that builds on the results here.

The model of this paper is quite general, and allows for straightforward extensions such as job search or employment while enrolled, or choice of educational path (e.g., academic versus vocational), that I may explore in future work. Its emphasis on uncertainty, academic and labor market shocks, and the option to re-enroll is appropriate for the South African context, where high grade repetition and unemployment make the school to work transition especially difficult. Nonetheless, the model may be relevant for other contexts as well, not only for other developing countries where youth face similarly difficult circumstances, but also developed countries in which educational choices for dropouts are an increasingly important part of human capital investment. In the United States, for instance, the prevalence of high school dropouts obtaining the GED or mid-career workers enrolling in community or for-profit colleges makes the option to re-enroll an important consideration for students.

The key findings of the model – that youth update their expected returns and enrollment decisions based on past schooling and labor market outcomes; that youth alter their enrollment behavior when the option to re-enroll after dropout is restricted; and that youth respond similarly to other changes in the costs and benefits of enrollment – suggest that a dynamic framework is essential for understanding the school to work transition in South Africa.

2.8 Figures and Tables

Table 2.1: Elements of observable state space X and exclusion restrictions

Variable	**Description**	**enrollment**	**passing**	**wage/ employment**
Time-invariant (X_0):				
$race$	$\mathbb{I}$(race=r) for r=black,coloured	x	x	x
g	$\mathbb{I}$(female)	x	x	x
s_0	$\mathbb{I}$(schooling$\geq$ 6 at $t = 0$)	x	x	x
z_q	$\mathbb{I}$(ability quartile=q)	x	x	x
y_q	$\mathbb{I}$(household income quintile=q) for q =1,2	x		
Time-varying (X_t):				
s	schooling	x	x	x
f_{tot}	total grades failed since $t = 0$	x	x	
f	$\mathbb{I}$(failed last grade enrolled)	x	x	
x_{tot}	work experience			x
x	$\mathbb{I}$(worked last period not enrolled)			x
HSG	$\mathbb{I}$(high school graduate)	x	x	
$(s-12)*HSG$	(schooling-12)*$\mathbb{I}$(high school graduate)	x	x	x
y_q*HSG	interactions between y_q and HSG	x		
ζ	$\mathbb{I}$(bad macro environment)†			x

†A bad macro environment is a year in which the employment/population ratio for 15-24 year olds is less than 25%. Between 1991-2006, this occurred in 2002-2003.
Note: "x" denotes inclusion in equation, by column. Enrollment refers to equation (2.8); passing refers to (2.10); wage/employment refers to (2.7) and (2.9).

Table 2.2: Summary statistics

Variable	**Obs.**	**Mean**	**Std. Dev.**
female	3,806	0.54	0.50
black	3,806	0.28	0.45
coloured	3,806	0.55	0.50
white	3,806	0.17	0.38
age of school entry	3,806	6.2	1.1
age last enrolled	3,806	17.9	2.3
completed schooling	3,806	10.9	2.2
grades failed	3,806	1.2	1.3
schooling at age 12	3,806	5.4	1.2
LNE score	3,806	0.23	0.98
ever re-enrolled	3,806	0.33	0.47
ever re-enrolled, grade 1-12	3,806	0.20	0.40
passed matric (high school)	3,806	0.50	0.50
ever enrolled in post-secondary	3,806	0.27	0.44
ever worked	3,806	0.51	0.50
work experience	3,806	1.1	1.3
wage (FTE, max.)	2,708	40,287	45,002

LNE score is from literacy and numeracy evaluation administered in CAPS Wave 1, normalized to mean zero and standard deviation one. Wage is annual full-time equivalent based on 160 monthly hours, maximum over all individual's observations in panel, in South African rand (real, base year 2002). Survey weights applied.

Table 2.3: Education and labor market outcomes, by selected characteristics

	school entry age	age last enrolled	completed schooling	grades failed	ever worked	ever re-enrolled	ever re-enrolled (pre-12)	wage (FTE, max.)
race								
black	6.9	18.5	10.3	1.7	0.31	0.39	0.31	22,453
coloured	6.0	17.2	10.6	1.2	0.64	0.27	0.16	36,125
white	5.9	19.3	12.8	0.5	0.40	0.43	0.14	70,412
gender								
male	6.3	17.8	10.6	1.4	0.55	0.33	0.20	42,551
female	6.2	18.0	11.1	1.1	0.47	0.34	0.20	38,067
ability								
low	6.5	17.3	9.8	1.7	0.49	0.28	0.22	27,616
high	5.9	18.5	11.9	0.8	0.53	0.39	0.18	51,082
schooling								
< 9	6.9	15.3	7.0	2.2	0.48	0.08	0.07	24,604
9-11	6.4	17.6	10.2	1.8	0.47	0.30	0.30	29,497
12	5.9	18.1	12.0	0.6	0.63	0.37	0.22	42,094
> 12	5.8	20.3	13.7	0.3	0.44	0.54	0.07	65,453

All values are survey-weighted means. Ability based on age-adjusted median from literacy and numeracy evaluation administered in CAPS Wave 1. Ever re-enrolled pre-12 refers to re-enrollment before completing grade 12. Wage is annual full-time equivalent based on 160 monthly hours, maximum over all individual's observations in panel, in South African rand (real, base year 2002).

Table 2.4: Choices and outcomes by completed schooling

	Completed Schooling			
	< 9	$9-11$	12	> 12
enrollment	0.80	0.60	0.24	0.54
failure	0.10	0.24	0.13	0.04
dropout	0.05	0.21	0.58	0.30
re-enrollment	0.04	0.15	0.12	0.26
employed	0.18	0.26	0.40	0.44
wages	19,982	25,557	33,167	59,422
school fee	840	2,412	9,556	12,902

Cells report survey-weighted fraction of relevant person-year observations in each category, based on completed schooling at beginning of year. Enrollment proportions are unconditional. Failure conditions on enrollment. Dropout conditions on enrollment in previous year. Re-enrollment conditions on non-enrollment in previous year. Employment conditions on non-enrollment. Wages conditions on employment. School fee conditions on enrollment. Wage is annual full-time equivalent based on 160 monthly hours, maximum over all individual's observations in panel, in South African rand (real, base year 2002). School fee in South African rand (real, base year 2002).

Table 2.5: Transition matrices between enrollment and non-enrollment

		t					
		Full sample		$<$ 12 years schooling		$\geq$12 years schooling	
		enrolled	not enrolled	enrolled	not enrolled	enrolled	not enrolled
	enrolled	0.82	0.18	0.86	0.14	0.54	0.46
	of which:						
	passed	0.87	0.13	0.93	0.07	0.55	0.45
$t-1$	failed	0.48	0.52	0.49	0.51	0.16	0.84
	not enrolled	0.12	0.88	0.11	0.89	0.15	0.85
	of which:						
	worked	0.05	0.95	0.01	0.99	0.10	0.90
	unemployed	0.15	0.85	0.14	0.86	0.18	0.82

Cells show transitions from state at time $t-1$ to state at time t. Survey-weighted means reported.

Table 2.6: Structural parameter estimates

	(1)	(2)	(3)	(4)
equation	pass	work	enrollment	wage
parameters	ϕ_e	ϕ_j	α	γ
black	-0.30	-0.99	0.21	-3.02
	(0.11)	(0.12)	(0.55)	(0.30)
coloured	-0.55	-0.14	-4.51	-2.40
	(0.10)	(0.12)	(0.49)	(0.31)
female	0.27	-0.37	0.78	-0.55
	(0.04)	(0.05)	(0.24)	(0.10)
$\mathbb{I}$(schooling$\geq$ 6, age 12)	0.58	-0.17	2.26	0.28
	(0.05)	(0.06)	(0.29)	(0.08)
ability quartile 2	0.49	-0.04	0.37	-0.04
	(0.05)	(0.07)	(0.32)	(0.07)
ability quartile 3	0.95	0.10	-0.07	0.17
	(0.06)	(0.08)	(0.35)	(0.09)
ability quartile 4	1.56	0.11	1.85	0.66
	(0.09)	(0.10)	(0.44)	(0.14)
HH income quintile 1			-1.37	
			(0.34)	
HH income quintile 2			-1.43	
			(0.34)	
schooling	-0.39	0.20	-0.23	0.13
	(0.01)	(0.02)	(0.04)	(0.02)
high school graduate	1.19	0.14	-16.31	-0.06
	(0.15)	(0.08)	(1.29)	(0.14)
(schooling-12)*	0.69	-0.09	-4.52	1.76
high school graduate	(0.16)	(0.08)	(1.17)	(0.42)
failed last grade	0.09		-6.73	
enrolled	(0.08)		(0.69)	
cumulative grades failed	-0.60		-3.04	
since age 12	(0.05)		(0.56)	
worked last period		3.08		0.07
not enrolled		(0.08)		(0.14)
work experience		-0.57		0.19
		(0.04)		(0.03)
HH income quintile 1*HSG			-1.28	
			(0.85)	
HH income quintile 2*HSG			-0.25	
			(0.17)	
$\mathbb{I}$(bad macro environment)		-1.61		-0.94
		(0.07)		(0.11)
constant	4.38	-2.33	11.32	3.70
	(0.14)	(0.22)	(0.71)	(0.38)
N	24,679	13,747	3,806	4,370
R^2				0.30
$\ln L$	-736,735.0	-532,639.9	-3,138,189.6	

Unit of observation in columns (1), (2) and (4) is the person-year, from ages 12-26. Unit of observation in column (3) is person. Columns (1), (2) and (3) are logits (column (3) includes full solution to dynamic programming problem); columns (4) is OLS regression. All estimates use survey weights. All standard errors robust. Wages are full-time equivalent, based on 160 hours of work per month, in tens of thousands of real South African rand per year (constant 2002 values). Model units for enrollment equation same as wage equation. Cumulative failure since age 12 top-coded at 3. Work experience top-coded at 4. “Bad macro environment” refers to employment/population ratio for 15-24 year olds below 25 percent.

Table 2.7: Distribution of completed schooling at age 20 under restricted re-enrollment (before high school completion) simulation

Completed schooling	unrestricted (1)	restricted re-enrollment (2)	difference (2)-(1)
<9 years	15.6	11.6	-4.0
	(0.1)	(0.1)	(0.1)
$9-11$ years	46.8	34.2	-12.6
	(0.2)	(0.1)	(0.2)
12 years	18.4	26.7	8.3
	(0.1)	(0.1)	(0.2)
>12 years	19.2	27.5	8.3
	(0.1)	(0.1)	(0.2)

Table shows distribution of completed schooling at age 20 in unrestricted model and under "no re-enrollment before high school completion" simulation. Percentage of sample within each completed schooling category shown (standard error in parenthesis). In simulation, 50 simulated histories are generated for each observation. Column (1) shows simulation results of unrestricted model. Column (2) shows simulation results in which re-enrollment restricted prior to high school completion. Column (3) shows difference is means, Column (1) minus Column (2). All results calculated using survey weights.

Table 2.8: Distribution of completed schooling at age 20 under compulsory schooling simulation

Final compulsory schooling age	unrestricted (1)	age 15 (2)	age 16 (3)
Panel (a): percent of sample in schooling category			
<9 years	15.6	13.1	11.3
	(0.1)	(0.1)	(0.1)
$9-11$ years	46.8	46.3	46.4
	(0.2)	(0.2)	(0.2)
12 years	18.4	18.9	18.7
	(0.1)	(0.1)	(0.1)
>12 years	19.2	21.7	23.7
	(0.1)	(0.1)	(0.1)
Panel (b): Difference (compulsory minus unrestricted)			
<9 years		-2.5	-4.4
		(0.1)	(0.1)
$9-11$ years		-0.5	-0.4
		(0.2)	(0.2)
12 years		0.5	0.3
		(0.2)	(0.2)
>12 years		2.6	4.5
		(0.2)	(0.2)

Table shows distribution of completed schooling at age 20 in unrestricted model and under compulsory schooling simulation. Percentage of sample within each completed schooling category (Panel [a]) or difference between compulsory schooling simulation and unrestricted model (Panel [b]) shown (standard error in parenthesis). In simulation, 50 simulated histories are generated for each observation. Column (1) shows simulation results of unrestricted model. Columns (2)-(3) shows simulation results in which schooling made compulsory through ages 15 and 16, respectively. All results calculated using survey weights.

Table 2.9: Distribution of completed schooling at age 20 under increased pass rate (10 percentage points) simulation

Completed schooling	**unrestricted** (1)	**increased pass rate** (2)	**difference** (2)-(1)
< 9 years	15.6	5.1	-10.5
	(0.1)	(0.1)	(0.1)
$9 - 11$ years	46.8	39.0	-7.8
	(0.2)	(0.1)	(0.2)
12 years	18.4	22.7	4.3
	(0.1)	(0.1)	(0.2)
> 12 years	19.2	33.2	14.0
	(0.1)	(0.2)	(0.2)

Table shows distribution of completed schooling at age 20 in unrestricted model and under increased pass rate simulation. Percentage of sample within each completed schooling category shown (standard error in parenthesis). In simulation, 50 simulated histories are generated for each observation. Column (1) shows simulation results of unrestricted model. Column (2) shows simulation results in which pass rates increased 10 percentage points at all grade levels (or set to 100% if initial pass rate exceeds 90%). Column (3) shows difference is means, Column (1) minus Column (2). All results calculated using survey weights.

Table 2.10: Proportion enrolling in post-secondary schooling by age 22 under post-secondary fee subsidy

	Predicted (no subsidy) (1)	**Predicted (25% subsidy)** (2)	**Difference** (2)-(1)
full sample	44.2	44.4	0.3
	(0.2)	(0.2)	(0.3)
household income quintile			
first (bottom)	24.0	24.4	0.3
	(0.3)	(0.3)	(0.5)
second	27.5	27.9	0.3
	(0.4)	(0.4)	(0.5)

Table shows percentage of sample enrolling in post-secondary education by age 22, by indicated characteristic, based on model simulation. In simulation, 50 simulated histories are generated for each observation. Column (1) shows simulation results of unrestricted model. Column (2) shows simulation results in which post-secondary fees reduced by 25%. Column (3) shows difference is means, Column (2) minus Column (1). All results calculated using survey weights. All results calculated using survey weights.

Table 2.A.1: Panel balance

	N	Proportion censored at age 18	20	22	24
full sample	3,806	0.00	0.27	0.58	0.83
black	1,714	0.00	0.26	0.55	0.81
coloured	1,660	0.00	0.27	0.57	0.82
white	432	0.00	0.31	0.65	0.92

Cells show number of observations (N) or percent of sample with missing enrollment information by age. Survey weights used in calculation.

Table 2.A.2: Mean spell duration from monthly calendar data

Type of spell	**all spells**	**full-time only**	**complete spells only**
school	20.0	14.3	10.2
labor force participation	14.1	14.4	7.2
work	13.8	13.7	7.5

Table shows mean spell duration, in months, by type of spell. Sample monthly calendar reports from estimation sample, CAPS Waves 1-4 (August 2002-December 2006). Labor force participation includes work or search. "Full-time only" refers to spells in which no other concurrent activity was reported. Survey weights used in calculation.

Table 2.A.3: School failure and non-affordability

	Full sample	**Completed schooling**	
		< 12 years	**12+ years**
Panel (a): Reason for failure			
Fail (any reason)	0.14	0.15	0.08
academic failure	0.07	0.08	0.01
withdrew	0.03	0.03	0.03
no grade/continued	0.03	0.04	0.00
unspecified	0.01	0.00	0.05
Panel (b): Can't afford school (new dropouts only)			
Full sample	0.07	0.07	0.06
Household income quintile			
first (bottom)	0.14	0.12	0.22
second	0.09	0.08	0.13
third	0.06	0.05	0.07
fourth	0.04	0.04	0.03
fifth (top)	0.01	0.01	0.01

Cells report survey-weighted fraction of observations in person-year panel, among school enrollees (reasons for failure) and new dropouts (can't afford school). "New" dropouts refers to person-years in which respondent is not enrolled in school after being enrolled in the previous year. "No grade/continued" not considered failure for completed schooling of 12+ years. Household income quintile refers to per capita household income, Wave 1.

Table 2.A.4: Dropout and re-enrollment in response to household shocks

	(1) dropout	(2) dropout	(3) re-enroll	(4) re-enroll
household shock	-0.01		-0.005	
	(0.03)		(0.006)	
household shock$(a-1)$		-0.04		0.003
		(0.08)		(0.009)
failed last grade enrolled	0.28	0.13		
	(0.05)***	(0.11)		
cumulated grades failed	0.11	0.28		
	(0.05)**	(0.13)**		
worked last period of non-enrollment			0.014	0.013
			(0.009)	(0.013)
work experience			0.012	0.001
			(0.007)*	(0.012)
Observations	4771	3019	9515	6215
R-squared	0.77	0.82	0.41	0.58

Sample is person-years for which household shock defined. All regressions include individual, age and schooling fixed effects. “Household shock” refers to death; serious illness or injury; job loss; business failure or bankruptcy; abandonment or divorce; theft, fire or property damage; or other major shock. Contemporaneous and lagged household shock coefficients not included in same regression because they are not separately identified, due to lack of observations with both variables for the same individual in the sample. Standard errors clustered by individual.

Table 2.A.5: Predicted school fees

black	-0.69
	(0.06)
coloured	-0.62
	(0.06)
schooling	0.06
	(0.01)
high school graduate	0.64
	(0.06)
constant	0.24
	(0.10)
N	3138
R^2	0.40

Table shows coefficients in school fee prediction equation. Unit of observation is person-year, where school fees measured in South African rand per year (ten thousands). In enrollment utility (2.8), school fee set to first percentile of school fee distribution if predicted school fee falls below minimum observed school fee in sample. Survey weights used in regression.

Table 2.A.6: Structural parameter estimates (no post-secondary schooling)

	(1)	(2)	(3)	(4)
equation	pass	work	enrollment	wage
parameters	ϕ_e	ϕ_j	α	γ
black	-0.11	-1.06	-3.69	-2.23
	(0.14)	(0.16)	(0.88)	(0.25)
coloured	-0.38	-0.19	-7.13	-1.63
	(0.13)	(0.15)	(0.79)	(0.25)
female	0.26	-0.37	0.19	-0.55
	(0.05)	(0.06)	(0.40)	(0.06)
$\mathbb{I}$(schooling$\geq$ 6, age 12)	0.46	-0.17	2.23	0.23
	(0.05)	(0.06)	(0.31)	(0.07)
ability quartile 2	0.40	-0.18	0.96	0.02
	(0.06)	(0.08)	(0.38)	(0.07)
ability quartile 3	0.75	-0.04	0.54	0.07
	(0.07)	(0.09)	(0.42)	(0.08)
ability quartile 4	1.27	0.10	1.00	0.41
	(0.08)	(0.09)	(0.47)	(0.11)
HH income quintile 1			-1.11	
			(0.47)	
HH income quintile 2			-1.26	
			(0.46)	
schooling	-0.43	0.20	0.05	0.15
	(0.01)	(0.02)	(0.96)	(0.02)
high school graduate		0.01		-0.04
		(0.09)		(0.11)
failed last grade	0.14		-6.59	
enrolled	(0.08)		(1.63)	
cumulative grades failed	-0.54		-3.09	
since age 12	(0.05)		(0.76)	
worked last period		3.17		-0.11
not enrolled		(0.09)		(0.08)
work experience		-0.61		0.20
		(0.04)		(0.03)
$\mathbb{I}$(bad macro environment)		-1.39		-0.84
		(0.08)		(0.07)
constant	4.33	-2.28	12.22	2.73
	(0.17)	(0.24)	(9.74)	(0.30)
N	17,845	11,963	2,979	3,728
R^2				0.21
$\ln L$	-588,726.6	-438,327.3	-1,421,917.0	

Unit of observation in columns (1), (2) and (4) is the person-year, from ages 12-26. Unit of observation in column (3) is person. Columns (1), (2) and (3) are logits (column (3) includes full solution to dynamic programming problem); columns (4) is OLS regression. All estimates use survey weights. All standard errors robust. Wages are full-time equivalent, based on 160 hours of work per month, in tens of thousands of real South African rand per year (constant 2002 values). Model units for enrollment equation same as wage equation. Cumulative failure since age 12 top-coded at 3. Work experience top-coded at 4. “Bad macro environment” refers to employment/population ratio for 15-24 year olds below 25 percent. Sample excludes those who enroll in post-secondary schooling.

Table 2.A.7: Enrollment parameter estimates (robustness checks)

	(1) $\beta = .9$	(2) $T = 52$	(3) $T = 65$
black	-0.58	0.36	0.15
	(0.32)	(0.53)	(0.67)
coloured	-3.18	-4.05	-4.72
	(0.29)	(0.46)	(0.54)
female	0.19	0.68	0.82
	(0.14)	(0.23)	(0.26)
$\mathbb{I}$(schooling$\geq$ 6, age 12)	1.08	1.99	2.38
	(0.17)	(0.28)	(0.34)
ability quartile 2	0.27	0.35	0.38
	(0.19)	(0.31)	(0.34)
ability quartile 3	0.13	-0.08	-0.06
	(0.20)	(0.33)	(0.39)
ability quartile 4	1.31	1.60	1.95
	(0.26)	(0.42)	(0.48)
HH income quintile 1	-0.89	-1.28	-1.42
	(0.20)	(0.33)	(0.37)
HH income quintile 2	-0.89	-1.34	-1.48
	(0.20)	(0.32)	(0.36)
schooling	0.03	-0.15	-0.26
	(0.03)	(0.10)	(0.23)
high school graduate	-11.34	-15.24	-16.79
	(0.70)	(0.88)	(3.60)
(schooling-12)*	-0.52	-3.88	-4.80
high school graduate	(0.71)	(1.09)	(2.26)
failed last grade	-6.01	-6.64	-6.77
enrolled	(0.41)	(0.67)	(1.16)
cumulative grades failed	-1.55	-2.76	-3.16
since age 12	(0.21)	(0.49)	(0.97)
HH income quintile 1*HSG	-0.70	-1.26	-1.29
	(0.56)	(0.77)	(1.13)
HH income quintile 2*HSG	-0.16	-0.36	-0.20
	(0.17)	(0.18)	(0.27)
constant	6.87	10.07	11.87
	(0.42)	(0.73)	(2.88)
N	3,806	3,806	3,806
$\ln L$	-2,193,919.5	-2,936,402.3	-3,229,274.8

Table shows enrollment utility parameter (γ) estimates under modifications to model, as indicated by column. All estimates use survey weights. All standard errors robust. Model units scaled to 10,000 South African rand per year (constant 2002 values). Survey weights used in all estimates.

Table 2.A.8: Truancy and child labor

Age	**Truancy**	**Employment (among truants)**
12	0.01	0.05
13	0.01	0.09
14	0.02	0.21
15	0.08	0.08

Table shows proportion of sample who are truants (not enrolled in school before completing grade 9) and employment rate among truants, by age. Survey weights used in calculation.

Table 2.A.9: Policy simulation results summary

Simulation **completed schooling**	**unrestricted** (1)	**no re-enrollment before high school completion** (2)	**compulsory schooling (age 15)** (3)	**compulsory schooling (age 16)** (4)	**increased pass rate (10 p.p.)** (5)
Panel (a): percent of sample in schooling category					
< 9 years	15.6	11.6	13.1	11.3	5.1
	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)
$9 - 11$ years	46.8	34.2	46.3	46.4	39.0
	(0.2)	(0.1)	(0.2)	(0.2)	(0.1)
12 years	18.4	26.7	18.9	18.7	22.7
	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)
> 12 years	19.2	27.5	21.7	23.7	33.2
	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)
Panel (b): Difference (policy minus unrestricted model)					
< 9 years		-4.0	-2.5	-4.4	-10.5
		(0.1)	(0.1)	(0.1)	(0.1)
$9 - 11$ years		-12.6	-0.5	-0.4	-7.8
		(0.2)	(0.2)	(0.2)	(0.2)
12 years		8.3	0.5	0.3	4.3
		(0.2)	(0.2)	(0.2)	(0.2)
> 12 years		8.3	2.6	4.5	14.0
		(0.2)	(0.2)	(0.2)	(0.2)

Table shows distribution of completed schooling at age 20 in unrestricted model and policy simulations. Percentage of sample within each completed schooling category (Panel [a]) or difference between indicated policy simulation and unrestricted model (Panel [b]) shown (standard error in parenthesis). In simulation, 50 simulated histories are generated for each observation. Column (1) shows simulation results of unrestricted model. Column (2)-(5) shows policy simulation results in which re-enrollment restricted before high school completion (column [2]); school enrollment compulsory until age 15 or 16 (columns [3]-[4], respectively); and grade advancement rates increased 10 percentage points (column [5]). All results calculated using survey weights.

Figure 2.1: Youth employment/population in South Africa and selected countries

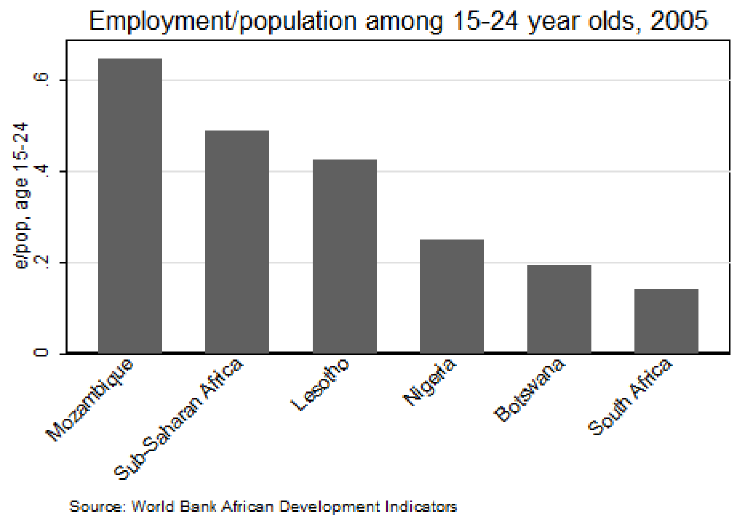

Figure 2.2: K-10 enrollment at age 20 in South Africa and selected countries

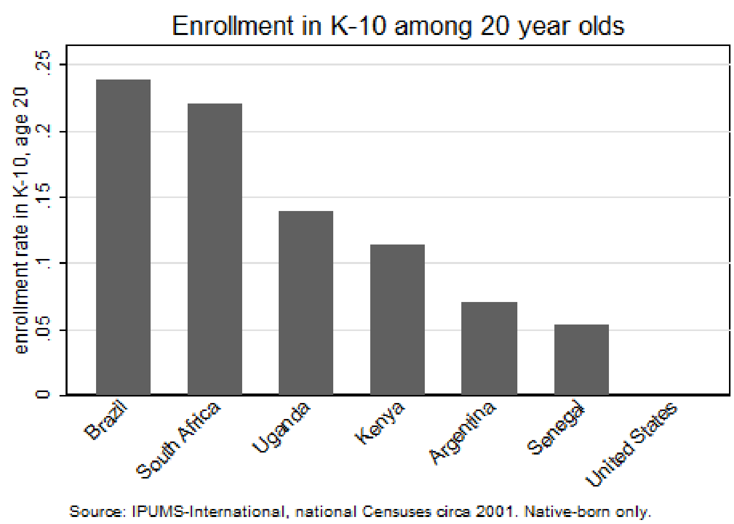

Figure 2.3: Simple model of re-enrollment

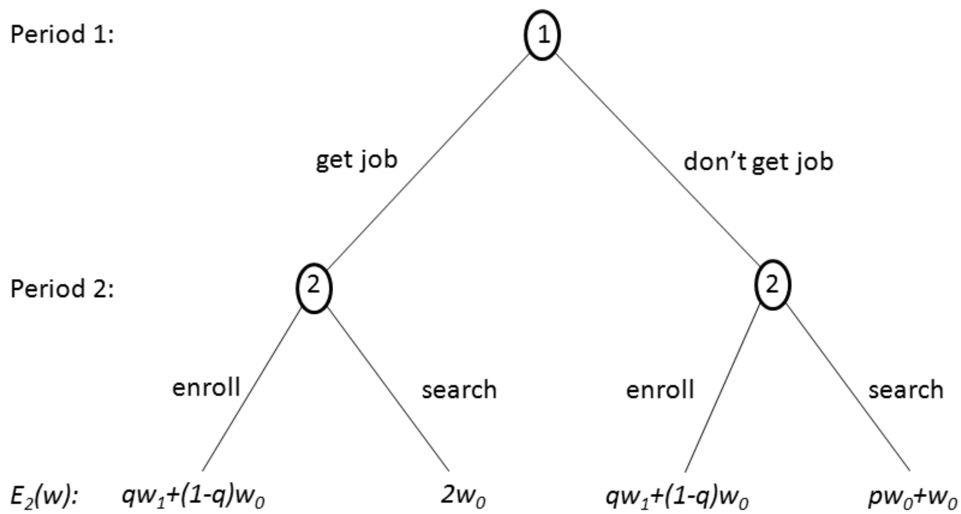

Numeral within each node refers to time period of model. Agents begin in labor force in period 1. Agents then choose to enroll in school or remain in labor force in period 2. All agents are in the labor force in period 3, with guaranteed employment. Payoffs at bottom refer to expected payoffs at outset of period 2.

Figure 2.4: Model timing, choices and payoffs

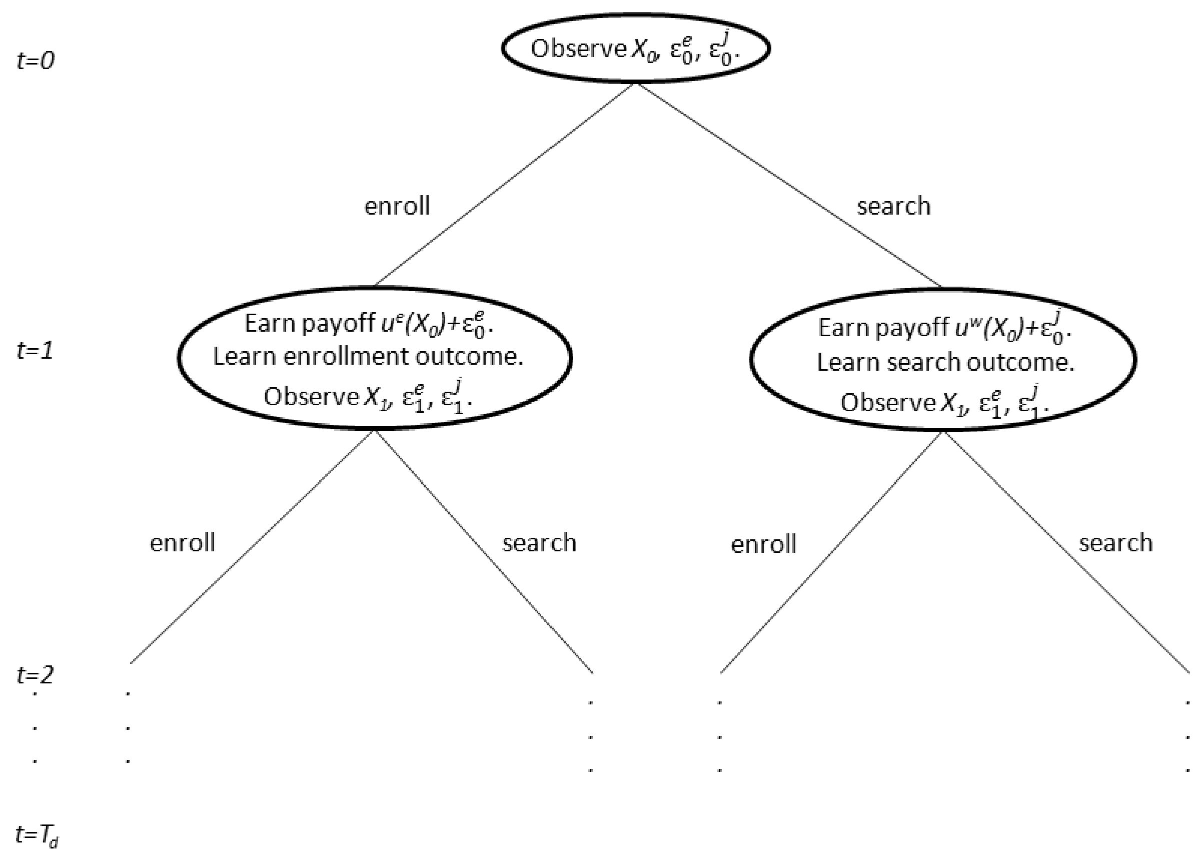

Figure 2.5: Schooling spell birth and death

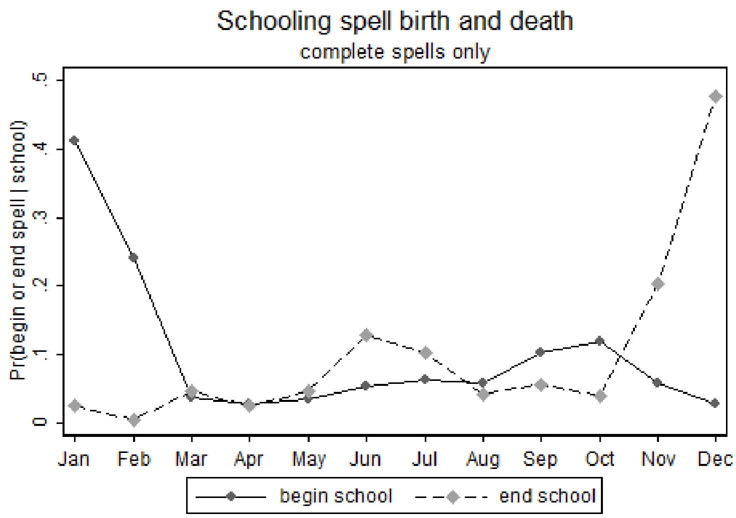

Figure shows proportion of students beginning or ending a schooling spell in each month, using CAPS monthly calendar data, Waves 1-4 (August 2002-December 2006). Censored spells excluded.

Figure 2.6: Failure and employment rates, by last outcome

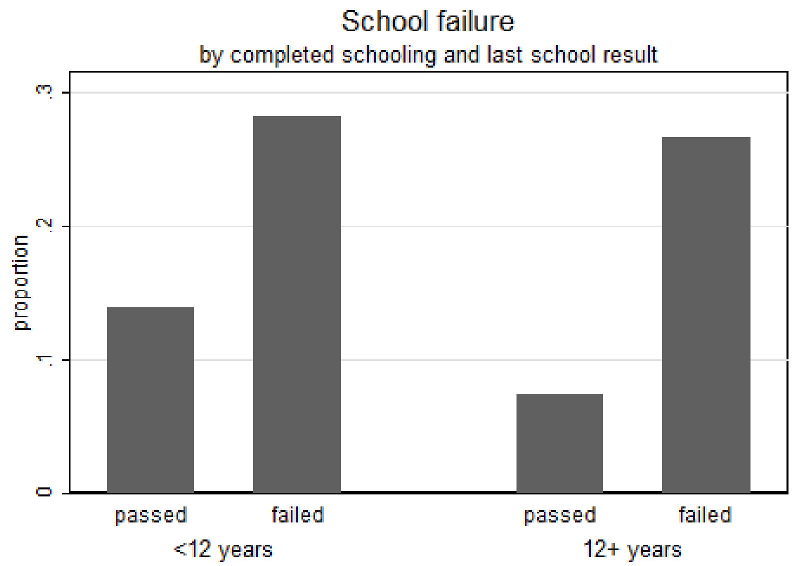

(a)

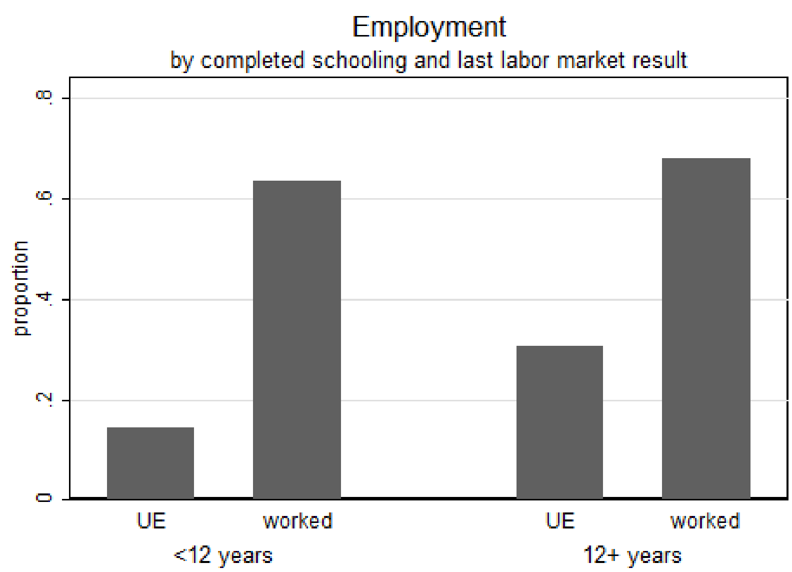

(b)

Panel (a) shows proportion of enrollees (person-year observations) who fail current grade, conditional on schooling outcome of last enrollment period and completed schooling category. Panel (b) shows employment rates among non-enrollees (person-year observations), conditional on labor market outcome of last non-enrollment period and completed schooling category.

Figure 2.7: Model fit: grade advancement and employment

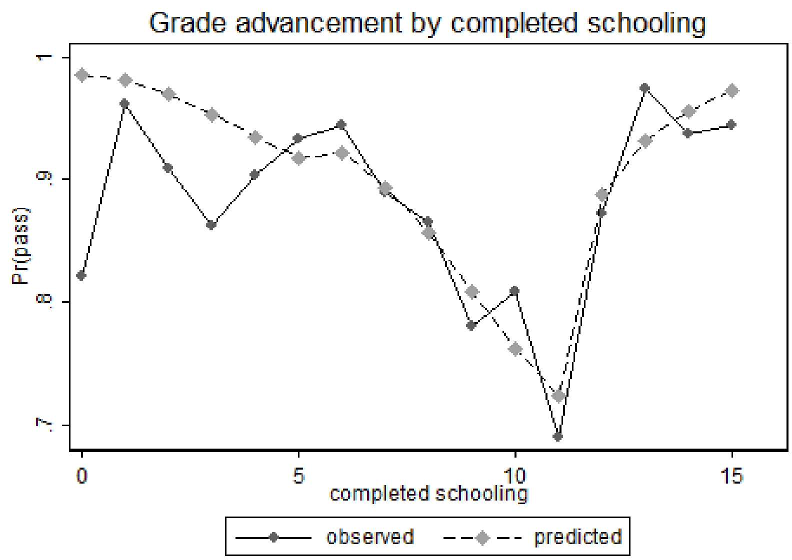

(a)

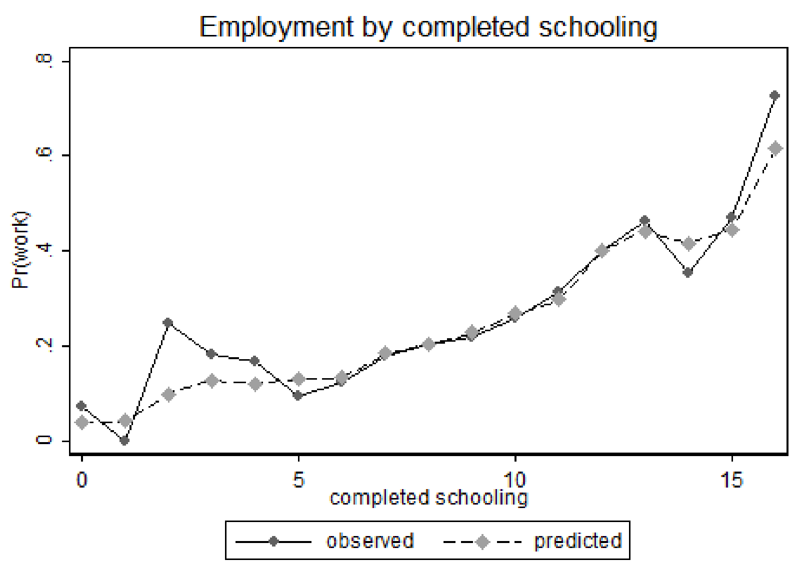

(b)

Figure shows labeled characteristics among person-year observations in sample, conditional on completed schooling, i.e., proportion advancing to next grade among person-year observations with s years completed schooling. Grade advancement rates condition on current school enrollment. Employment rates condition on current non-enrollment. Predicted grade advancement and employment calculated as $\int \Pr(y|S, \hat{\theta}) f(S) dS$ for $y =$ pass,work, using observed distribution of states S and survey weights.

Figure 2.8: Model fit: enrollment

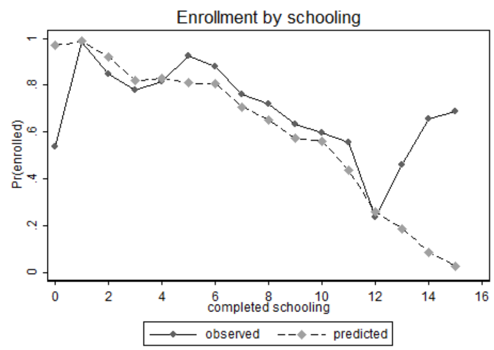

(a)

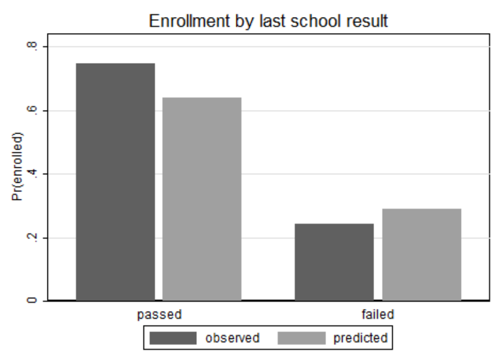

(b)

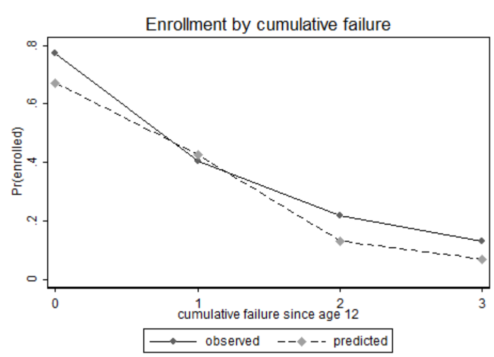

(c)

Figure shows enrollment rates among person-year observations in sample, conditional on labeled characteristics, i.e., proportion enrolled in school among person-year observations with s years completed schooling. Predicted enrollment calculated as $\int \Pr(d|S, \hat{\theta}) f(S) dS$, using observed distribution of states S and survey weights.

Figure 2.9: Model fit: dropout and re-enrollment

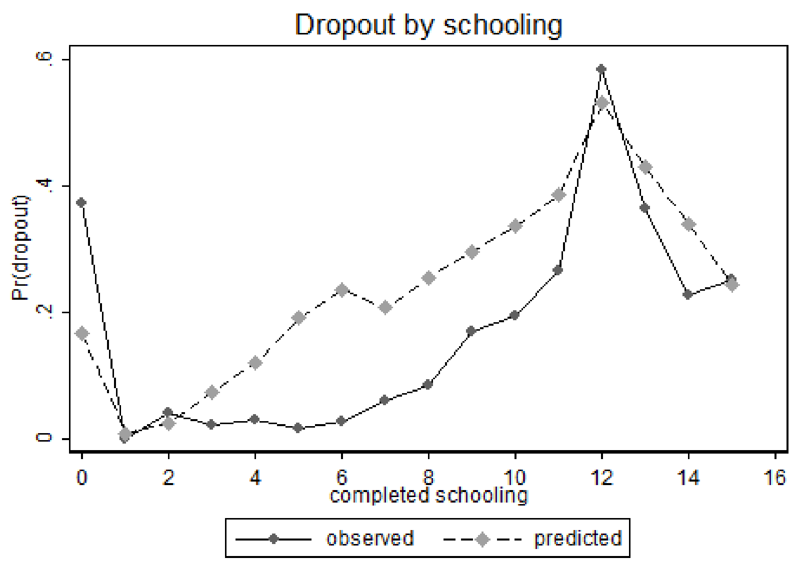

(a)

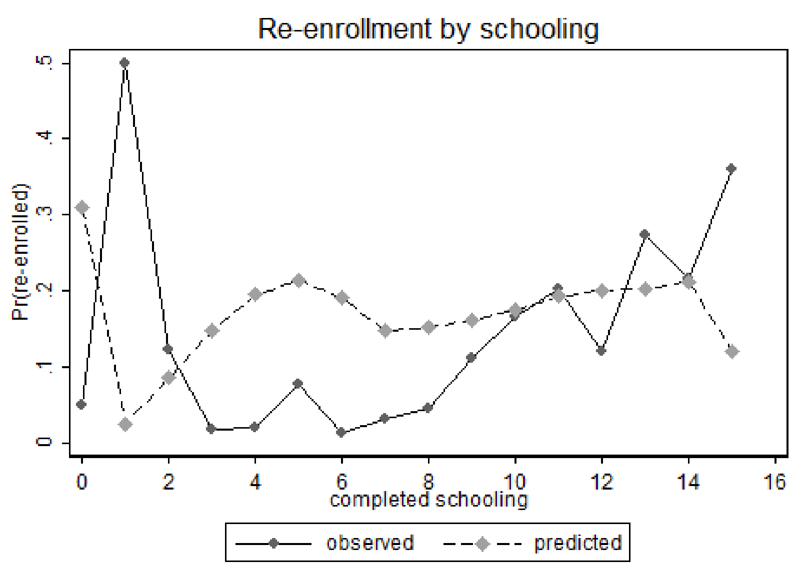

(b)

Figure shows dropout and re-enrollment among eligible person-year observations in sample, conditional on completed schooling. Dropout refers to disenrollment after a period of enrollment, i.e., Pr(not enrolled at $t|$enrolled at $t-1$, schooling $= s$). Re-enrollment refers to enrollment after at least one period of disenrollment, so that figure depicts Pr(enroll at $t|$not enrolled at $t-1$, schooling $= s$). Predicted values calculated as $\int \Pr(y|S,\hat{\theta})f(S)dS$ from simulation in which enrollment history simulated 50 times for each observation in the dataset, using the observed state at $t = 1$ (age 12) as the initial condition. Predicted dropout and re-enrollment are undefined for age 12 because no simulated prior history exists at that age.

Figure 2.10: Restricted re-enrollment simulation

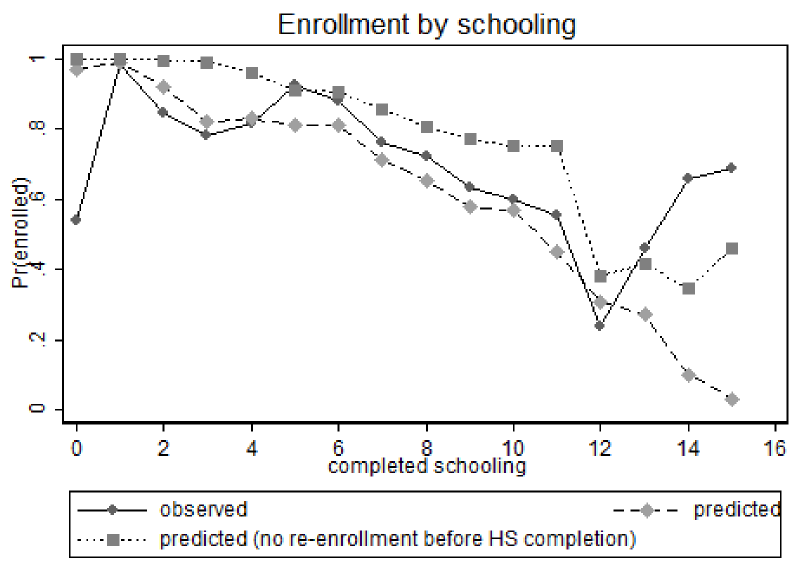

(a)

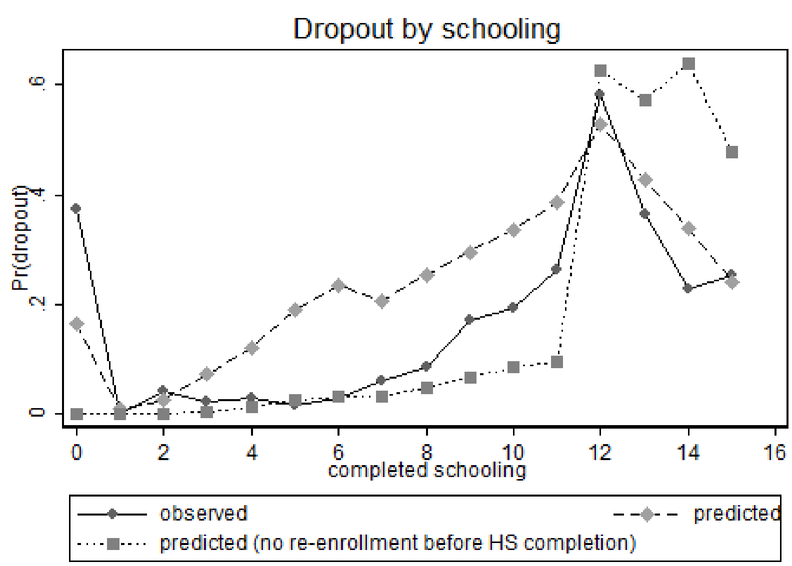

(b)

Figure shows enrollment and dropout among person-year observations in sample, conditional on completed schooling. Dropout refers to disenrollment after a period of enrollment, i.e., Pr(not enrolled at $t|$enrolled at $t-1$, schooling $= s$). Predicted values calculated as $\int \Pr(y|S, \hat{\theta}) f(S) dS$ from simulation in which enrollment history simulated 50 times for each observation in the dataset under “no re-enrollment before high school completion” restriction, using the observed state at $t = 1$ (age 12) as the initial condition. Predicted dropout is undefined for age 12 because no simulated prior history exists at that age.

2.A.1 Why students drop out (in context of re-enrollment model of Section 2.2.2)

The re-enrollment model of Section 2.2.2 conditions on agents already participating in the labor market. Here, I consider how such initial dropout and labor market participation was a rational choice in the context of the model. As in the model, ignore time discounting, assume that all agents have zero skill units ($s = 0$) at $t = 0$, and that one skill unit is the maximum possible. Further assume that agents with $s = 1$ obtain employment at wage w_1 with certainty. Agents with 0 skill units who were not employed in the previous period find a job with probability $p \in (0, 1)$. Let q denote the probability of passing the grade and obtaining one skill unit.

At $t = 0$, an agent will drop out and enter the labor market if:

$$\begin{aligned} \mathbb{E}_0[U(\text{labor market})] > \mathbb{E}_0[U(\text{enroll})] \Leftrightarrow \\ pw_0 + \mathbb{E}_1(\max\{\text{labor market,re-enroll}\}|s = 0) > q\mathbb{E}_1(\text{work}|s = 1) \\ + (1 - q)\mathbb{E}_1(\max\{\text{labor market,re-enroll}\}|s = 0) \end{aligned} \tag{2.22}$$

Denote $\mathbb{E}_1(\max\{\text{labor market,re-enroll}\}|s = 0) \equiv \mathbb{E}max$ to conserve notation. Noting that $\mathbb{E}_1(\text{work}|s = 1) = 2w_1$, we have:

$$\begin{aligned} \mathbb{E}_0[U(\text{labor market})] > \mathbb{E}_0[U(\text{enroll})] \Leftrightarrow \\ pw_0 + \mathbb{E}max > q(2w_1) + (1 - q)\mathbb{E}max \Leftrightarrow \\ pw_0 > q(2w_1 - \mathbb{E}max) \Leftrightarrow \\ \frac{pw_0}{2w_1 - \mathbb{E}max} \equiv \underline{q}_d > q \end{aligned} \tag{2.23}$$

The decision to drop out may be summarized by a threshold value of q, denoted $\underline{q}_d$, such that academic success is sufficiently unlikely relative to the expected gains from working. (This is analogous to analysis of the re-enrollment decision, which also depended on threshold values for q.) Here, the option value of re-enrollment $\mathbb{E}max$ also plays a role: the greater this option value, the greater must be the probability of advancement q to justify remaining enrolled in school.

2.A.2 Education in South Africa

The South African education system consists of General Education (grades 1-9), Further Education (grades 10-12) and Higher Education (grades 13 and above). Education is compulsory for youth aged 7-15 or through completion of grade 9 (whichever comes first), but I treat this regulation as non-binding due to the presence of non-enrolled youth in this age range in the data. To enter post-secondary schooling, students must pass a nationally standardized "matric" exam at the completion of grade 12. Post-secondary education includes both academic and vocational programs ("Technikons"). Additionally, students at the secondary level may enroll in vocational "National Training Certificate" (NTC) programs.[43] For simplicity, I do not distinguish between academic and vocational education in the model. Private schools serve less than 3% of the student population in South Africa, according to government statistics; in the Western Cape province the figure is 3.1% (Fiske and Ladd 2004).

The government subsidizes public education, but students must pay fees to attend, and such fees vary considerably among schools. Public schools are self-governing, and are free to set their own admissions policies and fees (subject to provincial government approval). Although admissions can not discriminate based on race, test scores,

[43] See Appendix 2.A.3 for information on mapping NTC programs to grade levels.

or ability to pay fees, prevailing patterns of residential segregation serve to maintain quality differences among schools. Moreover, despite legal prohibitions, Fiske and Ladd (2004) conclude that “there is little doubt that many schools consider a family’s likely ability to pay their fee when making admissions policy” (p. 143). Although low-income families may qualify for fee exemptions under a policy adopted in 1998, only 2.5% of primary school students and 3.7% of secondary school students receive the exemption (these figures rise to 4.1% and 5.7%, respectively, in historically white schools; Fiske and Ladd 2004). At the post-secondary level, the South African government offers subsidized loans to qualified students who pass a means test through the National Student Financial Aid Scheme (NSFAS), and individual institutions also offer financing. Private banks also offer student loans at market interest rates.

2.A.3 Data Definitions

This section discusses variable definitions and sample selection criteria. The data comes from retrospective life history data collected in Wave 1 of CAPS, augmented with life events recorded in Waves 2-4.[44] The Wave 1 retrospective life histories record events by youth’s age, where age refers to the age at which the event occurred in the case of living arrangements and marriage, and to age at the beginning of the calendar year in the case of enrollment and progression through school, labor force participation, and pregnancy. I follow this convention in mapping Wave 2-4 responses to youth’s age.

Schooling level covers grades 1-16, with National Training Certificate (NTC) I,

[44]The Cape Area Panel Study Waves 1-2-3 were collected between 2002 and 2005 by the University of Cape Town and the University of Michigan, with funding provided by the US National Institute for Child Health and Human Development and the Andrew W. Mellon Foundation. Wave 4 was collected in 2006 by the University of Cape Town, University of Michigan and Princeton University. Major funding for Wave 4 was provided by the National Institute on Aging through a grant to Princeton University, in addition to funding provided by NICHD through the University of Michigan.

II and III mapped to grades 10, 11, and 12, respectively.[45] Students enrolled in university or Technikon programs that include grade 12 are considered enrolled in grade 12. Reporting successful completion of the grade level or reporting enrollment in a higher grade level in a subsequent year is considered passing the level for grades 1-12. Beginning at grade 13, reporting successful completion of the grade level or "no grade/continuing" are considered passing the level, up to a maximum of 16 years completed schooling. This distinction is made because "no grade/continuing" is the modal response for those enrolled in the post-secondary education sector, indicating that most youth are continuing in their programs of higher education, whereas "passing" reports at these levels drop considerably. Unfortunately, I am unable to determine whether students are making satisfactory progress towards degree completion. All other results while enrolled are considered failure. I define "dropout" as disenrollment following a year of enrollment, and "re-enrollment" as enrollment following a year of non-enrollment. Schooling histories in which levels regress with age are re-coded so that such regression can not occur. Grades failed represent the accumulation of periods of enrollment in which the agent did not pass the grade, and therefore may include events such as withdrawal, illness or residential moves rather than outright academic failure.

Labor force participation variables (i.e., work and search) and wages are conditional on non-enrollment at a given age, where reporting enrollment supersedes reports of labor market participation. School fees are conditional on enrollment, and include total household expenditure on fees and other educational expenses, in real rand per year (base year 2002). Wages are full-time annual equivalent based on 160 working hours per month (those reporting monthly hours above 160 are considered

[45] NTC conversion based on coding in CAPS, derived variable *w1h_higrd*.

full-time and do not receive an adjustment). Wages and school fees are available only at the time of the interview, rather than as retrospective histories; predicted values are imputed based on observed characteristics for purposes of estimation. Work experience includes only those periods of simultaneous work and non-enrollment; I exclude work experience while enrolled in school.[46]

I make several sample restrictions. I keep only those observed until at least age 18. Those who report advancing two or more grades in a year, or without continuous information on enrollment, are dropped from the sample. I drop those who report entering school prior to age 4 or exiting school after age 24 (which effectively sets $T_d = 24$ as the decision horizon). I also drop those whose educational histories, by our definitions above, place them with more than 16 years of completed schooling. The restriction on observing a person until at least age 18 accounts for more than 60% of those dropped from the sample.

Other covariates are largely self-explanatory. Ability quartiles refers to in-sample rank of age-adjusted score on the literacy and numeracy evaluation (LNE) administered to all CAPS respondents in Wave 1. Unfortunately, because this ability measure was taken in Wave 1, when the sampled youth were at least age 14, it is not predetermined with respect to enrollment choices in the model, which begins at age 12. I include it, however, because it is a measure common to all in the sample, and therefore helps to distinguish between the role of ability and human capital investment in labor market returns. To mitigate bias in the LNE score due to age differences in Wave 1, I adjust for age as follows: using the estimation sample, I regress the standardized literacy and numeracy evaluation (LNE) score on age and age squared at Wave 1 (when the test was administered) and get predicted residuals.

[46] Work or search while enrolled never exceeds 3% of the sample at any grade level, and never exceeds 2% during grades 1-12.

I then sort observations into quartiles based on these residuals. Household income quintiles are derived from the distribution of household per capita income reported in Wave 1 of CAPS.[47]

I calibrate the wage-age profile for years following the decision horizon (i.e., from periods T_d+1 to T) using the 10% public use micro-sample of the 2001 South African Census. First, I define the estimation sample as native-born residents of urban areas in Western Cape province (which includes metropolitan Cape Town, from which CAPS respondents are drawn) who are ages 25-64, in the labor force, and classify themselves as one of the three major racial groups (white, black, or coloured). I also exclude the self-employed and unpaid workers, leaving a sample of $n = 111,772$. I then predict employment and income by running logit and OLS regressions, respectively, using as controls race and gender dummies, years of schooling, a high school graduate dummy, race-schooling and race-high school graduate interactions, age and age squared.[48] I then create an expected income variable for each observation as the product of these predicted values (i.e., $\hat{\mathbb{E}}(\text{income}) = \hat{\Pr}(\text{work}) \times \widehat{\text{income}}$), and regress expected income on the same set of controls. I save the coefficients on age and age squared from this regression for use in the wage equation of the structural model; the coefficients on age and age squared are 3,695.3 and -45.9, respectively.[49] The macro environment variable is based on the South African employment/population ratio for 15-24 year olds, from the World Bank Africa Development Indicators.

All variables measured in monetary values used in this paper are in real South

[47]Due to non-response, 7% of the sample uses imputed values for household income, based on multiple imputation conducted by CAPS.

[48]The included controls are the largest subset (other than age variables) of the controls used in the structural model that are available in the Census. Note that interactions of schooling and the high school graduate dummy are not included because the maximum years of schooling reported in the Census is 13. I use income rather than wages because the latter are unavailable in the Census. I adjust income to 2002 South African rand to be consistent with the base year of CAPS.

[49]Since I care only about the coefficients, inconsistent estimation of their standard errors due to the generated outcome variable is not problematic in this context.

African rand per year (base year 2002), unless otherwise noted. The South African rand traded at 10.3 per US dollar in August, 2002 when CAPS Wave 1 began.

CHAPTER III

The Role of Reservation Wages in Youth Unemployment in South Africa: A Structural Approach

3.1 Introduction

Unemployment is persistently high in South Africa, and has increased dramatically since the fall of apartheid. According to the standard International Labor Organization (ILO) definition, national unemployment among 16-64 year olds rose from 15.6 percent in 1995 to 26.7 percent in 2005. Using a broader definition, unemployment rose from 28.2 percent to 41.1 percent over the same period.[1] Youth are particularly likely to be unemployed: using the ILO definition, in 2005 the unemployment rate for 16-19 year olds was 56.6 percent, while for 20-24 year olds it was 52.3 percent. These rates far exceed those in developed countries such as the United States, as shown by the trends in employment/population ratios among 15-24 year olds shown in Figure 3.1. The immediate causes of such trends in unemployment are found largely in the substantial increases in labor force participation since the fall of apartheid, which has occurred for almost all groups but particularly among African women. The new entrants tended to be less skilled than those already in the labor force. At the same time that labor supply was increasing, labor demand stagnated,

[1] The ILO definition classifies "working age individuals as being in the labor force if during a week of reference they were employed or wanted to work and were available to start working within a week but also had actively looked for work during the past four weeks...The broader definition...[eliminates] the requirement of having actively searched for a job in order for an individual as to be classified as unemployed." (Banerjee et al. 2008, pp. 6).

particularly for the low-skilled (Banerjee et al. 2008).

Despite broad agreement on these proximate causes of the unemployment increase, its level and persistence remain a puzzle. Standard labor market models would predict that wages should decline to clear the market and reduce unemployment to more reasonable levels than those observed. Although there is evidence that real incomes have fallen since apartheid (Leibbrandt, Levinsohn and McCrary 2010), the failure of unemployment rates to fall frustrates conventional economic wisdom. The observed patterns suggest that there are substantial frictions in the labor market. One hypothesis regarding such frictions is that reservation wages among the unemployed are high relative to offered wages, leading job searchers to reject job offers as unacceptable (or leading firms to adapt by failing to make such offers in the first place). According to this reservation wage hypothesis, the fall of apartheid spurred a climate of increased economic expectations among previously disenfranchised groups, particularly blacks and coloureds. Such heightened labor market expectations for disadvantaged groups coincided with increased human capital investments, resulting in reservation wages that tended to exceed employers' willingness to pay. Thus the reservation wage hypothesis holds that the increase in South African unemployment is largely voluntary, resulting from an influx of workers unwilling to work for the prevailing wages offered by firms.

In this paper, we examine the reservation wage hypothesis by estimating a structural job search model applied to the Cape Area Panel Study (CAPS), a panel dataset of youth in Cape Town with detailed histories of education, job search and labor market behavior, including survey reports of the reservation wage. The structural search approach is appropriate for the context we study because it explicitly models the reservation wage as the optimal individual response to labor market frictions that lead to equilibrium unemployment. Although estimation of a structural search model can not determine whether reservation wages are "too high," as the reservation wage hypothesis contends, it can nonetheless determine what must be true of the model's parameters in order to reconcile the reservation wage reports with observed data on accepted wages and unemployment durations. The resulting structural parameter estimates offer a picture of the labor market consistent with a search model in which agents follow a reservation wage policy. Moreover, the structural approach provides a valid framework in which to conduct policy simulations, allowing us to analyze the effects of an employer wage subsidy for hiring unemployed youth, a policy that has been proposed to alleviate youth unemployment in South Africa (Pauw and Edwards 2006, Levinsohn 2008, Go, Kearney, Korman, Robinson and Thierfelder 2010, Burns 2010).

The data we use are particularly suited to our purpose since they focus on a group (urban youth) with extremely high unemployment rates, and contain survey reports of the reservation wage, which is typically unobserved. We estimate the parameters of a simple search model with this survey data on reservation wages, which allows us to assess the role of reservation wages under the restrictions implied by our model. To our knowledge, this is the first attempt to apply data on reservation wages from a developing country to a structurally estimated search model, and among a handful of studies in the broader job search literature that use survey measures of reservation wages. Our model incorporates measurement error in reported wages and observed heterogeneity in the structural parameters, and makes use of survey data on reservation wages in a novel procedure to recover the full distribution of net search costs in the sample. For comparative purposes, we also estimate our model using alternate reservation wage measures that could be obtained in the absence of survey reports.

We find that inclusion of reservation wage data as an input to our model implies a labor market in which job offers are relatively frequent but at wages that tend to be too low to be accepted. These results stand in stark contrast to those obtained using the traditional method of estimating reservation wages from the accepted wage distribution or by maximum likelihood, which imply less frequent offers that are accepted with higher probability. Our interpretation of the results using reservation wage data is that youth have a high "implicit refusal" rate for jobs they think are available but would find unacceptable, rather than being the recipients of many actual job offers. Using the model's results to estimate individual-specific net search costs provides insights on individual heterogeneity relevant to search behavior, confirming the model's predictions about the relationship between search costs and labor market outcomes. Counterfactual policy simulation of an employer wage subsidy shows that youth increase their reservation wages in response to the subsidy, but by an amount modest enough for the subsidy to increase accepted wages and reduce the probability of lengthy unemployment spells.

To use the terminology of Eckstein and Berg (2007), our model is a standard "classical job search" model. As such, it is a partial equilibrium model, in that it models only the worker's optimal search policy in a dynamic setting, leaving the firm's behavior as exogenous; and it is a "wage posting" model, in that firms post wages which potential workers must either accept or reject (in contrast to "bargaining" models, in which workers and firms bargain over the wage after a match has been made). Flinn and Heckman (1982) provide an extensive discussion of parameter identification in such models. Christensen and Kiefer (1991) present a model of this type that is quite similar to ours, develop its likelihood function, and discuss parameter identification. Our model follows Wolpin (1987) and Eckstein and Wolpin (1995) in its focus on the transition from school to work, and is among the small number of papers (such as Lancaster and Chesher 1983, Lynch 1983, Berg 1990) to use survey data on the reservation wage in a structurally estimated search model.

This paper also is part of a vast literature on unemployment in South Africa. For our purposes, the most relevant is the recent literature on search and reservation wages in Cape Town. Using the CAPS, Lam et al. (2009) document the lengthy unemployment spells faced by Cape Town youth who exit school. Nattrass and Walker (2005) analyze data from the Khayelitsha/Mitchell's Plain (KMP) survey conducted in 2000-2001, which sampled working-age adults from a Cape Town working-class district. They use a reservation wage report similar to that used in this paper, and find that it is generally consistent with other reports of labor search behavior reported in the survey, with the vast majority reporting reservation wages below their predicted wages. They conclude that elevated reservation wages are not a major contributor to adult unemployment in this Cape Town district. Using the same KMP data, Schoer and Leibbrandt (2006) find that several different search strategies prevail in the data. They classify individuals as "non-searchers," "exclusive active searchers," "exclusive passive searchers" and "mixed strategy searchers," and find that observable characteristics have strong correlations with the choice of search strategy. Their results suggest that in Cape Town, search is not a monolithic activity, as most search models imply. We nonetheless model search as a simple process in this paper, though future work may attempt to differentiate between search strategies.

The remainder of the paper is structured as follows: the next section presents the model and discusses its estimation and identification. Section 3.3 describes the data and Section 3.4 presents results. Section 3.5 discusses results from estimation of search costs, and Section 3.6 presents results of the policy simulation of an employer wage subsidy. Section 3.7 concludes.

3.2 Model, Estimation and Identification

3.2.1 Model and Estimation

We consider the infinite-horizon, continuous-time dynamic programming problem of an unemployed worker searching for a job, who faces a known wage offer distribution with cumulative distribution function $F_W(w)$ and Poisson job offer arrival rate q. When unemployed, the searcher's flow value of leisure[2] is b and she/he discounts the future by discount factor δ. If accepted, a job pays constant wage w, but the worker faces an exogenous probability of job separation p. Once rejected, wage offers may not be recalled. The corresponding continuous-time Bellman equations for the value of search and employment (V^s and V^e, respectively) are:

$$(1-\delta)V^s = b + qE[\max\{0, V^e(w') - V^s\}] \tag{3.1}$$

$$(1-\delta)V^e(w) = w + p[V^s - V^e(w)] \tag{3.2}$$

where w' denotes a future draw from F_W. The reservation wage w^* makes the agent indifferent between accepting the job offer and continued search, i.e., it solves: $V^e(w^*) = V^s$. Manipulation of the above Bellman equations lead to the following standard expression for the reservation wage w^*:

$$w^* = b + \frac{q\delta}{(1-\delta)+p}\int_{w^*}^{\infty}(w - w^*)dF_W(w) \tag{3.3}$$

Given values of b, δ, p and the parameters characterizing F_W, one may solve for w^*

[2]The flow value of leisure may also be viewed as the net search cost. In this paper, I will use the terms "flow value of leisure," "net search cost," and "search cost" interchangeably. All refer to the model parameter b.

through policy function iteration using the above.[3,4]

The model implies a joint distribution of accepted wages and unemployment durations, $f(w, d|w \geq w^*)$, which will form the basis of the likelihood function and whose parameters we seek to recover. Since the model assumes that offer arrivals are independent of wage draws, this joint distribution may be factored as the product of the marginal distributions of accepted wages and unemployment durations, leaving us with $f(w, d|w \geq w^*) = f_W(w|w \geq w^*) \times f_D(d|w \geq w^*)$. We consider estimation of each in turn.

According to the model, no agent accepts a wage below the reservation wage, allowing us to use the truncation of the wage distribution from below at w^* to recover the parameters of the wage offer distribution, since $f_W(w|w \geq w^*) = \frac{f_W(w)}{1-F_W(w)}$. In practice, however, wages are measured with error, so that some reported wages may fall below the reservation wage. Suppose classical measurement error, such that $w_o = w + \epsilon$, where w_o denotes observed wages and $\epsilon \sim N(0, \sigma_\epsilon^2)$ is independent of w. Although the support of the measurement error distribution is unbounded, we may bound realized draws of ϵ by noting that no true accepted wage may fall below w^*, i.e., $\Pr(w < w^*) = 0$.[5] Therefore we have:

$$\begin{aligned} w = w_o - \epsilon \geq w^* \Leftrightarrow \\ \epsilon \leq w_o - w^* \equiv \bar{\epsilon} \end{aligned} \tag{3.4}$$

[3]Note that as a partial equilibrium model, we do not model how firm behavior helps to determine F_W in equilibrium. Although this restricts the realism of the model, it allows us to maintain our focus on youth labor supply. Moreover, the leading method for structurally estimating a search model in general equilibrium, the Burdett-Mortensen model (as exemplified by Berg and Ridder 1998) assumes wage offer and accepted wage densities that are increasing in the wage, which is squarely contradicted by our data.

[4]We also do not account for institutional features of the labor market such as minimum wages or union wage-setting. We feel this is justified because several studies have found low enforcement rates of minimum wages in South Africa (Hertz 2005, Yamada 2007, Dinkelman and Ranchhod 2010), and in the CAPS data, only 2% of employed respondents reported being union members (Wave 2). Youth facing these constraints in particular occupations should be able to switch sectors with relative ease.

[5]This approach to bounding the measurement error distribution follows Christensen and Kiefer (1994), although they do not assume that the measurement error is normally distributed, as we do.

The corresponding density of observed wages is:

$$f_W(w_o|w \geq w^*) = \int_{-\infty}^{\bar{\epsilon}} f_W(w_o|w \geq w^*, \epsilon)\phi\left(\frac{\epsilon}{\sigma_\epsilon}\right) d\epsilon \tag{3.5}$$

where $\phi(\cdot)$ is the standard normal density.[6]

Now consider the density of unemployment durations, $f_D(d)$. Under the assumption of Poisson offer arrivals, the hazard rate of unemployment exit, h, is a (constant) product of the offer arrival rate and the probability that a wage draw exceeds the reservation wage, i.e., $h = q(1 - F_W(w^*))$. Accordingly, unemployment durations are distributed exponentially with parameter h, so that $f_D(d) = h\exp(-hd)$. In practice, however, some unemployment spells will be right-censored, so that observed duration $d = \min\{d^*, d_c\}$, where d^* is the true duration and d_c is the duration observed when the spell was censored. Let $c = \mathbb{I}\{d = d_c\}$ be an indicator for censored spells. Then the density of observed unemployment durations, $g_D(d)$, is:

$$g_D(d) = f_D(d)^{1-c}[1 - F_D(d)]^c \tag{3.6}$$

We observe a sample of accepted wages and (possibly right-censored) unemployment durations. By definition, we do not observe accepted wages for those with right-censored durations, and an additional subset of observations with completed unemployment spells may also have missing wage data. Let $m = \{0, 1\}$ be an indicator for missing wage data. Therefore, the vector of observed data for each observation

[6] Allowing instead for measurement error in reservation wages rather than accepted wages would not change the results of our model. To see this, suppose (without loss of generality) that reservation wages are measured with error, such that $w_o^* = w^* - \epsilon$, where w_o^* is the observed reservation wage and ϵ is distributed $N(0, \sigma_\epsilon^2)$, as above. Then we would have:

$$\begin{aligned} w \geq w^* = w_o^* + \epsilon \Leftrightarrow \\ w - w_o^* \equiv \bar{\epsilon} \geq \epsilon \end{aligned}$$

This leads to the same upper bound on ϵ, and thus the same accepted wage density as the case with measurement error in wages. The only difference would arise in the interpretation of the placement of the measurement error, but estimation results would be identical.

is $Y = (w, d, c, m)$, and the corresponding log likelihood function is:[7]

$$L(\theta|Y) = \sum_{i=1}^{N} (1 - m_i) \ln f_W(w_{o_i}|w_i \geq w^*; \theta) + \ln g_D(d_i; \theta) \tag{3.7}$$

We estimate (3.7) using quasi-Newton techniques, with starting values chosen from initial estimates obtained from separate, preliminary estimation of the observed wage and unemployment duration distributions. We parameterize the wage offer distribution as exponential with parameter λ, so that the model parameters estimated by the likelihood function are $\theta = (q, \lambda, \sigma_\epsilon)$.[8] We describe estimation of the reservation wage w^* in the following subsection.

3.2.2 Identification

Identification of the model parameters depends crucially on the reservation wage. In addition to determining the policy function of the theoretical search model, the reservation wage plays a key role in empirical parameter identification in the likelihood function. By providing the truncation point of the accepted wage distribution, the reservation wage, in conjunction with the dispersion of accepted wages around it, serves to identify the underlying wage offer distribution. Additionally, its role in truncating the accepted wage distribution helps to identify the measurement error variance by placing an upper bound on the measurement error for all observed wages. Moreover, by entering into the expression for the hazard rate of unemployment exit, the reservation wage helps to identify the offer arrival rate by reconciling variation in observed unemployment durations with the probability of offer acceptance.

We estimate the preferred version of the model using survey data on the reservation wage, since the main purpose of this paper is to describe the South African

[7] Appendix 3.A.1 describes the derivation and form of the likelihood function in greater detail.

[8] To restrict our estimated parameters to the positive domain, as implied by theory, we actually estimate each parameter as exponentiated functions of observable characteristics, e.g., $q = \exp(\phi' X)$. Note that the parameters (b, δ, p) of the theoretical model are not identified by the likelihood function.

youth labor market as implied by the reservation wage reports. Because the CAPS data we use in this paper has the rare advantage of survey reports of the reservation wage, we use the median reservation wage (within cells defined by included covariates) as model inputs. The median reservation wage, rather than individual reservation wage reports, is used because under the model all agents face identical structural parameters and therefore must have an identical reservation wage.[9]

However, for comparative purposes, we also estimate the model under alternative measures of the reservation wage, and report how results change under each. Under the model assumptions, the minimum accepted wage in the data is a consistent estimator of the reservation wage (Flinn and Heckman 1982). However, under the assumption that wages are measured with error, this estimator will be susceptible to outliers in the left tail of the observed wage distribution, so instead we use the 5th percentile of observed wages, which is also a consistent estimator of the reservation wage (Flinn and Heckman 1982, Eckstein and Berg 2007).[10]

The theoretical model also provides a means to identify the reservation wage in a manner that is fully structural. In doing so, several problems typically arise. The first is the reliance of the reservation wage estimate on the calibration of several model parameters (in particular, b, δ, and p) which are not identified by the likelihood function alone. Moreover, as the truncation point of the accepted wage distribution, the reservation wage may not be estimated by maximum likelihood, because it is a boundary value. However, because our model assumes that measurement error in the reservation wage may lead some observed wages to fall below the reservation wage, the boundary value problem is eliminated, and the reservation wage may indeed be

[9] We could also choose the mean reservation wage or other measure of central tendency, but chose the median because it is less sensitive to outliers. Parameter estimates obtained using mean reservation wages are qualitatively similar to those obtained under the median.

[10] Flinn and Heckman (1982) and Eckstein and Berg (2007) note that any fixed order statistic of the accepted wage distribution consistently estimates w^*.

estimated as an additional model parameter in a conventional maximum likelihood framework.

3.3 Data

We use data from the Cape Area Panel Study (CAPS), a longitudinal study of youth in metropolitan Cape Town, South Africa (Lam et al. 2008). CAPS sampled about 4,800 youth aged 14-22 in Wave 1 (August-December 2002) and currently contains four waves, the most recent conducted in 2006. For our purposes, the most relevant features of the data are its monthly histories (for a period of 52 months from 2002-2006) of education, search and employment activity, as well as its questions on reservation wages. We focus only on those youth who have left school,[11] are observed for at least 12 months in the calendar sample, and have a valid response to the reservation wage question. Additionally, those outside the 1st and 99th percentiles of the accepted wage distribution are dropped to limit the influence of outliers in the estimation.[12] This leaves $N = 1,430$ individuals in the sample. Key variables are described in Appendix 3.A.2.

Table 3.1 presents summary statistics for the full sample. Among the notable features are the high durations and rates of unemployment: mean duration to first job since school exit is nearly 12 months, while 42% of the sample is unemployed for at least one year. Observed search behavior appears low: only 19% of the the time till first job (or censoring) is spent in search, and 35% report never searching since leaving

[11] We define school exit as being out of school for at least 3 consecutive months. In our sample, 6% report returning to school in at least one month after leaving school permanently according to our definition, but none of these have returned to school full-time (i.e., they always report searching or working concurrently with re-enrollment in school). Although a companion paper (Pugatch 2011) finds higher rates of school re-enrollment using CAPS, this paper considers only data from the concurrent waves of the panel, whereas Pugatch (2011) augments the concurrent waves with retrospective life histories. The shorter time series of the panel used in this paper accounts for the lower rate of school re-enrollment.

[12] Estimation results using the untrimmed sample are qualitatively similar to those with trimming for most variants of the model. However, the model using maximum likelihood estimation of the reservation wage produces several coefficients with inconsistent sign using the untrimmed sample.

school. Nonetheless, few youth are returning to school: only 6% report returning to school before obtaining their first job (or censoring), and none returned to school full-time (i.e., all report searching or working concurrently with re-enrollment in school). Of those who find work, most (77%) are employed full-time.[13] Wages and reservation wages are measured in real South African rand per month (base month August 2002, at which time the South African rand/US dollar exchange rate was 10.59).

Table 3.2 breaks down unemployment durations and rates by observable characteristics. The trends follow the expected patterns: unemployment is more prevalent and prolonged for coloureds and blacks, females, the young, and the low-skilled (both in terms of low schooling and low ability). The levels can be quite striking, however, even for the most advantaged groups: 21% of whites and 15% of those with at least some post-secondary education are unemployed for at least one year since school exit, for instance. Another surprising result is the post-school labor market experience of those who report never searching: of this group, only 36% are censored, meaning that the remaining 64% obtain a job, despite reporting to never have searched. This suggests that "search," at least as understood by the survey respondents, is not necessary to obtain employment, and thus many youth who may appear to be non-participants in the labor market may in fact be searching passively, or at least prepared to accept a job should an acceptable offer arrive.[14] Table 3.3, which shows how youth obtained their first job since school exit, provides more supporting evidence for the prevalence of passive search: more than 60% of the sample obtained their first job through informal networks.

Given the high prevalence and duration of unemployment in the sample, the ques-

[13] Our model results are qualitatively similar when excluding part-time workers from the sample.

[14] Our definition of "never searched" excludes those who report obtaining employment immediately after leaving school. Although such youth do not report searching between school exit and employment, we expect that many in fact did actively search for work prior to obtaining work, and therefore exclude them from the "never searched" group so as not to bias results.

tion of what youth are doing with their time after leaving school naturally arises. Table 3.4 seeks to answer this question with data from more recent waves, for which the most post-school observations are available. The overall attrition rate is 12%. Only 0.1% are dead, suggesting fatal illnesses such as AIDS are not immediately afflicting this age group, although 7% do report serious illness. Although only 6% are married and 4% currently pregnant (including males who report their partners as pregnant), 18% are caring for their own children. A large percentage, 78%, remain co-resident with at least one parent, with 18% living in a household with a pensioner, suggesting that many youth may still have access to intra-household resource transfers. Less than 10% engage in unpaid work, suggesting that informal or underground employment does not explain the lack of wage employment in the sample.

Because reservation wage reports will be used in the main version of the model, it is worth pausing to consider the quality of the reservation wage data. Our reservation wage measure is the minimum monthly wage for which the youth reported to be willing to accept full-time work, measured at the latest wave prior to obtaining a job after permanent school exit (or censoring).[15] Table 3.1 shows that 24% of those with completed spells and non-missing wage data report reservation wages that exceed their reported wage; Figure 3.2 is a graphical depiction of the same, with points below the 45-degree line indicating observations for which $w^* > w$. While this is troubling, the model can account for such discrepancies through its estimates of the distribution of measurement error in wages. Table 3.5 presents regressions of the reservation wage on a set of observable characteristics. Although few coefficients are statistically significant, they generally enter with the expected sign: reservation wages are lower among females, blacks and coloureds, who likely face more labor

[15] Appendix 3.A.2 contains additional details on the construction of the reservation wage measure.

market disadvantages than similarly-skilled males and whites; lower (convexly) as a function of age, suggesting that older youth are less patient in their search; higher for the more skilled, as proxied by schooling and ability; higher for those with employed fathers or with co-resident parents, likely due to the greater availability of intra-household transfers; lower for those whose parents want them more strongly to work; and lower for those with their own children in the household, who have greater need to accept paid work. A notable exception is the negative coefficient on pension receipt by a household member, which contradicts the conventional wisdom that availability of pension-related resources increases reservation wages, although the coefficient is significant only at the 10% level. The regression results suggest that, despite some discrepancies between observed wages and reservation wages, the reservation wage data from the survey are generally internally consistent when considering correlations with observable attributes.

A major assumption of our model is the constant arrival rate of wage offers, which (in combination with the assumption that all other structural parameters are time-invariant) implies that the reservation wage is also constant. Given the high prevalence of observations for which the reservation wage report is inconsistent with search theory (i.e., for which $w^* > w$), it is reasonable to wonder whether the reservation wage declines over time. Because CAPS asks about reservation wages in each wave of the panel, we can test this hypothesis by regressing the reported reservation wage on unemployment duration. By including individual fixed effects, we can separate (time-invariant) unobserved heterogeneity from duration dependence; evidence of the latter, in the form of a negative coefficient on unemployment duration, would be evidence against our assumption of constant reservation wages. Table 3.7 presents results. Column (1), which restricts the sample to the first unemployment spell (following permanent school exit) only, has a positive but statistically insignificant coefficient on unemployment duration. Column (2) extends the sample to multiple spells, and finds a positive (and marginally significant) coefficient on unemployment duration. Thus we find no evidence that the reservation wage declines over the course of an unemployment spell, giving us confidence that our assumption of constant reservation wages is plausible.[16]

Finally, we consider the adequacy of our distributional assumptions used to form the likelihood function. Figures 3 and 4 show kernel density estimates of accepted wages and first unemployment spells, respectively; recall that both distributions are assumed exponential for purposes of estimation.[17] Although the empirical distributions from the full sample may mask considerable heterogeneity and thus can not show that our distributional assumptions are correct, observable patterns consistent with the exponential distribution (e.g., monotonically decreasing with a long right tail) will at least suggest that our estimates may fit the data well. The accepted wage distribution (Figure 3.3, panel (a)) does exhibit the left tail mode and long right tail that is characteristic of the exponential distribution; in our model, measurement error may account for the increasing density in the far left tail. The unemployment duration density (for completed spells; Figure 3.3, panel [b]) also exhibits these patterns, and appears to be consistent with our assumption of a constant hazard rate of unemployment exit, in the aggregate.[18]

[16]Moreover, the leading methods for incorporating time-varying reservation wages in structurally estimated search models make unpalatable assumptions: assuming a finite search horizon (as in Wolpin 1987) seems unsuited to youth seeking their first job following school exit, and allowing structural parameters (typically the unemployment benefit, as in Berg 1990) to evolve over time in a known fashion seems at odds with the South African context.

[17]Under exponential wage offers, the density of accepted wages will also be exponential, with a rightward shift of the offer distribution by the amount of the reservation wage.

[18]Although the kernel density is increasing in the far left tail, the empirical mode is 1 month (the minimum allowed, by assumption), so the empirical density does have its mode at the left tail of the distribution.

3.4 Results

3.4.1 Parameter Estimates

Table 3.8 presents estimates of our model, using the median reservation wage (within included covariate cell) from survey reports as the measure of w^*. Observed heterogeneity is incorporated by modeling (the natural log of) each parameter as a linear function of a parsimonious set of covariates: dummies for black, coloured, high school graduate, at least some college, high ability,[19] and previous work experience; the omitted group is low-ability whites with less than a high school education and no previous work experience. The reservation wage is calculated within groups defined by these covariates; for reference, Table 3.6 reports regressions of w^* and w_{q_5} on the covariates. The measurement error variance is estimated as a single parameter for the entire sample, however.[20]

Consider first the results for q, the job offer arrival rate: the "baseline level" reported in the first row is the exponentiated value of the constant term, and may be interpreted as the monthly probability of receiving a job offer for the omitted group.[21] The baseline monthly probability of a job offer is 27%. The reported coefficients on $\ln q$ represent the marginal effect, in log points, on the offer arrival rate. We see that blacks and coloureds face offer arrival rates that are .8 and .4 log points (or approximately 80% and 40%) lower, respectively, than those for whites. High school graduation and post-secondary schooling generate large returns on offer arrivals (coefficients of .48 and .69, respectively), while high ability and previous work experience also increase the offer arrival rate considerably (coefficients of .27 and .37,

[19] We define "high ability" as above the median literacy and numeracy evaluation score within the estimation sample.

[20] Although in principle we could have treated the measurement error as heteroskedastic by allowing its variance to vary according to observable characteristics, in practice the measurement error coefficients were rarely significant in such models, and frequently led to numerical instability in the parameter estimates.

[21] When the estimate exceeds unity, the parameter may also be interpreted as the predicted number of job offers per month.

respectively). The estimates imply that a black, low-ability high school dropout with no previous work experience has a monthly offer probability of just 12%, but that high ability, previous work experience and some college education nearly quadruple this probability, to 46%.

Now consider the results for λ, the wage offer distribution parameter, whose baseline represents the mean (and standard deviation) of the wage offer distribution; coefficients are marginal effects in log points, as before. The estimated baseline wage offer, at 710 rand, is quite low relative to the mean accepted wage of 2,486 rand.[22] Not surprisingly, the model predicts that only 29% of wage offers are accepted.[23] As with the offer arrival rate, the estimates imply considerable labor market disadvantages for black and coloured youth (coefficients -.32 and -.13, respectively). Schooling, ability and previous work experience generate large returns, however, with the coefficient of .73 on previous work experience particularly notable (although this coefficient may be picking up a number of omitted factors that are correlated with experience, such as motivation or access to employment networks). Comparing model estimates again for black, low-ability high school dropouts with no previous work experience to their high ability, college-educated and experienced counterparts, we find that the former face a mean wage offer of 513 rand, while the latter receives offers more than four times as large, at 2,113 rand. The estimated measurement error standard deviation, σ_ϵ, implies that measurement error accounts for 27% of the standard deviation in accepted wages.[24]

Table 3.9 repeats the estimates of Table 3.8, and presents parameter estimates

[22]Such a comparison must be interpreted with caution, however, as the baseline wage offer is for the omitted category of white, low ability high school dropouts without previous work experience, while the mean accepted wage is for the full sample.

[23]We calculate the probability of offer acceptance, $\Pr(w \geq w^*)$, as the mean over the distribution of the full sample, i.e., $\Pr(w \geq w^*) = \int \Pr(w \geq w^*|x) f(x) dx$.

[24]Bound and Krueger (1991) found that measurement error accounts for 18% of the variance in reported annual earnings for men in the US.

for two additional models that vary by the reservation wage used in estimation (as indicated at the top of each column): w^* is the median reservation wage from survey reports; w_{q_5} is the 5th percentile of accepted wages; and w^*_{MLE} leaves the reservation wage as a parameter to be estimated.[25] As in Table 3.8, the reported baseline represents the estimated level for each parameter for the omitted group, while the coefficients represent marginal effects, in log points. Results are qualitatively consistent regardless of the reservation wage used in estimation, with expected signs on all coefficients.

Turning first to results for q, the job offer arrival rate, we see that baseline offer arrivals are estimated to be more frequent under w^* than the other models: a monthly job offer probability of .27, versus .07 and .15 under w_{q_5} and w^*_{MLE}, respectively. Although the differences between the models shrinks for some groups when coefficients are factored in, the generally higher offer arrival rates of column (1) are consistent with higher reservation wages under w^*: youth who face more frequent offers will be more selective about which to accept.

Differences between the models' estimates of λ, the wage offer distribution parameter, are also quite striking. The baseline mean wage offer of 1,445 rand in the model with w_{q_5} (Table 3.9, column 2) is more than double that of the model with w^*. The baseline offer of 899 rand in the model with w^*_{MLE} (column 3), while not nearly as high, still exceeds the baseline under w^* by more than 20%. Again, certain coefficients mitigate these differences somewhat, but the generally lower level of wage offers in the model with w^* comes through clearly in the estimated probabilities of offer acceptance: 29% under w^*, versus 59% and 44% under w_{q_5} and w^*_{MLE}, respectively. Considered in conjunction with the offer arrival rate results, the estimates

[25]In the estimation, w^*_{MLE} is restricted to be $w^* = \bar{w} - \lambda$, corresponding to the truncation of the exponential accepted wage distribution at w^*.

offer a contrasting picture of the labor market: under w^*, wage offers are relatively frequent but low, while under w_{q_5} offers are infrequent but high.

This arrival/wage offer tradeoff is how the model reconciles different reservation wages using the same data on unemployment durations and accepted wages. Accordingly, the probability of offer acceptance ($\Pr(w \geq w^*)$) implied by the models suggest that if youth behave according to their reservation wage reports, they are less than half as likely to accept a wage offer than under w_{q_5}; we will return to this discrepancy and suggest possible explanations shortly. Results for the model with w^*_{MLE} fall somewhere in between the other two, with intermediate offer arrivals and wage offers for most subgroups, as may be expected when we "let the data speak" to find the best fit.

The estimated measurement error standard deviation, σ_ϵ, is greatest in the model with w^* and smallest in the model with w_{q_5}. This is unsurprising: recall that the measurement error parameter serves to reconcile the density of observed wages below the reservation wage, and hence should be largest in the model with w^*, since reservation wages are highest (on average) in that case. Finally, the coefficients on w^*_{MLE} in column (3) follow a qualitatively similar pattern to those on the alternative reservation wage measures presented in Table 3.6. As expected, black and coloured youth have lower reservation wages relative to whites, while reservation wages are increasing in schooling and ability. Interestingly, the negative coefficient on previous work experience suggests that youth who have already engaged in paid work are willing to work for less than their inexperienced peers, although this coefficient is imprecisely estimated.

The relatively frequent offer arrivals and low job acceptance probability in the model with w^* begs the question, "If the South African youth labor market is so bad, why are youth turning down so many jobs?" Our answer is that it is quite unlikely that youth are actually receiving, and refusing, job offers with the frequency implied by our estimates. Instead, we consider it more likely that low-wage jobs are more abundant than the unemployment data may suggest, but such low-wage matches are made infrequently. "Search" is not necessarily an active process for this group, as the 64% of our sample who obtained employment without ever reporting search activity suggests. Thus the high frequency of offer arrivals and refusals we estimate are more likely to represent "implicit refusals" of low-wage offers that are available in principle, but that are not literally made by employers to unemployed youth. The matching costs incurred by both sides may exceed the surplus generated by these low-wage matches.

3.4.2 Model Fit

The structural search model generates predictions for the distributions of unemployment durations and accepted wages, and estimates of these distributions may be compared to their empirical counterparts to assess model fit. Before considering formal tests of model fit, we first offer a more qualitative assessment of how well our estimates account for some features of the data.

Consider first the distribution of unemployment durations till obtaining the first job. Because some durations are right-censored, it will be convenient to work with the survivor function for unemployment, or the probability that an unemployment spell d exceeds some value d_0 (i.e., $S(d_0) = \Pr(d \geq d_0)$). Table 3.10 shows, in column (1), the empirical survivor function at various monthly durations, along with model estimates according to the reservation wage value in columns (2)-(4). Perhaps the most noteworthy aspect of the results is that, beginning at a duration of 24 months, the predicted survivor function weakly exceeds its empirical counterpart for all estimated models. This means that youth are experiencing shorter unemployment spells than our model predicts at the right tail of the distribution.

Now consider the distribution of accepted wages. Recall that by incorporating measurement error in the reported wage, our model estimates the distribution of *observed* accepted wages, which is therefore directly comparable to empirically observed accepted wages. In Table 3.11, we compare this empirical distribution with its estimated counterparts at their respective means, standard deviations, and selected quantiles. All estimated models have mean and standard deviation that fall somewhat below those of the empirical distribution. The reported quantiles suggest that the reason may be the longer right tail of the empirical distribution: the 75th and 90th quantiles of all estimated models are below those of the empirical distribution, and such a longer right tail in the empirical distribution will increase its mean and standard deviation relative to the estimated models.

To test the model formally, we conduct LM tests separately for the unemployment duration and accepted wage distributions of each model.[26] We reject the null hypothesis that the model is correctly specified in all cases. Moreover, no model appears to offer an unambiguously better fit than the others, leaving no clear reason to favor one method of measuring reservation wages over another.

3.5 Search Cost Estimation

The model estimation described in preceding sections used values for the reservation wage defined within each covariate cell; thus, all coloured high school graduates with low ability and previous work experience were assumed to have identical reservation wages, for instance. This is consistent with our structural model, under which agents facing identical structural parameters must have identical reservation

[26]Appendix 3.A.3 describes details of these tests.

wages.[27] However, our data includes survey reports of each individual's reservation wage, which in general do not coincide with the reservation wages used in estimation. One way to reconcile these individual reservation wages with the underlying structural model is to assume that one or more structural parameters faced by the individual, but not included in the likelihood function used for estimation, generated the reported reservation wage. In our model, the agent's flow value of leisure or net search cost (b), discount factor (δ), and probability of job separation (p) determine behavior but do not explicitly enter estimation. We use individual reservation wage reports to shed light on one of these parameters, the net search cost (b).[28] The results allow us to learn about individual heterogeneity in our sample in ways that are (arguably) richer than the standard approach of estimating a mixture distribution (Heckman and Singer (1984)), which requires a finite number of types (typically two or three) for tractable estimation.[29]

We estimate b as follows: for each individual, we use our maximum likelihood estimates of (λ, q); calibrate p according to observed job separations in the data (a monthly rate of approximately .04); choose $\delta = .95$ annually; and then choose $\hat{b}$ to match w^* to the individual's reservation wage report (by inverting the reservation wage function derived from the search model). This generates the distribution of $\hat{b}$ in our sample in a way that makes use of numerous sources of information, including the restrictions of our structural model, the distributions of accepted wages and unemployment durations on which our maximum likelihood estimates are based, and the individual heterogeneity incorporated in each agent's reported reservation

[27] If we used individual reservation wage reports directly in the estimation, we would essentially be estimating the parameters of *individual-specific* accepted wage and unemployment duration distributions using just one observation for each, which is intractable.

[28] We choose b rather than δ or p because we think it the most likely source of individual-specific heterogeneity: reasonable priors allow us to calibrate δ, and p may be calibrated to match data on job separations within our sample.

[29] Note that our approach is possible due only to the availability of reservation wage reports; structurally estimated search models lacking such data would still have to use the Heckman-Singer approach, or some variant, to incorporate unobserved heterogeneity.

wage. To our knowledge, this is the first use of reservation wage data to shed light on individual-specific search costs in this manner in the literature.[30]

We find the distribution of $\hat{b}$ under each variation of reservation wages used in estimation of the model (w^*, w_{q_5} and w^*_{MLE}). We then use these estimates of individual-specific search costs to test the predictions of our model. Specifically, our model predicts that those with lower net search costs (i.e., higher b) will have higher reservation wages, and therefore experience longer unemployment durations and receive higher accepted wages, all else equal. We can test these predictions by regressing these labor market outcomes on our estimates of search costs, while also controlling for the covariates included in our structural estimation. If our estimates of search costs accurately capture aspects of individual heterogeneity relevant to search behavior, then we should see a positive correlation between unemployment durations, the probability of a censored unemployment spell (i.e., the probability of failing to obtain a job by the end of the sample), accepted wages and search costs.

This is (partially) confirmed in Table 3.12, which presents results of regressions of unemployment durations, the censoring indicator and accepted wages on $\hat{b}$, our search cost estimate for each individual (all regressions also include the covariates included in the structural model, and standard errors are bootstrapped to account for sampling variation in $\hat{b}$). We find that the coefficient on $\hat{b}$ is positive in all regressions, as predicted, regardless of the variant of the reservation wage used in the underlying structural estimation (although statistically significant coefficients are obtained only for accepted wages). The results are not very large in magnitude, however: for example, for accepted wages (columns (7)-(9)), an increase of 100 rand in the value of leisure implies just a 5 rand increase in accepted wages. This suggests that search

[30]Eckstein and Wolpin (1995) conduct a conceptually similar exercise, using their structural model to recover search costs after estimating the remaining parameters. However, since they lack individual data on reservation wages, they are limited to using their reservation wage estimates defined within the cells of their model.

costs play a relatively unimportant role in labor market outcomes in our sample when compared with job arrival rates and wage offers. Nonetheless, our procedure to recover individual-specific search costs coincides with our theoretical model, and illustrates the value of using survey data on reservation wages to reveal information on heterogeneity in search behavior that would otherwise remain unobserved.[31]

3.6 Policy Simulation: Employer Wage Subsidy

Because the parameters of the structural model represent the primitives of the search model and are therefore invariant to policy, our model may be used to simulate counterfactual outcomes of various policies. One such policy to consider is an employer wage subsidy, which we may model as an exogenous increase in the mean wage offer. Therefore, a subsidy s to hiring unemployed youth would truncate the wage offer distribution from below at s, leaving all other structural parameters unchanged.[32] One may think of the subsidy as a voucher, with nominal value s, that employers may apply towards a youth's wage. We may then calculate how various features of the model, such as the quantiles of the accepted wage and unemployment duration distributions and the proportion of offers accepted, change from the baseline case to that under the subsidy.

One complication that arises, however, in such simulation is calculation of w^*. Under the search model, a change in the wage offer mean (or any structural parameter) will change w^*, and hence the simulation results will depend crucially on how

[31] Note that such an exercise would not be possible using the Heckman-Singer approach to unobserved heterogeneity, which recovers type-specific structural parameters and type proportions, but can map heterogeneity in parameters to particular observations only in a probabilistic sense. Our procedure, by contrast, uses survey data on reservation wages to map heterogeneity in search costs to individuals in the sample, and thus allow for more severe tests of our model predictions.

[32] Note that in our partial equilibrium framework, we do not model any effect the wage subsidy may have on the frequency of offers or on the destruction of jobs. Moreover, by assuming that the wage offer distribution becomes truncated below by s, we implicitly assume that the subsidy is fully passed through to job seekers in the form of wage offers, which would generally not be the case if employers have market power in the youth labor market. In this sense, our simulation results present a best-case scenario of the effect of the subsidy on employee welfare.

the model accounts for the agent's updated w^* in response to the policy change. When w^* is estimated structurally, the approach is straightforward: merely update the structural estimate of w^* under the new wage offer distribution. However, when w^* is estimated from the data, we must update w^* by calibrating some elements of θ that we did not observe nor estimate in our baseline specification. In our simulation, we update w^* in the same fashion as in estimation of the search cost distribution described in the previous section. That is, we calibrate the model parameters not estimated by our model (b,δ,p) such that they reproduce the value of w^* used in the baseline estimation. As in the previous section in which we estimated search costs, we use our maximum likelihood estimates of (λ, q); calibrate p according to observed job separations in the data; choose $\delta = .95$ annually; and then choose b to match w^* to the data (by inverting the reservation wage function, as described in the previous section). We then update w^* by varying the subsidy value s, holding all other parameters fixed.

Figures 3.4 shows reservation wages and mean accepted wages, respectively, under a range of employer wage subsidy values.[33] The subsidy $s = 0$ corresponds to the baseline estimates discussed in the preceding sections, and s increases to 1,000 rand in increments of 100 along the horizontal axis. The figures show that both reservation wages and mean accepted wages increase (approximately) linearly in the amount of the subsidy, by about 60 rand per 100 rand increment in s,[34] showing that the benefits (in terms of increased mean accepted wages) of the subsidy recover only about 60% of its costs.[35] Reservation wages are uniformly greatest in the model with w^*, while

[33] In Figures 5-9, the lines labeled `wrhat=wr` correspond to the model estimated with w^*; `wrhat=wp5` to w_{q_5}; and `wrhat=wrmle` to w^*_{MLE}.

[34] This equality is a consequence of the assumption of exponential wage offers, because the corresponding accepted wage distribution is shifted to the right by exactly the reservation wage.

[35] This ignores any benefits of the subsidy on reduced employment durations, which are considered later in this section.

reservation wages in the model with w^*_{MLE} are the next greatest. Results for mean accepted wages as a function of the subsidy (Figure 3.4, panel [b]) show a similar linear increase across all models.

The greater selectivity of youth in the model with w^* is also shown in Figure 3.5, which plots the probability of wage offer acceptance (i.e., $\Pr(w \geq w^*)$) for each model. The probability of wage offer acceptance is nearly one half as low under w^* than under w_{q_5} for all subsidy values considered. Moreover, as the subsidy grows from 0 to 1,000 rand, the acceptance probability under w^* increases by only about 15 percentage points, while in the other models it grows by 20 percentage points or more.

Finally, Figure 3.6 plots the unemployment survivor function, or the probability that a youth experiences an unemployment spell of a given duration, for spells of 12 and 24 months. The figure shows that the probability of such lengthy unemployment spells is lowest in the model with w^* for all subsidy values, due to the higher offer arrival rates under that model. Moreover, the subsidy causes the likelihood of lengthy unemployment spells to fall the most in the model with w^* compared to the other models. Overall, the subsidy appears quite effective at reducing lengthy unemployment spells; the probability of experiencing an unemployment spell of at least 12 months decreases by a high of 15 percentage points in the model with w^*, and a low of 10 percentage points under w_{q_5}, as the subsidy increases from 0 to 1000 rand. Whether such a reduction produces 1,000 rand in social benefits (or at least enough social benefit when paired with increased accepted wages to exceed costs) is unclear and requires more formal analysis, however.

Our simulation of an employer wage subsidy shows that youth respond to the increased opportunities resulting from the subsidy by raising their reservation wages.

However, the reservation wage increases are modest enough for the subsidy to have beneficial effects on accepted wages and unemployment durations. It is unclear, however, whether these benefits exceeds the subsidy's costs.

3.7 Conclusion

In this paper, we have presented a simple, standard search model that incorporates observed heterogeneity and measurement error in wages to explain the role of reservation wages in the high observed unemployment rates and durations among Cape Town youth. Using data on accepted wages and unemployment durations for school leavers who found their first job, we estimated the parameters of the model using survey reports of the reservation wage, as well as alternate measures including the 5th percentile of accepted wage offers and maximum likelihood estimation, for comparative purposes. Results using survey data on reservation wages suggested that searchers received job offers frequently, but at wages that were typically unacceptably low. In contrast, results using the 5th percentile of observed accepted wage offers and maximum likelihood estimation suggested less frequent offers, but a higher probability of offer acceptance. Accounting for observed heterogeneity revealed that, as expected, the frequency and quality of labor market opportunities are generally worse for disadvantaged groups, such as blacks, coloureds and the less skilled.

We used the results of the model, in combination with individual reservation wage reports, to estimate the full distribution of search costs in the sample. Correlations between our estimates of individual-specific search costs and labor market outcomes confirmed our model's predictions, at least with respect to accepted wages. Thus our model allows for insights into individual-specific heterogeneity relevant to search behavior that may not be inferred from the data alone, nor may it be captured in the standard approach of estimating a mixture distribution over unobserved types.

Finally, in a policy simulation of the effect of an employer wage subsidy, we found that although the subsidy has the unsurprising effect of increasing reservation wages, it nonetheless may have substantial positive benefits on accepted wages and unemployment durations. However, because we have assumed that firms will pass the subsidy along in full to employees, such positive effects may be considered an upper bound. A more complete model of firm response to the wage subsidy may find less beneficial effects for youth job seekers.

Returning to our initial motivating inquiry on the role of reservation wages in Cape Town youth unemployment, we found that implied wage offer acceptance rates are indeed substantially lower under the survey reports of the reservation wage than alternative measures. However, to reconcile these low acceptance probabilities with the observed data, the model estimates a correspondingly lower average wage offer. Moreover, if youth behave according to our model and their stated reservation wages, offers appear to arrive with much greater frequency than under alternative measures. The true role of reservation wages therefore depends on which picture of the Cape Town youth labor market–frequent but low offers, versus infrequent but high offers–is more accurate. While the latter picture is consistent with popular perception and is the one that would emerge from the data in the absence of reservation wage reports, the availability of reservation wage data allows us to suggest an alternative view of the youth labor market that is equally consistent with search theory. Although the high frequency of (relatively low) wage offers implied by our estimates may not literally be occurring in the South African youth labor market, our results are consistent with a labor market that is inefficient at matching employers to employees, leading to a high "implicit refusal" rate.

Given the simplicity of our model in its current form, there is much scope for further work. For instance, our model conditions on exit from school, when in fact this decision may also be viewed in light of dynamic optimization. A companion paper (Pugatch 2011) endogenizes the decision to exit school and enter the labor market that we model in this paper.

3.8 Figures and Tables

Table 3.1: Summary statistics

Variable	N	Mean	Std. Dev.	Min	Max
female	1430	0.53	0.50	0	1
black	1430	0.26	0.44	0	1
coloured	1430	0.62	0.49	0	1
white	1430	0.12	0.32	0	1
age	1430	19.5	2.1	14	26
schooling	1430	10.7	2.1	0	16
ability score	1430	0.18	0.91	-2.97	2.01
wage	977	2486.4	1859.9	346.6	11642.3
reservation wage	1430	1594.2	1801.8	48.7	36645.8
$\mathbb{I}(w^* > w)$	977	0.24	0.43	0	1
first UE spell	1430	11.7	11.2	1	50
UE spell$\geq$1yr	1430	0.42	0.49	0	1
censor	1430	0.24	0.43	0	1
previously worked	1430	0.34	0.48	0	1
full-time	1027	0.77	0.42	0	1
never searched	1430	0.35	0.48	0	1
return to school (ft)	1430	0.00	0.00	0	0
return to school	1430	0.06	0.23	0	1

Sample is youth who have left school (absent at least 3 consecutive months after attending school at least one month in calendar sample), observed for at least 12 months in calendar sample after school exit, and with valid reservation wage data. Age and schooling measured at time of school exit. Ability score is z-score from literacy and numeracy evaluation administered in Wave 1. Wage is first reported wage following completion of first unemployment spell. Reservation wage is last reported reservation wage before first completed unemployment spell or censoring. Observations below 1st percentile and above 99th percentile of accepted wages dropped. Wages and reservation wages in real rand per month, base month August 2002 (South African rand/US dollar exchange rate at base=10.59). $\mathbb{I}(wr > w)$ is indicator that reservation wage exceeds reported accepted wage. Previously worked refers to work experience in calendar history prior to school exit. Full-time is average of at least 35 hours per week of work in last month. Never searched excludes those who obtain employment immediately after school exit. Statistics calculated using sample weights ($weightyr$).

Table 3.2: Unemployment, by observable characteristics

	First UE spell	UE spell$\geq$1yr	UE spell$\geq$2yrs	UE, month 12	censored
male	10.2	0.35	0.23	0.47	0.19
female	13.0	0.49	0.34	0.56	0.28
African	17.2	0.66	0.52	0.72	0.38
coloured	10.2	0.36	0.20	0.48	0.20
white	7.7	0.21	0.24	0.28	0.14
age:					
$\leq$18	13.9	0.50	0.35	0.59	0.33
19-22	10.9	0.39	0.25	0.50	0.20
$\geq$23	7.4	0.27	0.18	0.34	0.11
schooling:					
$\leq$9	16.3	0.59	0.43	0.70	0.38
10 or 11	12.7	0.48	0.28	0.55	0.28
12	9.2	0.32	0.19	0.42	0.15
>12	5.0	0.15	0.15	0.29	0.07
low ability	14.3	0.54	0.37	0.63	0.31
high ability	8.7	0.29	0.19	0.39	0.16
previously worked	15.1	0.57	0.41	0.66	0.37
never worked before	5.2	0.14	0.05	0.25	0.00
some search	10.2	0.36	0.26	0.47	0.18
never searched	14.5	0.55	0.33	0.61	0.36

Age and schooling measured at time of school exit. "Low" and "high" ability refer to below and above within-sample median literacy and numeracy evaluation score. "Some search" is reported search in at least one month prior to completion of first UE spell or censoring. "Previously worked" means work experience reported in calendar history prior to school exit. Never searched excludes those who obtain employment immediately after school exit. First unemployment spell measured in months; all other statistics are means of indicator variables. "UE, month 12" refers to employment at month 12 following school exit. All statistics weighted by sample weights.

Table 3.3: How obtained first job since school exit

	Full sample	Black	Coloured	White
informal network (household)	13.6	12.1	15.2	7.4
informal network (non-household)	46.5	49.6	47.2	36.3
formal	30.7	32.0	28.8	39.2
past work for firm	1.7	1.0	2.2	0.0
self-employed/family	3.9	3.3	3.8	5.4
other/don't know	3.6	2.1	2.7	11.7

Table shows method of obtaining first job, proportion by race. Sample weights used in calculation.

Table 3.4: Post-school activities

Variable	*N*	**Mean**
Wave 4 (2006)		
dead	1430	0.001
moved	1430	0.04
attrited	1430	0.12
married	1278	0.06
pregnant (inc. males)	1278	0.04
own child in HH	1278	0.18
live with at least one parent	1278	0.78
pension recipient in HH	1278	0.17
Wave 3 (2005)		
seriously ill	1170	0.07
unpaid work	1170	0.09

Variables for each wave calculated only for those who had left school by time of interview. "Pregnant" includes males who report their partners as pregnant. "Seriously ill" refers to self-reported inability to perform normal activities. All statistics calculated using sample weights.

Table 3.5: Reservation wage regressions

	(1) w_i^*	(2) w_i^*
female	-89.8	-102.9
	(107.2)	(114.9)
black	-754.3	-827.7
	(244.3)***	(233.8)***
coloured	-507.4	-449.6
	(241.3)**	(247.5)*
age	-109.6	-63.5
	(183.9)	(176.1)
age^2	3.9	3.2
	(4.7)	(4.5)
schooling	90.3	93.8
	(31.9)***	(31.0)***
ability score	281.9	303.8
	(74.4)***	(75.9)***
pensioner in HH		-181.1
		(106.0)*
father employed		69.1
		(128.5)
ill		117.4
		(190.1)
parents want youth to work		-79.9
		(25.4)***
co-resident with parent		180.8
		(79.0)**
own child in HH		-274.1
		(138.5)**
N	1430	1430
R^2	0.09	0.13

Robust standard errors in parentheses: * significant at 10%; ** significant at 5%; *** significant at 1%. Reservation wage w_i^* is individual-specific survey report, as defined in Appendix 3.A.2. Age and schooling measured at time of school exit. Pensioner in HH, father employed, ill, parents want to work, co-resident with parent, and own child in hh variables measured at time of reservation wage, where reservation wage is last report prior to job acceptance or end of calendar sample. "Ill" refers to self-reported illness that prevents normal activities. "Parents want youth to work" measured on self-reported 1-5 scale, with 5 being strongest. All regressions include fixed effects for wave at which w^* measured.

Table 3.6: Regressions of w^* and w_{q_5} on covariates used in model estimation

	(1) w^*	(2) w_{q_5}
constant	1575.1	1279.0
	(33.0)***	(122.8)***
black	-797.4	-814.5
	(33.9)***	(126.0)***
coloured	-592.5	-700.2
	(31.4)***	(104.6)***
HS grad	318.1	251.9
	(10.2)***	(35.9)***
at least some college	628.8	615.2
	(31.1)***	(59.1)***
high ability	264.7	78.8
	(10.5)***	(50.1)
previously worked	-119.6	51.7
	(14.0)***	(39.1)
N	1430	1423
R^2	0.88	0.57

Robust standard errors in parentheses: * significant at 10%; ** significant at 5%; *** significant at 1%. w^* is median reservation wage by cell defined by included covariates. w_{q_5} is 5th percentile of accepted wages, by cell defined by covariates. "High ability" is indicator for above median literacy and numeracy evaluation score within sample. "Previously worked" means work experience reported in calendar history prior to school exit.

Table 3.7: Regressions of w^* on unemployment duration

	(1) w^*	(2) w^*
unemployment duration	140.2	56.7
	(89.1)	(33.5)*
N	1126	1582
R^2	0.64	0.64
Individual fixed effects	x	x
Wave fixed effects	x	x
First UE spell only	x	
Spell fixed effects		x

Robust standard errors in parentheses, clustered by individual: * significant at 10%; ** significant at 5%; *** significant at 1%. Sample is person-years from estimation sample. Reservation wage measured in (real) rand per month; unemployment duration in months. All regressions use survey weights and include individual and wave fixed effects. Regressions including multiple spells include fixed effects for spell number.

Table 3.8: Parameter estimates, using reservation wage survey reports

Parameter	$\ln q$ **(offer arrival rate)**	$\ln \lambda$ **(wage offer parameter)**	$\ln \sigma_\epsilon$ **(measurement error s.d.)**
baseline level	0.27	710.58	495.11
constant	-1.30	6.57	6.20
	(0.29)	(0.17)	(0.05)
black	-0.80	-0.32	
	(0.32)	(0.19)	
coloured	-0.40	-0.13	
	(0.26)	(0.15)	
HS grad	0.48	0.27	
	(0.13)	(0.07)	
at least some college	0.69	0.54	
	(0.19)	(0.11)	
high ability	0.27	0.15	
	(0.12)	(0.07)	
previous work	0.37	0.73	
	(0.12)	(0.09)	
N		1430	
$\ln L$		-1,055,884	
$\Pr(w \geq w^*)$		0.29	
σ_ϵ (measurement error s.d.) as percentage of observed accepted wage s.d.		0.27	

Robust standard errors in parentheses. Estimation is by maximum likelihood, with reservation wage as median reservation wage from survey within covariate cell. Starting values taken from converged estimates of sequential estimation of wage offer and unemployment duration distributions. Optimization algorithm alternates between BFGS and BHHH. "Baseline level" refers to value of exponentiated constant term for each parameter, and may be interpreted as parameter level for left-out category (white high school dropouts of low ability, with no previous work experience). $\Pr(w \geq w^*)$ calculated as mean over distribution of full sample, i.e., $\Pr(w \geq w^*) = \int \Pr(w \geq w^*|x) f(x) dx$.

Table 3.9: Parameter estimates, using alternate reservation wage measures

	(1)	(2)	(3)
Reservation wage	w^*	w_{q_5}	w^*_{MLE}
$\ln q$ **(offer arrival rate): baseline**	0.27	0.07	0.15
constant	-1.30	-2.64	-1.88
	(0.29)	(0.20)	(0.24)
black	-0.80	-0.51	-0.74
	(0.32)	(0.20)	(0.23)
coloured	-0.40	-0.12	-0.33
	(0.26)	(0.18)	(0.19)
HS grad	0.48	0.54	0.43
	(0.13)	(0.09)	(0.13)
at least some college	0.69	0.92	0.73
	(0.19)	(0.18)	(0.20)
high ability	0.27	0.25	0.10
	(0.12)	(0.10)	(0.13)
previous work	0.37	1.13	0.78
	(0.12)	(0.09)	(0.12)
$\ln \lambda$ **(wage offer parameter): baseline**	710.58	1445.88	899.51
constant	6.57	7.28	6.80
	(0.17)	(0.15)	(0.12)
black	-0.32	-0.53	-0.33
	(0.19)	(0.16)	(0.13)
coloured	-0.13	-0.34	-0.17
	(0.15)	(0.13)	(0.10)
HS grad	0.27	0.22	0.27
	(0.07)	(0.07)	(0.08)
at least some college	0.54	0.49	0.55
	(0.11)	(0.11)	(0.12)
high ability	0.15	0.11	0.20
	(0.07)	(0.07)	(0.09)
previous work	0.73	0.41	0.61
	(0.09)	(0.07)	(0.09)
$\ln \sigma_\epsilon$ **(measurement error s.d.): baseline**	495.11	262.09	322.73
constant	6.20	5.57	5.78
	(0.05)	(0.09)	(0.07)
$\ln w^*$**: baseline**			1304.30
constant			7.17
			(0.11)
black			-0.64
			(0.10)
coloured			-0.44
			(0.09)
HS grad			0.20
			(0.06)
college			0.40
			(0.10)
high ability			0.09
			(0.06)
previous work			-0.09
			(0.07)
N	1430	1430	1430
$\ln L$	-1,055,884	-1,055,534	-1,052,301
$\Pr(w \geq w^*)$	0.29	0.59	0.44
σ_ϵ (measurement error s.d.)	0.27	0.14	0.17
as percentage of observed accepted wage s.d.			

Robust standard errors in parentheses. Reservation wages at top row refer to inputs of maximum likelihood estimation: w^* is median reservation wage from data; w_{q_5} is 5th percentile reservation wage; and w^*_{MLE} is maximum likelihood estimate (all by cell defined by included covariates). Estimation is by maximum likelihood, with starting values taken from converged estimates of sequential estimation of wage offer and unemployment duration distributions. Optimization algorithm alternates between BFGS and BHHH. "Baseline" refers to value of exponentiated constant term for each parameter, and may be interpreted as parameter level for left-out category (white high school dropouts of low ability, with no previous work experience). $\Pr(w \geq w^*)$ calculated as mean over distribution of full sample, i.e., $\Pr(w \geq w^*) = \int \Pr(w \geq w^*|x) f(x) dx$.

Table 3.10: Empirical and predicted unemployment survivor functions

		$\mathbf{Pr}(d \geq d_0)$		
	(1)	(2)	(3)	(4)
Reservation wage	**empirical**	w^*	w_{q_5}	w^*_{MLE}
UE duration (months)				
3	0.69	0.75	0.75	0.75
6	0.58	0.60	0.60	0.60
12	0.42	0.42	0.43	0.42
24	0.16	0.25	0.25	0.25
36	0.04	0.15	0.16	0.16
χ^2		424.7	399.3	430.7
p-value		0.00	0.00	0.00

Each cell reports value of survivor function at UE duration in left-hand column, i.e., each cell gives the proportion of the unemployment duration distribution that is at least as great as the value in the left-hand column. Column (1) is empirical survivor function observed in the sample, while columns (2)-(4) give predicted survival function for models using the indicator reservation wage inputs. χ^2 statistic is from auxiliary regression of ones on sample moments; statistic is NR^2 from this regression, and is distributed $\chi^2(m)$, where $m = 1$ is the number of moments; see Cameron and Trivedi (2005, p. 261-2). Appendix 3.A.3 describes this test in greater detail.

Table 3.11: Moments and quantiles of empirical and predicted accepted wage distributions

		Accepted wage		
	(1)	(2)	(3)	(4)
Reservation wage	**empirical**	w^*	w_{q_5}	w^*_{MLE}
mean	2486.4	2346.4	2336.0	2295.2
std. dev.	1859.9	1356.6	1682.5	1529.5
quantiles				
0.1	902.0	886.9	709.6	866.6
0.25	1299.9	1341.2	1087.4	1224.6
0.5	1835.2	1969.7	1760.6	1789.8
0.75	3108.0	2899.8	2915.4	2753.2
0.9	4961.0	4278.9	4676.0	4282.1
χ^2		221.7	204.0	196.8
p-value		0.00	0.00	0.00

Each cell reports corresponding moment or quantile of observed accepted wages for empirical wage distribution (column 1) and predicted wage distribution by reservation wage input used in model estimation (columns 2-4). χ^2 statistic is from auxiliary regression of ones on sample moments; statistic is NR^2 from this regression, and is distributed $\chi^2(m)$, where $m = 1$ is the number of moments; see Cameron and Trivedi (2005, p. 261-2). Appendix 3.A.3 describes this test in greater detail.

Table 3.12: Regressions of labor market outcomes on estimated search costs

Panel A: Unemployment duration	(1)	(2)	(3)
$\hat{b}$	0.034	0.022	0.028
	(0.053)	(0.058)	(0.062)
N	1430	1430	1430
R^2	0.24	0.24	0.24
Panel B: Censored duration	(4)	(5)	(6)
$\hat{b}$	0.002	0.001	0.002
	(0.002)	(0.002)	(0.002)
N	1430	1430	1430
R^2	0.19	0.19	0.19
Panel C: Accepted wage	(7)	(8)	(9)
$\hat{b}$	0.049	0.048	0.051
	(0.016)***	(0.017)***	(0.018)***
N	977	977	977
R^2	0.29	0.28	0.28
w^* used	w^*	w_{q5}	w^*_{MLE}

Robust standard errors in parentheses: * significant at 10%; ** significant at 5%; *** significant at 1%. All regressions include covariates used in structural model estimation: dummies for black, coloured, HS grad, at least some college, high ability and previous work experience. $\hat{b}$ calculated by calibrating search cost so that w_i^* matches w^* from structural model, with discount factor $\delta = .95$ annually and separation probability p calibrated from observed separations from first job in sample. $\hat{b}$ measured in thousands for regressions with unemployment duration and censoring indicator as outcomes. All regressions use survey weights. Standard errors calculated by bootstrap (500 replications).

Figure 3.1: Youth employment/population in the US and South Africa

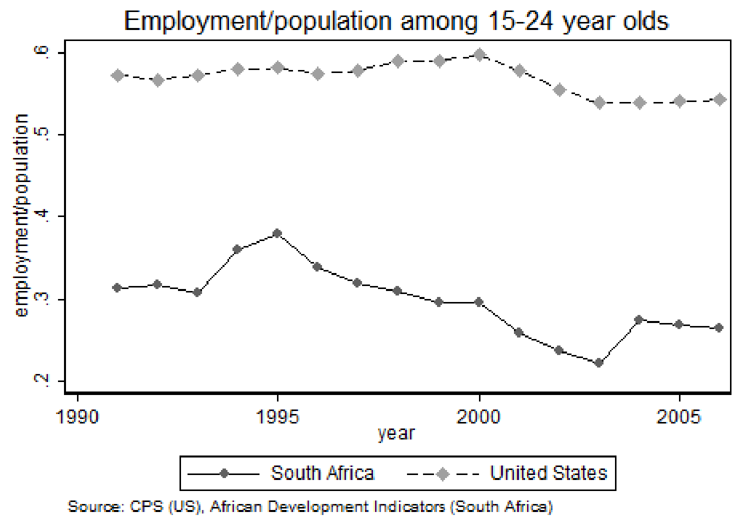

Figure 3.2: Wages and reservation wages

Full-time equivalent wages based on 160 hours of work per month.

Figure 3.3: Density of accepted wages and first unemployment spell

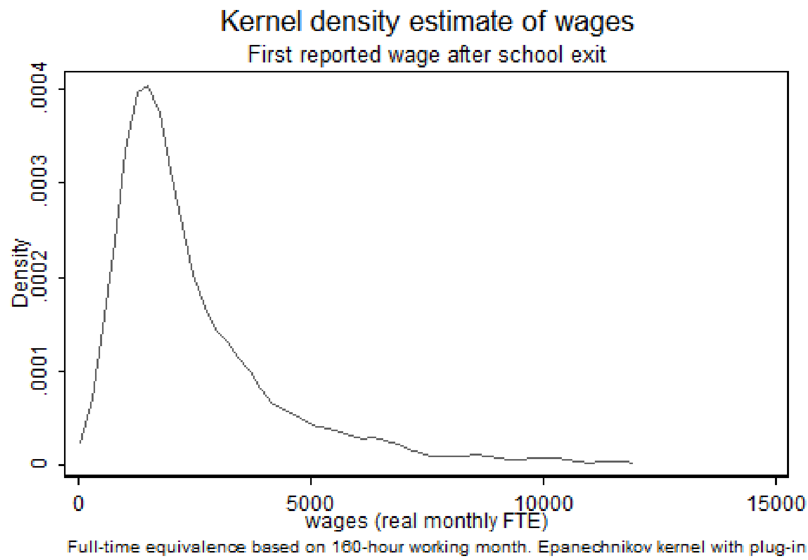

(a)

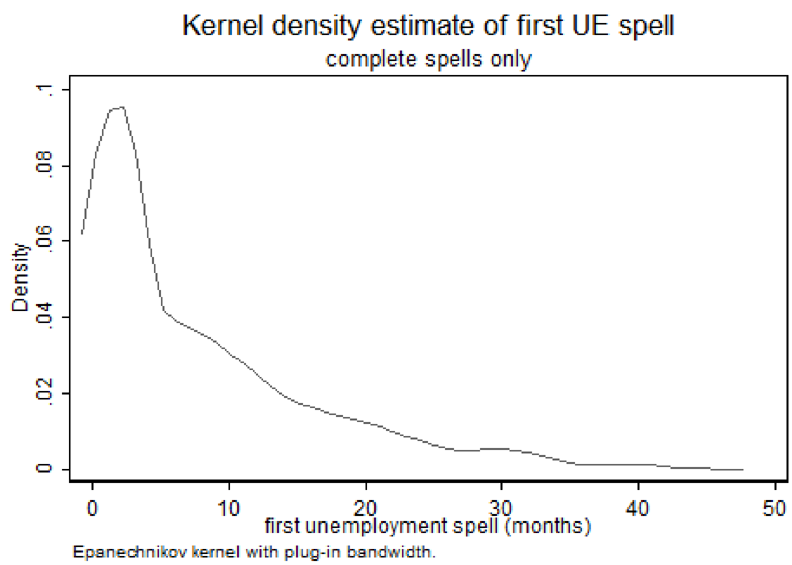

(b)

Figure 3.4: Reservation wages and accepted wages under employer wage subsidy

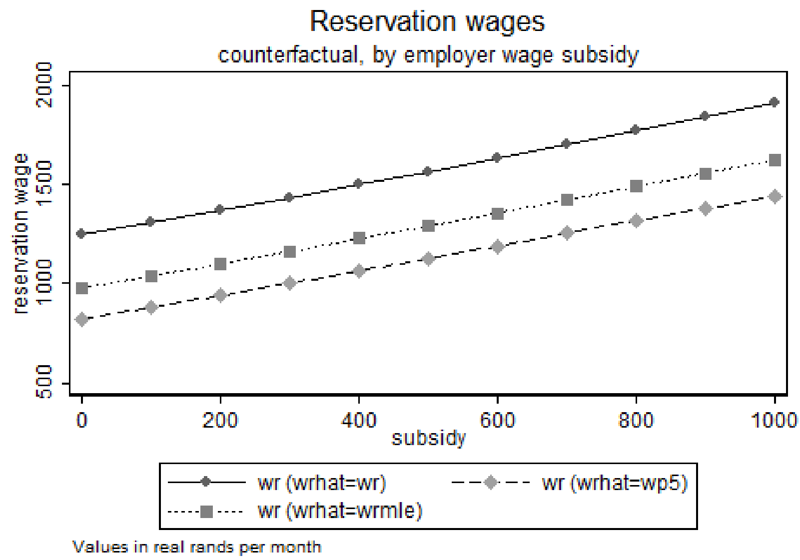

(a)

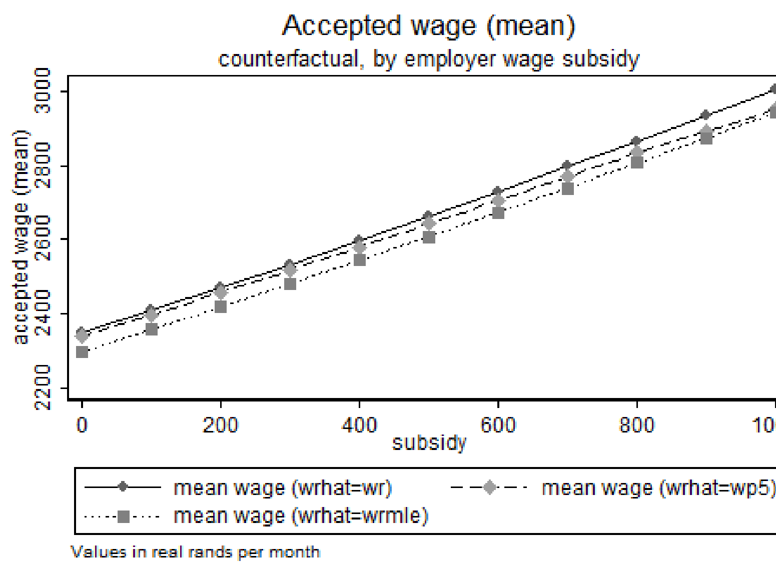

(b)

Figure 3.5: Probability of offer acceptance under employer wage subsidy

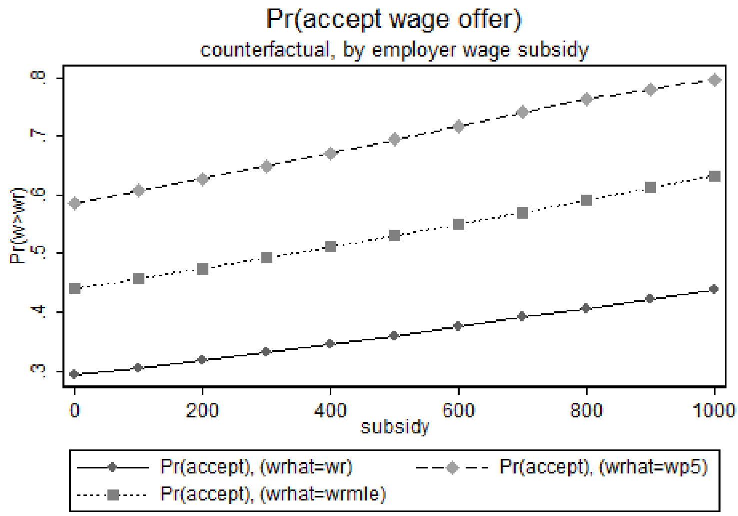

Figure 3.6: Unemployment survivor function under employer wage subsidy: 12 and 24-month UE spell

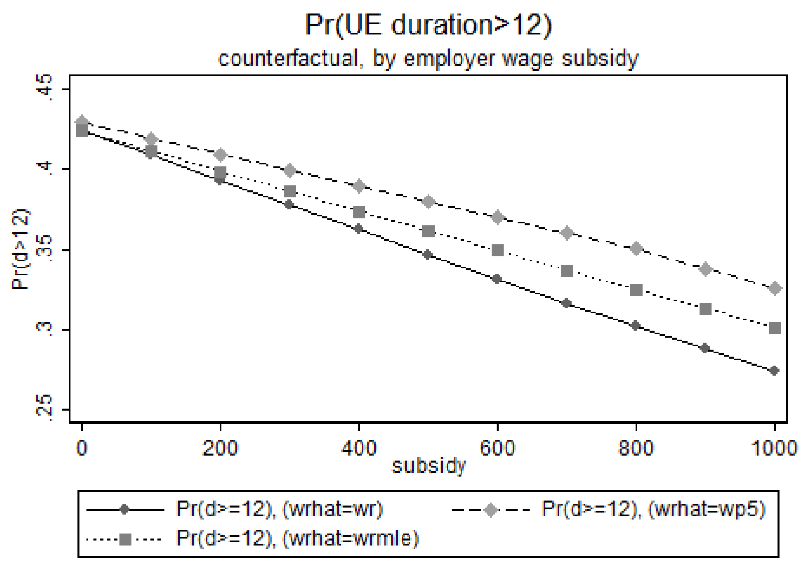

(a)

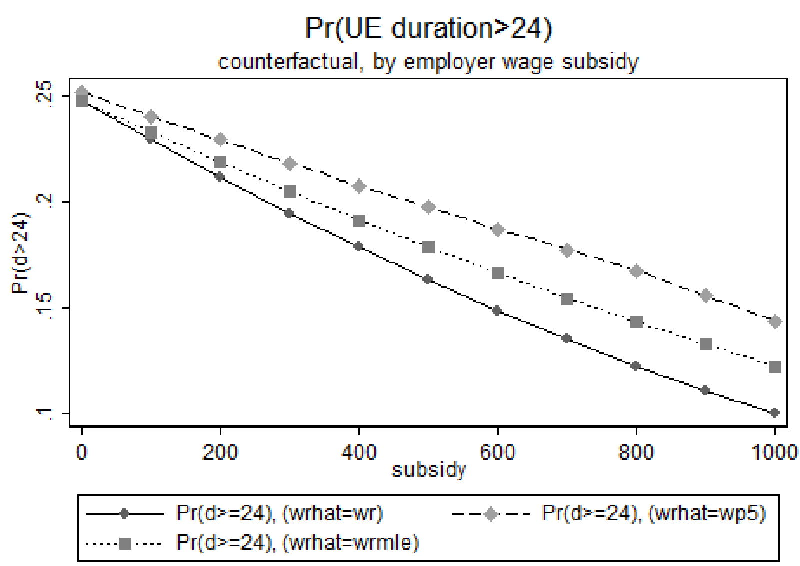

(b)

3.A.1 Derivation of Likelihood Function

This appendix provides more detail on the derivation and form of the likelihood function used in model estimation. The likelihood function is composed of two additively separable parts that follow from the search model: the accepted wage distribution and the unemployment duration distribution. We consider each in turn:

Accepted wage distribution. Under our assumption that wage offers are distributed exponential(λ), the accepted wage distribution is:

$$\begin{aligned} f_W(w|w \geq w^*) &= \frac{f_W(w)}{1 - F_W(w^*)} \\ &= \frac{1}{\lambda} \exp\left(-\frac{w - w^*}{\lambda}\right) \end{aligned}$$

Because we also assume that wages are measured with error such that $w_o = w + \epsilon$, where w_o is the observed accepted wage and ϵ is distributed $N(0, \sigma_\epsilon^2)$, we have the following distribution of observed accepted wages:

$$\begin{aligned} f_W(w_o|w \geq w^*) &= \int_{-\infty}^{\bar{\epsilon}} f_W(w_o|w \geq w^*, \epsilon)\phi\left(\frac{\epsilon}{\sigma_\epsilon}\right) d\epsilon \\ &= \int_{-\infty}^{\bar{\epsilon}} \frac{1}{\lambda} \exp\left(-\frac{w_o - \epsilon - w^*}{\lambda}\right) \phi\left(\frac{\epsilon}{\sigma_\epsilon}\right) d\epsilon \\ &= \exp\left(\frac{-2w_o\lambda + 2w^*\lambda + \sigma_\epsilon^2}{2\lambda^2}\right) \times \frac{1}{\lambda}\phi\left(\frac{w_o - w^*\lambda + \sigma_\epsilon^2}{\lambda\sigma_\epsilon}\right) \end{aligned}$$

where $\phi(\cdot)$ is the standard normal distribution, and $\bar{\epsilon} = w_o - w^*$ is the upper bound on the distribution of ϵ.

Unemployment duration distribution. Under our assumption of Poisson offer arrivals, the hazard of unemployment exit h is the (constant) product of the offer arrival rate q and the probability that the offer will be accepted, i.e., $h = q(1 - F_W(w^*))$. Accordingly, unemployment durations d are distributed exponentially with parameter h, so that $f_D(d) = h\exp(-hd)$. Because some unemployment spells are right-censored, the observed duration $d = \min\{d^*, d_c\}$, where d^* is the true duration and d_c is the duration observed when the spell was censored. Let $c = \mathbb{I}\{d = d_c\}$ be an indicator for censored spells. Then the density of observed unemployment durations, $g_D(d)$, is:

$$\begin{aligned} g_D(d) &= f_D(d)^{1-c}[1 - F_D(d)]^c \\ &= [h\exp(-hd)]^{1-c}[\exp(-hd)]^c \end{aligned}$$

Finally, let $m = \{0, 1\}$ be an indicator for missing wage data (either due to a censored unemployment spell or otherwise). The individual's likelihood contribution is the (log) sum of the observed accepted wage and unemployment duration densities:

$$L(\theta) = (1 - m)\ln f_W(w_o|w \geq w^*; \theta) + \ln g_D(d; \theta)$$

for $\theta = (q, \lambda, \sigma_\epsilon)$.

3.A.2 Data Definitions

The sample is all young adults in CAPS who began as enrolled students at the inception of the monthly calendar data (August 2002) but have exited school; are observed for at least 12 months since leaving school in the monthly calendar data; and have non-missing reservation wage data (reservation wage measure defined below). Additionally, those below the 1st and above the 99th percentiles of accepted wages are dropped. School exit is defined as at least 3 consecutive months of school absence in the calendar data (only 6% report returning to school after a minimum 3-month absence, none of them full-time). Time is calculated relative to month of school exit, so that month 1 is the first of the minimum 3 consecutive months of school absence that define school exit.

Unemployment duration is calculated relative to month of school exit, so that the minimum unemployment duration is one month. An unemployment spell ends when the youth reports working in any job in a calendar month, where work is defined as employment for pay, in-kind benefits or "family gain." Censored observations are those that had not completed their first unemployment spell by the end of the observation period (December 2006).

The observed wage is the first reported wage after school exit across Waves 1-4, adjusted for monthly CPI (base is August 2002, the first month of calendar data) at the time of interview and scaled to full-time monthly equivalent based on 160 working hours per month (those reporting monthly hours above 160 are considered full-time and do not receive an adjustment). Wages reported in Waves 2-4 are the sum of wages reported across all jobs held.

When the reservation wage is based on survey data, it is the value from the most recent interview before conclusion of the first unemployment spell since exiting school. For Wave 1, the reservation wage $w^* = w^*_{moft}$, where w^*_{moft} is the response to the question, "What is the lowest monthly wage you would accept for full-time work?" For Waves 2-4, the reservation wage is defined as $w^* = \min\{w^*_{moft}, w^*_{revealed}\}$, where $w^*_{revealed}$ is the lowest wage associated with an affirmative response to the series of questions, "Would you accept a job doing occupation x at monthly wage w?" Reservation wages are adjusted for monthly CPI (August 2002 base) at the time of interview. For those with a censored first unemployment spell, the reservation wage is the last reported reservation wage in the panel.

Search is defined as a positive response to the "Searched for work in this month?"

question in the calendar data. The job separation probability is calibrated as total number of separations from the first job divided by total months employed in first job since leaving school for all observations in the sample.

Age is age in years at school exit. Schooling is years of completed schooling at school exit. The ability proxy is the z-score from the literacy and numeracy evaluation (LNE) administered by CAPS in Wave 1. The "previously worked" variable is an indicator for whether the youth worked for pay (i.e., reported a non-zero wage) in the panel prior to school exit. Full-time work is defined as an an average of at least 35 hours per month. The survey weight is the young adult sample weight, which is adjusted for the sample design plus household and young adult non-response.

3.A.3 Tests of Model Fit

This appendix discusses the formal test of model fit we use to compare our predicted unemployment duration and accepted wage distributions to the data. For continuous data, Cameron and Trivedi (2005, pp. 261-2) propose a variation of the Lagrange Multiplier (LM) test using the sample moments and scores from the estimated model.[36] Let $\hat{m}_i = m(x_i, \hat{\theta})$ be the sample moment(s) for observation i evaluated at the estimated parameters $\hat{\theta}$. For instance, for exponential wage offers we would have $\hat{m}_i = w_i - (\hat{\lambda} + w^*)$. Let $\hat{s}_i = s(x_i, \hat{\theta}) = \frac{\partial \ln L_i}{\partial \hat{\theta}}$ be the score vector for observation i evaluated at $\hat{\theta}$. Under the null hypothesis that the model is correctly specified, $E(m) = E(s) = 0$. Cameron and Trivedi propose the following auxiliary regressions:

[36] Although many researchers use the Pearson χ^2 test to evaluate the fit of structural models, Cameron and Trivedi (2005, pp. 266) note that the test is invalid if the data are not generated from a multinomial distribution. Since our outcomes of interest (duration and wages) are continuous, we use the LM test described above.

$$1 = \hat{m}_i'\delta + \hat{s}_i'\gamma + u_i$$

$$1 = \hat{m}_i'\delta + u_i$$

where 1 is a vector of ones and the second auxiliary regression is valid in the case where $\frac{\partial m}{\partial \theta} = 0$, as it is in our case. The corresponding test statistic is then:

$$M = NR_u^2$$

where R_u^2 is the uncentered R^2 from the auxiliary regression. Under the null, M is distributed $\chi^2(h)$, where h is the dimension of m (i.e., h is the number of moments).[37]

[37] Another test of model fit that could be applied in our context is the Kolmogorov-Smirnov test, which is a nonparametric test for the equality of two distributions. However, when the parameters of one distribution are estimated using data from the other, the test statistic may not be asymptotically distributed according to the Kolmogorov distribution, invalidating the test.

CHAPTER IV

The Impact of Mexican Immigration on U.S. Labor Markets: Evidence from Migrant Flows Driven by Rainfall Shocks

4.1 Introduction

The effect of immigration on the labor market outcomes of natives is among the most controversial topics in both labor economics and policy circles. Despite wide variation in immigrant presence at the state and local levels, a major barrier to identifying the effects of immigration on the labor market is the endogeneity of the immigration decision: if immigrants choose their locations (at least in part) according to labor market opportunities, or if places with high and low immigrant presence differ systematically in unobserved ways, then estimates of immigration's impact will be biased. This paper seeks to overcome this barrier to identification by using a novel approach that relies on weather shocks in Mexico (the largest source of migrants to the U.S.) and variation in regional migration patterns from Mexico to the United States.

The existing literature on the effects of immigration on the labor market outcomes of natives is sharply divided. Borjas, Freeman and Katz (1996) distinguish between two main methodological approaches: "area analyses," which use variation in immigrant presence in local labor markets (usually the state or metropolitan statistical area) to identify effects on native employment and wages; and "factor proportions analyses," which treat immigrants as an input into an aggregate production function and examine their effects on similarly-skilled natives. The two approaches tend to find strikingly different results. Area analyses (such as Card 1990, Card and Lewis 2007) often find little or no adverse effect of immigration on natives,[1] while factor proportions analyses (such as Borjas 2003, Aydemir and Borjas 2006, Borjas and Katz 2007) often show that native wages decline significantly in the presence of similarly-skilled immigrants.[2] Borjas (2001) and Cadena (2007) argue that endogenous migrant location choice explains the divergent findings of the two approaches: if immigrants choose locations with better-performing local labor markets (e.g., with differentially more positive trends in wages and employment rates), area analyses will underestimate the immigrant impact on natives. Despite numerous studies based on each approach, the debate remains unresolved.

The central challenge for research on this topic is finding a convincing, exogenous source of variation in immigration that can be exploited to identify causal impacts on labor market outcomes. The endogeneity of local immigrant penetration and skill composition is widely acknowledged in the literature, a key concern being that areas with positive trends in wages and employment may also attract more immigrants. This effect builds in a positive correlation in the observed relationship between immigrant flows and wages (as noted by Borjas and Katz 2007, among others), potentially obscuring any negative causal impacts that might exist.

The main identification strategy used in previous work has been to construct instruments for immigration from recent national immigration flows interacted with past population shares of immigrants in particular geographies (U.S. states or metropoli-

[1]Early literature reviews by Friedberg and Hunt (1995) and Borjas (1994) also conclude that wage impacts of immigration are small in studies using the area approach.

[2]A highly related literature seeks to establish the degree of substitutability between immigrants and natives within skill groups. See, for example, Peri (2010), Ottaviano and Peri (2008), Borjas, Grogger and Hanson (2008), and Peri and Sparber (2009).

tan areas), as in Card (2001), Lewis (forthcoming), Card and DiNardo (2000), and Card (2009). This identification strategy aims to isolate "supply-push" factors driving immigration, and assumes that the factors determining the recent national supply of immigrants are exogenous vis--vis trends in labor market conditions in particular localities. Because this strategy exploits network effects that lead new immigrants to settle in enclaves settled previously by immigrants from the same country, we refer in this paper to the corresponding instrument as the "network instrument."

The main concern with this instrument (as pointed out by Card 2001, Card 2009, among others) is that the central assumption of exogeneity of the recent national immigrant supply (in total, or in variously-defined national-level subgroups) would be violated if past labor market trends that led to high concentrations of immigrants in particular U.S. destinations were persistent over time, continuing to affect local labor market conditions in the present day while also continuing to attract migrants to the same destination areas.[3] A related but even simpler violation of exogeneity would be if shocks to labor market conditions in particular U.S. localities were important determinants of the recent national immigrant flow. For example, if differentially higher wage growth in historically important immigrant destinations in the U.S. (e.g., California, Texas) was an important driver of current national immigrant flows, then the network instrument would lead to positively-biased estimates of the wage impact of immigration.

The contribution of this paper is to identify a source of variation in immigration to particular U.S. states that we argue is more plausibly exogenous to U.S. labor market conditions than previously-used instruments (e.g., the network instrument), and then to use this instrument to identify the labor market impacts of immigration.

[3]This point has been made by Card (2001, footnote 23).

In particular, our instrument isolates exogenous variation in Mexican immigration to specific U.S. states driven by weather variation (specifically, rainfall) in the Mexican states from which migrants to given U.S. states have historically originated. Rainfall affects economic conditions in Mexico, which then affect propensities for affected Mexicans to migrate to the U.S. Crucially for the analysis, migrants from different Mexican states have historically been destined for different U.S. states. This fact, combined with within-year and cross-year variation in Mexican rainfall, allows us to construct an instrumental variable for variation in the Mexican share of the labor force across U.S. states (within year) as well as across years (within U.S. states). Using this approach, we estimate the impact of the Mexican share of the labor force using annual U.S. labor market data over three decades, controlling for year fixed effects, U.S. state fixed effects, and U.S.-state-specific linear time trends.

In taking this approach, our paper follows the area analysis approach to research on the impact of immigration. The focus on migration from Mexico is sensible, since Mexico is the largest source country for U.S. immigrants, accounting for 31 percent of the foreign-born population of the U.S. in 2004 (Hanson 2006). In previous work, Munshi (2003) has shown that Mexican rainfall shocks (particularly, droughts) lead to increased migration to the U.S. Other studies that find effects of weather shocks on international migration are Kugler and Yuksel (2008) and Akee (2007), who find that hurricanes drive emigration from Central America and Micronesia, respectively. Hanson and Spilimbergo (1999) and Orrenius and Zavodny (2005) provide additional evidence that negative Mexican economic shocks stimulate emigration to the U.S., finding that Mexican migration to the U.S. increases following declines in the Mexican real wage. Munshi (2003) also shows that migrants leaving Mexico due to drought tend to operate through existing migrant networks at the destination when making their location choice and in seeking employment. Hanson and McIntosh (2007) also find that variation in Mexico-to-U.S. migration driven by relative cohort size is stronger along historically important migration pathways. The identification strategies in Saiz (2003), Saiz (2007), and Kugler and Yuksel (2008) are similar to ours, exploiting shocks in immigrant-source countries (interacted with past U.S. immigrant destinations by origin country) as predictors of current immigration to examine impacts of immigration on economic outcomes in the U.S..

We find that exogenous increases in the Mexican share of the labor force lead to declines in wages and increases in unemployment among non-Mexicans. Our instrumental variables (IV) estimates are substantially larger in magnitude than the corresponding ordinary least squares (OLS) estimates, and also larger than other estimates in the existing literature. The magnitude of our estimates is likely due to the fact that previous estimates suffer from two types of biases. First, OLS estimates in general are likely to be biased in a positive direction because migrants choose to locate in U.S. states with more positive trends in labor market conditions. In addition, previous estimates suffer from attenuation bias due to classical measurement error in the immigrant share of the workforce (as shown by Aydemir and Borjas 2011). Our IV estimates deal with both these biases and thus yield substantially larger coefficient estimates.

The main concern one might raise about our identification strategy is that the Mexican rainfall variables we construct and use as instrumental variables might be correlated with wages or unemployment at the U.S. state-year level via channels other than changes the Mexican labor force share. Formally, this would constitute a violation of the IV exclusion restriction. While it is impossible to conclusively rule out all possible alternative channels for the effects we find, we provide evidence from additional empirical analyses that helps rule out the three most plausible alternative channels: U.S. rainfall (which might be correlated with rainfall in a U.S. state's historical Mexican migration source areas), U.S. exports to Mexico, and U.S. remittances to Mexico.

In the next section, we describe the methodology used to estimate exogenous changes in the Mexican labor force in the U.S. based on Mexican rainfall shocks. Section 4.3 describes the data and provides evidence in favor of our identification strategy. Section 4.4 presents the main results, and Section 4.5 presents robustness checks and additional analyses. Section 4.6 concludes.

4.2 Methodology

The endogeneity of immigration to labor market outcomes is well understood, as migrants will tend to locate in areas with stronger economic conditions. Consistent estimation of the effect of migration on labor market outcomes therefore depends on isolating plausibly exogenous variation in immigrant presence within labor markets. Our identification strategy relies on regional weather variation in Mexico and its differential impacts on migration to destination states in the U.S.

There are several reasons to expect that weather shocks, such as abnormal levels of rainfall, will affect migrant flows between Mexico and the United States, although the direction of such effects is ambiguous. Because traditional agriculture is rain-fed, rainfall levels can affect agricultural output and employment, as well as other sources of income and employment in agricultural communities. One possibility is that low rainfall creates adverse local economic conditions, leading to an increase in migration from Mexico to the U.S.; this is consistent with the findings of Munshi (2003). Another possibility is that potential migrants are credit-constrained, and must amass sufficient savings to cover fixed migration costs. If this is the case, then low rainfall decreases potential migrants' savings, and reduces migration from Mexico to the U.S., the opposite effect of the first scenario.[4]

However, regional variation in weather shocks in the source country, even if uncorrelated with economic conditions at the destination, is insufficient to identify exogenous changes in migrant stocks across destination regions. For identification, it must be the case that migrants differ across Mexican source regions in their propensity to choose particular U.S. destinations. That migrants tend to locate in areas populated by other migrants from their country or region has been well documented (Bartel 1989, Card 2001, Saiz 2003, Saiz 2007, Cortes 2008), including across Mexican communities (Munshi 2003). Because particular source regions in Mexico will tend to favor particular destination regions in the U.S. (due to distance, transportation costs, and migrant networks), spatial variation in weather shocks in Mexico will induce variation in migrant flows across U.S. regions. If such weather shocks cannot be anticipated and are uncorrelated with labor demand shocks in U.S. destination regions, this variation is exogenous. With data on the source regions of Mexican migrants within each U.S. destination, we could then identify exogenous changes in the Mexican migrant population by exploiting the fact that weather shocks in a particular Mexican locality should have differential effects on Mexican migration to historically important U.S. destinations for the given Mexican locality.

In practice, our approach involves weighting weather shocks in Mexico according to the historical propensity of residents of particular Mexican states to migrate to particular U.S. states. The time series variation in weather shocks at the Mexican

[4]The reverse scenario would also hold under each story, i.e., abnormally high rainfall could create favorable economic conditions and reduce emigration from Mexico to the U.S., or high rainfall could generate income and savings that allow credit-constrained potential migrants to emigrate. The observed correlations between rainfall and migration would be identical under each scenario in the cases of both low or high rainfall.

state level, combined with the cross-sectional variation in migrant networks across U.S. states, allow us to apply our method to panel data on labor market outcomes at the U.S. state level. If weather shocks affect migration from Mexico to the U.S., and if historical network ties influence residents of particular Mexican states when choosing among U.S. destination states, then our instrument will help predict changes in the Mexican share of the labor force in U.S. states. Furthermore, if such shocks and networks are uncorrelated with contemporaneous labor market outcomes across U.S. states, then the instrument will be valid. We discuss the construction of the instrument in more detail below.

Let M_{ut} be the flow of migrants from Mexico to U.S. state u at time t. This flow is composed of migrants from various Mexican states indexed by m,[5] leading to the accounting identity:

$$M_{ut} = \sum_{m=1}^{32} M_{umt} \tag{4.1}$$

where M_{umt} is the flow of migrants from Mexican state m to U.S. state u at time t. Since migration flows are endogenous, we instrument for M_{umt} with weather shocks at the Mexican state level. M_{ut} is observed, while data on M_{umt} does not exist.

As it turns out, absence of data on M_{umt} is immaterial. We construct weights ω_{um} to reflect the importance of migrant flows from Mexican state m to U.S. state u and reformulate M_{ut} as:

$$M_{ut} = \sum_{m=1}^{32} \omega_{um} M_{mt} + \epsilon_{ut} \tag{4.2}$$

where the stochastic error term ϵ_{ut} reconciles any discrepancy between the time-invariant weights ω_{um} and time-specific migrant flows from m to u.

[5] Mexico has 31 states plus the Federal District (Mexico City), for a total of 32 geographic units.

A key challenge for this approach is that there is no representative historical dataset with information on the source regions of Mexican migrants in particular U.S. states that could be used to construct ω_{um}. We therefore proxy for ω_{um} as follows. We start with measures of regional migration patterns that are predetermined with respect to t. Although historical measures of migrant flows between U.S.-Mexico state pairs are unavailable, a number of studies (such as Cardoso 1980, Massey, Durand and Malone 2002, Woodruff and Zenteno 2007) have documented the emergence of migration patterns between Mexican and U.S. regions connected by railroads at the beginning of the 20th century, as U.S. employers would travel by rail to Mexico and return with recruited laborers. Foerster (1971) provides data on a sample of Mexican migrants passing through 3 different U.S. border crossings (labeled San Antonio, El Paso, and Los Angeles[6]) in April 1924, indexed by Mexican state of origin.[7] From this sample, we denote the (annualized) flow of migrants from Mexican state m through border crossing region r as F_{mr}. We next assign U.S. states to a corresponding region according to the configuration of major rail lines at this time.[8] We then construct weights ρ_{ru} representing the share of Mexican immigrants in U.S. state u in the region r total in the 1930 U.S. Census. Thus ρ_{ru} may be interpreted as the probability that a migrant from the Foerster (1971) sample who crosses the border into region r settles in state u, and the product $F_{mr} \times \rho_{ru}$ is the (predicted) flow of migrants from Mexican state m to U.S. state u at a time when regional migrant networks were forming. Finally, because our Stage 1 outcome measure of the (male) Mexican labor force in each U.S. state will be scaled by the size of the U.S. state's

[6] Although San Antonio and Los Angeles are not located on the U.S.-Mexico border, we understand Foerster's terminology to refer to border crossings in eastern Texas (such as McAllen and Brownsville) and San Diego, respectively.

[7] Other papers in the literature that have used the Foerster (1971) data to construct measures of historical Mexico-to-U.S. migration flows include Hanson and McIntosh (2007) and Woodruff and Zenteno (2007).

[8] We based this classification on inspection of rail lines depicted in the *Rand McNally World Atlas* (1927).

(male) labor force, we divide this flow by a predetermined measure of the U.S. state labor force to construct ω_{um} as:

$$\omega_{um} = \frac{F_{mr} \times \rho_{ru}}{L_u} \tag{4.3}$$

where L_u is the size of the state u native labor force from the 1970 U.S. Census.

The next step is to interact these weights with Mexican state-level rainfall measures. Let the vector $\boldsymbol{R}_{mt-1}$ be a set of rainfall variables in Mexican state m in time $t-1$, to be defined below.[9] Combining our weights and rainfall variables, and including fixed effects to account for all elements common to U.S. states, time, and that vary linearly within U.S. states over time, we arrive at the Stage 1 regression:

$$P_{ut} = \phi + \beta' \sum_{m=1}^{32} \omega_{um} \boldsymbol{R}_{mt-1} + \gamma_u + \lambda_t + \delta_u t + \nu_{ut} \tag{4.4}$$

where P_{ut} is the Mexican share of the male labor force in state u at time t, γ_u and λ_t are state and year fixed effects, respectively, and δ_u is a state-specific time trend. ϕ is a constant term, and ν_{ut} is a mean-zero error term. The summation term takes the weighted average of each rainfall variable in time t across Mexican states m, where weights ω_{um} are time-invariant and specific to U.S. state u and Mexican state m, as described previously. β' is a vector of coefficients, one for each weighted rainfall variable.

In specifying rainfall variables, we seek to allow for the fact that, in addition to the level of rainfall mattering for migration, extreme deviations from normal may have impacts different from what would be implied when fitting a simple linear relationship. The rainfall variables included in $\boldsymbol{R}_{mt-1}$ are therefore as follows:

[9]Empirically, rainfall variables in year t (contemporaneous with the measure of the Mexican share of the labor force) do not have economically or statistically significant relationships with the Mexican labor force share, so we focus on lagged values of rainfall.

1. The normalized deviation (z-score) r_{mt-1} from long-run mean rainfall.[10] Thus, a value equal to 1 corresponds to rainfall one standard deviation above state m's long-run mean, -1 is one standard deviation below the long-run mean, etc.

2. An indicator for very high rainfall ($HIGH_{mt-1}$), which takes the value of 1 if the rainfall z-score is equal to or greater than 2 (two standard deviations or more above the mean), and 0 otherwise.

3. An indicator for very low rainfall (LOW_{mt-1}), which takes the value of 1 if the rainfall z-score is equal to or less than -2 (two standard deviations or more below the mean), and 0 otherwise.

4. The interaction of (1) and (2) ($r_{mt-1} * HIGH_{mt-1}$). This will allow the linear relationship between the rainfall z-score and the Mexican labor force share to change at very high rainfall realizations.

5. The interaction of (1) and (3) ($r_{mt-1} * LOW_{mt-1}$). This allows the linear relationship between the rainfall z-score and the Mexican labor force share to be different at very low rainfall realizations.

In practice, we will also examine up to 7 lags of the weighted rainfall variables in the regression, for a total of 35 instruments (5 in each year times 7 years).

Our Stage 2 regressions of U.S. labor market outcomes on the Mexican labor force share follow the form:

$$Y_{ut} = \alpha_0 + \alpha_1 \hat{P}_{ut} + \gamma_u + \lambda_t + \delta_u t + \varepsilon_{ut} \tag{4.5}$$

where Y is a U.S. labor market outcome (such as wages or unemployment). This is the prevailing Stage 2 specification in the literature, with the exception of the

[10]We start with state-level annual rainfall measured in milliliters, subtract the long-run mean, and divide by the long-run standard deviation.

state time trend, which controls for any unobserved state-specific factors that vary (linearly) over time, but is rarely included. We will therefore present results both with and without this state time trend. $\hat{P}_{ut}$ is the predicted Mexican share of the labor force from the Stage 1 regression described above. Our coefficient of interest is α_1, which measures the effect of changes in the Mexican male labor force share on male non-Mexican labor market outcomes.

4.3 Data

There are three main types of data used in this paper: 1) data on weather shocks in Mexico; 2) data on U.S.-Mexico regional migration pathways that are predetermined with respect to our outcomes of interest; and 3) data on U.S. labor market composition and outcomes. We will describe each in turn, and then provide evidence for the durability of Mexican migration networks over time.

The rainfall data used in the analyses are derived from global gridded datasets of monthly precipitation produced by the University of Delaware's Center for Climatic Research. This dataset has many advantages, in particular its high resolution (rainfall grid points are at 0.5-degree intervals). Gridded datasets such as the one used here are created using a combination of "station-level" data (i.e., time series data from actual weather stations) and an algorithm for interpolating rainfall in intervening grid points. The interpolation algorithm combines several primary sources of rainfall data to render the final result. Detailed documentation of this procedure can be found on the center's website.[11] Using GIS software, state identification variables were added to all grid points appearing in U.S. or Mexican states; then, all monthly rainfall values were added up over each year to produce annual rainfall estimates for every grid point. Finally, state-level average annual rainfall figures were generated by

[11] http://climate.geog.udel.edu/˜climate/html_pages/Global2_Ts_2009/README.global_p_ts_2009.html

taking the mean of the annual rainfall values of all grid points within a given state. To construct the normalized rainfall variable (z-score), we demean state-level rainfall by the long-run (1949-2008) mean in each state, and then divide by the within-state standard deviation of rainfall over the same time period.

As mentioned in the previous section, data on historical migration pathways comes from Foerster (1971), who reports on April 1924 migrant flows, by Mexican state,[12] to three U.S. border crossings, which we combine with data on the distribution of Mexicans within rail regions from the 1930 U.S. Census. Data from the 1970 U.S. Census are also used in the calculation of the weights (to scale the weights by U.S. state population).

Data on U.S. labor market composition and outcomes comes from the Current Population Survey (CPS) Merged Outgoing Rotation Group (MORG), which includes respondents in their final month of survey participation in each of the two years they appear in the CPS.[13] We chose this sample because of its annual frequency, large size (approximately 320,000 individuals annually), and inclusion of questions on usual weekly earnings and hours. We aggregate the data by state (including the District of Columbia) and year for the period 1979-2008.[14] Following Borjas (2003), we focus on males aged 18-64 in the civilian labor force. "Mexican" refers to self-reported ethnicity, because country of birth is not available in the CPS until 1994; we retain the self-reported definition throughout in order to maintain

[12]In 1924, Mexican states Baja California Sur and Quintana Roo did not exist; Foerster (1925) only reports on migrant flows from Baja California and Yucatan (which at that time included Baja California Sur and Quintana Roo respectively). We impute 1924 migrant flows in these cases by apportioning Foerster's migrant counts from Baja California and Yucatan (into the four present-day states of Baja California, Baja California Sur, Yucatan, and Quintana Roo) in proportion to 1920 Mexican census population figures for these four geographic units.

[13]CPS respondents are interviewed for 4 consecutive months, ignored for 8, then interviewed again for 4 consecutive months. The MORG is a sample of respondents in months 4 (prior to their hiatus from the survey) and 8 (prior to their permanent exit) of survey participation.

[14]An alternative geographical unit is the metropolitan statistical area (MSA), which is often used in the immigration literature. In our case, however, the use of distant-past migration data would present challenges to using the MSA as the unit of analysis, as MSA definitions and borders have changed substantially over time.

consistency.[15] While Borjas (2003) restricts attention to individuals not enrolled in school, we do not do so in our main results because information on whether an individual is enrolled in school is not available in the CPS for the entire sample period.[16]

Our labor market outcomes of interest are mean log hourly wages and unemployment rates, where hourly wages are calculated as weekly earnings (real, 1983 USD) divided by usual weekly hours, excluding the self-employed. We examine wage and unemployment outcomes for all non-Mexicans (rather than native-born individuals), again because place of birth is unavailable prior to 1994. In addition to presenting results for non-Mexicans in the aggregate, we decompose the non-Mexican sample into skill groups based on level of completed schooling: high school dropout, high school graduate, some college (1-3 years) and college (4+ years).

Table 4.1 presents summary statistics for key variables from all these data sources. Of note is the substantial variation in 1924 migration flows between U.S.-Mexico state pairs, which have a mean of 75 but a standard deviation of 504. Data on U.S. states from the CPS replicate the well-known fact that the Mexican labor force in the U.S. is primarily low-skilled.

Figure 4.1 presents graphical evidence that migration networks between source regions in Mexico and destination regions in the U.S. have endured in the long run. The top panel of the figure shows a color-coded map of migrants to the El Paso border crossing region in 1924 by Mexican state. Not surprisingly, migrants through this crossing (which is located at the westernmost point of Texas, on the Texas-New Mexico border) come disproportionately from the center of the country, which

[15]Although using U.S. Census data would allow us to distinguish the Mexican-born from self-identified Mexicans, as well as provide larger sample sizes, the annual frequency of the CPS allows for a lengthier panel, which is better suited to the high frequency of our weather shock data.

[16]We show below that results are robust to restricting the analysis to the subset of years (1984 onwards) when enrollment information is available and to focusing on wages and unemployment for only the non-enrolled population.

closely follows historical rail links to the U.S. The map in the bottom panel shows migrant counts to the same region[17] from Wave 1 (1993-1994) of the Northern Border Migration Survey (EMIF, for its acronym in Spanish), a Mexican government survey of migrants to various crossings spanning the U.S.-Mexico border. Migrants through El Paso tend to come from the same Mexican states 70 years after the 1924 sample. Analogous maps (not shown) for the other two border crossing regions in the 1924 sample, San Antonio and Los Angeles, show a similar pattern.

Figure 4.2 shows how we assign U.S. states to the three border crossing regions in the 1924 data. Definition of the three rail regions is based on our inspection of the major rail lines in the U.S. in 1927. Figure 4.3 shows predicted migrant flows, by Mexican state, to California and Texas, respectively, based on the 1924 migration data and the distribution of Mexicans within each U.S. rail region in the 1930 Census (i.e., it shows $F_{mr} \times \rho_{ru}$ from the previous section). The pattern is sensible, with migrants to California coming primarily from states in western Mexico, while migrants to Texas come from states in central and eastern Mexico.

Of course, these border crossing patterns do not demonstrate that migrants from particular Mexican states choose U.S. destination states in the same patterns over the long run, as migrants crossing through the same border location may nonetheless choose different destinations over time. However, despite the diffusion of Mexican migrants to non-traditional U.S. destinations in recent years (Massey et al. 2002, Card and Lewis 2007), historical migration networks remain significant predictors of the location choices of Mexican migrants. Table 4.2 shows results of a regression of migration (as share of 1930 Mexican state population) by U.S.-Mexico state pair in the EMIF (aggregated across all waves 1993-2005) on our predicted historical

[17] We aggregate migrants sampled in Ciudad Juarez and Nogales, Mexico to form the El Paso region.

migration rates. The correlation is positive and robust to inclusion of U.S. state and Mexican state fixed effects. The durability of these migrant networks, in conjunction with rainfall variation, will provide identification of exogenous changes in the Mexican labor force in U.S. states.[18]

4.4 Results

4.4.1 Effect of rainfall in Mexico

Before examining the impact of Mexican rainfall on Mexican presence in U.S. labor markets, it is useful to first confirm that Mexican rainfall has an impact on economic outcomes in Mexico. In Table 4.3, we report coefficients on the set of rainfall variables as defined above in regressions of log GDP (total, agricultural, and non-agricultural) at the Mexican state level for the period 1993-2006 on rainfall, where we include Mexican state and year fixed effects. For each dependent variable, we report regressions without and with the interactions between the rainfall z-score and the high and low rainfall indicators. The bottom row of the table reports the p-value of the F-test of joint significance of the right-hand-side rainfall variables included in the regression.

Focusing first on log total GDP (columns 1 and 2), the coefficient on the rainfall z-score is positive and significantly different from zero (at the 10% level) in both specifications. The p-value of the test of joint significance of the set of rainfall variables is substantially lower in column 2 than in column 1 (0.069 vs. 0.202), weighing in favor of inclusion of the full set of rainfall variables in the regression. Coefficients in column 2 indicate that extremely low rainfall reduces state-level economic output (the coefficient on the low rainfall indicator is negative and significant at the 10%

[18] A similar point about the long-term persistence of migration propensities from Mexican states is made by Hanson (2006). What we add here in our analysis is the point that this long-term persistence also extends to the U.S. destinations of Mexican migrants.

level), and that variations in rainfall are less positively correlated with GDP when rainfall is already very low (the coefficient on the low rainfall * z-score interaction is negative and significant at the 10% level). The magnitudes of these effects are reasonable and non-negligible. In column 2, for example, the coefficient on the rainfall z-score indicates that 1-standard deviation higher rainfall is associated with 0.5 log-points higher GDP (roughly 0.5%).

Results in the remaining columns of the table indicate that the full set of rainfall variables also is statistically associated with both agricultural and non-agricultural GDP (the p-value of the joint test of the significance of the rainfall variables rejects the null at conventional significance levels in columns 4 and 6). Results for non-agricultural GDP are very similar to results for overall GDP. None of the coefficients in column 4 for agricultural GDP are individually significantly different from zero, even if they are jointly significant, but the pattern of coefficients indicates that high rainfall leads to higher output, with rainfall having a more positive relationship with output when starting from low levels of rainfall.

4.4.2 Stage 1 estimates

We now turn to our Stage 1 regressions of Mexican male labor force share on the instruments, which are the U.S.-Mexico state pair migration weights times the set of Mexican state rainfall variables, summed over all Mexican states. (We will hereafter refer to these as the "rainfall instruments," though they rely on both rainfall shocks and historical migration patterns for identification.) We include rainfall lags 1-7 in our Stage 1 specification, as the effects of rainfall may persist for several years. As described in Section 4.2, the effect of weather shocks on migration is theoretically ambiguous, making our Stage 1 results of independent interest.

Table 4.4 presents Stage 1 results. Odd-numbered column present results with

U.S. state and year fixed effects only, while even-numbered columns also include U.S. state-specific linear time trends. To reveal the general pattern of results, it is first useful to reduce the number of coefficients to inspect and include only the 7 lagged rainfall z-score variables the columns 1 and 2. In column 1, the effect of rainfall is negative across the 2nd through the 7th lags: low weighted rainfall leads to higher Mexican share of the male labor force in U.S. states. The effects of rainfall in lags 2-4 are statistically significantly different from zero. The F-statistic on the instruments is 83.97. Results are quite similar when including state-specific time trends. Column 2 shows the negative and statistically significant coefficient pattern on lags 2-7 of the rainfall z-score. With inclusion of state time trends, the F-statistic remains high, at 80.45. The magnitudes of these coefficients are reasonable. For example, in column 2, the coefficient on the 3rd lag of the weighted rain z-score (-0.23) indicates that a 0.01 increase in this variable (which is between 1 and 2 standard deviations) leads to a reduction in the Mexican share of the labor force of 0.0023 (0.23 percentage points).

Regressions in columns 3 and 4 include the full set of 35 rainfall instruments (5 rainfall variables each in lags 1 through 7). Patterns in the large number of coefficients are somewhat harder to summarize, but the general pattern is still that higher Mexican rainfall leads to lower Mexican share of the male labor force in the U.S.: nearly all coefficients on the weighted rainfall z-scores remain negative in sign in both columns. Coefficients on the weighted interaction terms with extreme rainfall (both high and low extremes) are generally negative, indicating that the negative impact of Mexican rainfall is magnified when rainfall is equal to or greater than two standard deviations away from the historical mean. In columns 3 and 4, the F-statistics of the test of the joint significance of the instruments are high, at 379.37

and 343.17 respectively.

The amount of variation in the Mexican share of the male labor force in U.S. states generated by Mexican rainfall variation implied by these coefficients is reasonable. To see this, for each U.S. state-year observation we multiplied each of the coefficients in column 4 with the value taken by the corresponding rainfall variable for that state-year, and then took the sum of these products for the state-year observation. This sum represents the change in the Mexican share of the male labor force in the given U.S. state-year implied by our estimates and by Mexican rainfall realizations in the 7 years prior to the observation. The mean of this variable is very close to zero (as should be expected) and has a standard deviation of 0.0063, a 5th percentile of -0.0057 (a 0.57 percentage-point reduction in the Mexican share of the male labor force) and a 95th percentile of 0.0049 (a 0.49 percentage-point increase).

We conclude from these Stage 1 results that rainfall in Mexican migration source areas, operating through historical migration pathways, is negatively associated with the Mexican share of the state-level U.S. male labor force. We are unable to distinguish whether this effect operates primarily by driving new emigrants out of drought-afflicted Mexican states, or by discouraging return migration from the U.S. to these drought-affected areas. Moreover, the average negative effect of rainfall shocks may mask important dimensions of heterogeneity, whereby some Mexican states (or subgroups of potential migrants) may be more likely to face credit constraints when financing migration. We leave these important questions about the mechanisms by which rainfall shocks affect migration to future work, however, and proceed to a discussion of Stage 2 results, using columns (3) and (4) of Table 4.4 as the first stage regressions for our instrumental variable (IV) estimates. With F-statistics above 300, these instruments easily exceed the Stock and Yogo (2002) critical values for the weak instrument test based on TSLS size.[19]

To provide further guidance as to interpretation of the variation in the Mexican share of the labor force induced by Mexican rainfall shocks, in Appendix Table 4.A.1 we present separate regressions analogous to columns 1 and 2 of Table 4.4 but where the dependent variables across columns are Mexicans in particular skill groups (as share of the total male labor force in the U.S. state-year).[20] The magnitude of the coefficients in columns 2 through 5 provides insight into the skill level of Mexican immigrants whose presence in the U.S. is being affected by Mexican rainfall. It appears that Mexican high school graduates are the most responsive to Mexican rainfall shocks: for each given lag of the Mexican rainfall variable the regression coefficient for this skill group (in column 3) is nearly always substantially larger than the corresponding coefficients for any of the other skill groups. That said, effects of the Mexican rainfall variable are non-negligible and statistically significant for the other skill groups as well. All told, the Mexican rainfall variables we use as instruments should be interpreted as influencing the Mexican labor force share in the U.S. for all skill groups, with some weight towards the high school graduate group in particular.

4.4.3 Stage 2 (IV) estimates

We now investigate whether the exogenous variation in the Mexican labor force share identified in Stage 1 affects labor market outcomes for natives. We examine (log) hourly wages and unemployment rates as outcomes, for the entire male labor force as well as split into four skill groups according to completed schooling (high school dropouts, high school graduates, some college, and college). We also

[19] See Figure 2 in Stock and Yogo (2002) for K_2=35, critical value at 5% significance, $n = 1$ and size=0.1.

[20] By construction, the coefficients across skill groups will sum to coefficients reported in columns 1 and 2 of Table 4.4, which are repeated in the first column of Appendix Table 4.A.1.

present OLS results for comparison with our IV results, and run all specifications both without and with state time trends.

Table 4.5 shows results for (log) hourly wages. Panel A shows OLS results. In specifications without state-specific time trends, the migrant share coefficient is not statistically significantly different from zero when all skill groups are pooled (column 1) as well as for all skill groups separately. The coefficient is negative in sign in all cases except for college graduates, for which the coefficient is positive but very close to zero. When state time trends are included in the regressions, the coefficients for the most part turn positive though again none are statistically significant, with the exception of the high school dropout regression (column 4), where the coefficient is negative in sign and significant at the 10% level. (This is the only OLS coefficient that is statistically significantly different from zero at conventional levels.)

IV results in Panel B are strikingly different. In all regressions the coefficients are negative, and much larger in magnitude than the corresponding OLS coefficients. When individuals are pooled across skill groups, the coefficient is statistically significantly different from zero at the 5% level when state-specific time trends are not included (column 1); when the state-specific time trends are included in the regression (column 2) the coefficient becomes larger in magnitude and is statistically significantly different from zero at the 1% level. The coefficient in column 2 (-1.080) indicates that a 0.01 (1 percentage point) increase in the Mexican share of the labor force is associated with a 0.0108 (roughly 1%) decline in the log wage for non-Mexicans. A row at the bottom of the table reports the elasticities implied by the coefficient estimate (percent change in non-Mexican wages associated with a percent change in Mexican labor supply).[21] The coefficient of -1.08 in column 2 of Panel B implies an elasticity very close to unity (-1.006).

IV estimates by skill group are also nearly all significant at conventional levels (with the exception of the coefficient in the high school dropout regression without state-specific time trends). Coefficient estimates typically become larger in magnitude with the addition of the state-specific time trend to the regression. In regressions with state-specific time trends, effects of the Mexican share of the labor force on wages are largest for the lowest skill group, high school dropouts (where the coefficient is about -1.5), and smaller for the other skill groups (where coefficients are all around -1).

Table 4.6 shows a similar set of Stage 2 regressions with the unemployment rate as the dependent variable. The general pattern here is consistent with the wage impacts presented in the previous table. The IV results of Panel B reveal a positive effect of the Mexican share of the labor force on unemployment of non-Mexicans. In column 1 (the specification for all skill groups combined without inclusion of state-specific linear time trends), the point estimate implies that a 1 percentage point increase in the Mexican labor force share leads to a 0.215 percentage point increase in non-Mexican unemployment. The effect is larger (0.349 percentage points) when state time trends are included in column 2. Both coefficients are statistically significant (the latter at the 5% level), and substantially larger than their OLS counterparts.

In regressions by skill group, coefficient estimates are statistically significantly different from zero at conventional levels in regressions with or without state-specific

[21] This uses an elasticity formula similar to the one used in Borjas (2003):

$$\frac{\partial \log w_{ut}}{\partial m_{ut}} = \frac{\theta}{1 + m_{ut}^2} \tag{4.6}$$

where $\log w_{ut}$ is the mean of the log wage in state u and year t, m_{ut} is the percentage increase in labor supply attributed to Mexican presence (Mexicans divided by non-Mexicans), and θ is the regression coefficient on the Mexican share of the labor force (from Table 4.5, Panel B). We use $m_{ut} = 0.036$ (the mean in our state-year sample).

linear time trends, with the exception of the college skill group, where coefficients are very close to zero and insignificant in either specification. For either specification, the effect of the Mexican labor force share on unemployment is largest for the lowest skill group, high school dropouts, and declines nearly monotonically the higher the skill group. For high school dropouts (in the specification with state-specific time trends), a 1 percentage point increase in the Mexican labor force share leads to a 0.654 percentage point increase in native unemployment. Again, all IV coefficients are substantially larger in magnitude than their OLS counterparts.

In sum, we find that exogenous increases in the Mexican share of the labor force leads to declines in wages and increases in unemployment among non-Mexicans. The relative size of the IV estimates on both wages and unemployment, compared to OLS estimates, is exactly what we would expect if OLS estimates are biased towards zero because migrants choose to locate in U.S. states with more positive trends in labor market conditions.

4.5 Discussion and additional analyses

In this section we conduct additional analyses and robustness checks. First, we show that an alternative approach to identification used in related literature does not obtain the same results when applied to the same CPS data and sample period. Third, we describe additional tests for potential violations of the IV exclusion restriction.

4.5.1 Comparison to results using "network" instrument

Several studies in the immigration literature (Card 2001, Saiz 2003, Saiz 2007, Cortes 2008, Farre, Luna and Ortega 2009, Luna and Ortega 2009) have also exploited the durability of migrant networks in identifying the effect of migration on labor markets, but have not used source-region shocks explicitly as we do. Instead, they rely on contemporaneous national migrant inflows multiplied by past shares of migrants from each source at a particular destination, thereby constructing a measure of the contemporaneous migrant flow predicted by historical propensities to locate at the destination. We can conduct a similar exercise for Mexicans in the U.S. and compare the results with those presented in the previous section.

Specifically, let θ_u be the share of all Mexican migrants in the labor force living in U.S. state u from the 1970 U.S. Census, and let M_t be the number of Mexicans in the U.S. labor force at time t. Since θ_u may be interpreted as the (predetermined) probability that a Mexican migrant will locate in state u, we may construct an instrument for the Mexican labor force share in state u at time t as:

$$Z_{ut} = \frac{\theta_u M_t}{L_{ut}} \tag{4.7}$$

where L_{ut} is the labor force. Stage 1 is then a regression of the observed Mexican labor force share P_{ut} on Z_{ut}. Since this instrument relies on historical migrant networks, we refer to it as the network instrument.

Appendix Table 4.A.2 shows Stage 1 results for the network instrument. The instrument is strong when only U.S. state and year fixed effects are included (F-statistic 33.44), but weak when state time trends are added (F-statistic 0.12). The fact that the network instrument no longer remains valid once state time trends are included in the first-stage regression is an important difference vis--vis our rainfall instrument, whose strength is very similar in specifications with and without the state time trends.

Even though the network instrument is inferior to the rainfall instrument in this regard, it is still of interest to examine the Stage 2 results generated by the network instrument in specifications with U.S. state and year fixed effects only. Appendix Table 4.A.3 shows Stage 2 results for the network instrument, with native labor market outcomes identical to those presented for the rainfall instrument (we present here only specifications without state-specific linear time trends). The network instrument does generate a negative and statistically significant effect of the Mexican labor force share on wages across all skill groups, with a 1 percentage point increase in the Mexican labor force share leading to a 0.58 percent reduction in wages overall (column 1). The effect is largest at the low end of the skill distribution, and statistically significant for all except the college group. In contrast when the dependent variable is the unemployment rate, the network instrument finds no statistically significant effect overall (column 1), a (surprisingly) negative and statistically significant effect for high school dropouts, and a positive and statistically significant effect for the some college group.

Our findings using the rainfall instrument contrast somewhat with those of the network instrument. The comparison between the two types of instruments is only possible for the less-conservative specification where state-specific time trends are excluded from the regression, because with inclusion of state-specific time trends the network instrument has essentially zero ability to predict the Mexican share of the labor force across U.S. states and years. When the comparison between the two types of instruments is possible, coefficient magnitudes are roughly similar to those found with the rainfall instrument for wages, but diverge when the dependent variable is unemployment (in which case the network instrument yields a zero effect on overall unemployment). Another important difference between the two types of instruments is that only the rainfall instrument is robust to inclusion of state time trends. This gives us more confidence in its power to isolate exogenous changes in the Mexican labor force presence in the U.S. The network instrument's non-robustness to the inclusion of state time trends means that empirical results using that instrument are more open to concern about bias due to long-running trends in labor market outcomes affecting immigration flows.

4.5.2 Tests for violations of the IV exclusion restriction

It is important to consider – and, when possible, formally test – whether use of Mexican rainfall shocks as instruments for Mexican labor force share might violate the IV exclusion restriction. To be specific, the concern is that the weighted Mexican rainfall variables we construct might be correlated with wages or unemployment at the U.S. state-year level via channels other than changes the Mexican labor force share. We consider here three alternative channels: U.S. state rainfall, U.S. state exports to Mexico, and U.S. remittances to Mexico.

U.S. state rainfall

The first potential concern might be that changes in labor market outcomes of non-Mexicans in particular U.S. states might be due to rainfall *in the given U.S. state*, if U.S. state rainfall is correlated with rainfall in that U.S. state's historical migrant origin states in Mexico. We test whether this is a likely channel in regressions of Appendix Table 4.A.4, where we gauge robustness of our original results (repeated in columns 1 and 2 for reference) to inclusion of controls for current and lagged rainfall in the given U.S. state. We specify U.S. rainfall variables that are exactly analogous to the Mexican rainfall variables (rainfall z-score, high and low indicators, and interactions between the z-score and high and low indicators), and include the variables for the current period (t) and all lags up to the 7th ($t-7$) for a total of 40 control variables. For space reasons we do not show the individual coefficients on the U.S. rainfall variables, but we report at the bottom of the table the F-statistic and p-value of the test of their joint significance. For both dependent variables and in both specifications the U.S. rainfall variables are jointly statistically significantly different from zero at the 1% level.

If inclusion of U.S. state rainfall controls led our original coefficients to fall substantially in magnitude, we would be concerned about this channel representing a potential violation of the IV exclusion restriction. As it turns out, our original results are robust to inclusion of controls for current and lagged U.S. state rainfall. The coefficient on the instrumented Mexican labor force changes little in the wage regressions (Panel A). In the specification without state-specific time trends, the coefficient falls slightly from -0.659 to -0.574 when the U.S. rainfall controls are added, while in the specification with state-specific time trends, the corresponding change is from -1.080 to -0.956. In the unemployment regressions (Panel B), inclusion of the U.S. rainfall controls has almost no effect on the coefficient in the specification without state-specific time trends. In the specification with state-specific time trends, the corresponding change in the coefficient is a bit more substantial (falling from 0.349 to 0.252), but it remains three-quarters its original size and maintains its original 5% level of statistical significance. The robustness of the original results to inclusion of controls for current and lagged U.S. state-level rainfall indicates that our results do not simply reflect the impact of U.S. rainfall on labor market outcomes (that might be correlated with Mexican rainfall in U.S. states' historical migrant origin areas).

U.S. state exports to Mexico

A second potential violation of the IV exclusion restriction would be if higher weighted rainfall in Mexico led to higher U.S. state exports to Mexico. This might come about if higher rainfall in a U.S. state's historical migrant origin areas in

Mexico led to higher demand for U.S. goods (e.g., due to positive income effects), if this higher demand for U.S. goods was biased towards goods from the migrants' historical U.S. destination states, and if higher U.S. state exports positively affected U.S. state-level labor market conditions. We investigate this hypothesis in Appendix Table 4.A.5. The empirical test here is similar to that in Appendix Table 4.A.4: we gauge the extent to which inclusion in the regression of controls for (current and lagged) U.S. state-level exports to Mexico attenuates the estimated coefficient on the Mexican labor force share in wage and unemployment regressions. However, data on U.S. state-level exports to Mexico are only available for a subset of years, 1998-2008.[22] This requires us to limit analysis to CPS data from years 1999-2008 (so that we can include one year of lagged U.S. exports in the regression).[23] The first two columns of the appendix table first show that our original findings from the wage regression are very similar in this restricted sample of years, but not in the unemployment regressions. In the unemployment regressions, the coefficient is small and insignificant in the specification without state-specific time trends, and actually acquires an unexpected negative sign in the specification with state-specific linear time trends.

Columns 3 and 4 of the table show that inclusion of controls for current and lagged log exports from the U.S. state to Mexico by sector (agriculture, manufacturing, and other) has very little effect on the estimated coefficient on the Mexican labor force share.[24] There is therefore no indication that U.S. state-level exports to Mexico are an alternative channel for the effects we find in the paper.

[22] U.S. trade data comes from the WISERTrade dataset.

[23] The instruments remain strong in this restricted sample, as evidenced by high F-statistics for the test of joint significance of the instruments in the first stage, reported at the bottom of the respective panels.

[24] The instruments remain strong in this restricted sample, as evidenced by high F-statistics for the test of joint significance of the instruments in the first stage, reported at the bottom of the respective panels.

U.S. remittances to Mexico

A final channel that is worth considering is remittance and associated labor supply responses by migrants in the U.S. Other research (for example, Yang and Choi 2007, Yang 2008) has found that remittances appear to serve as insurance for households in migrants' home countries, responding countercyclically to weather shocks in migrants' home areas. A concern for this paper's analysis might be that when rainfall is low in Mexican migrants' home states, they send more remittances home and simultaneously reduce their consumption in the U.S.. Reduced consumption in the U.S. could lead, via reduced demand for local goods and services, to lower wages and higher unemployment among non-Mexicans. In principle, such an effect could occur entirely on the basis of responses on the part of Mexican migrants already situated in the U.S., and be independent of any new migration flows responding to Mexican rainfall. Unfortunately, we are not aware of any data source for flows of remittances from U.S. states to Mexico, which prevents us from conducting an analysis analogous to Appendix Tables 4.A.4 and 4.A.5.

However, we are able to address this potential channel indirectly by examining data from the Mexican government on the impact of Mexican state-level rainfall on remittance receipts by Mexican states. Annual data on remittances received at the Mexican state level are available for the six-year period 2003 to 2008. Regressions using these data help to rule out this potential alternative channel. The unit of analysis is the Mexican state-year, yielding 192 observations (32 states times 6 years). The analysis involves regressing log annual remittances received on the state's current and lagged annual rainfall z-score, including state and year fixed effects in the regression. A finding that higher rainfall led to lower remittances received would raise concerns that changes in remittances in response to Mexican rainfall might constitute an alternative channel for our results.

As it turns out, rainfall if anything appears to be positively associated with remittances received at the Mexican state level. Appendix Table 4.A.6 presents regressions where various lags of the Mexican state-level rainfall z-score are included separately and in combination. When statistically significant coefficients appear, they tend to be positive and in the more recent lags (years t and $t-1$). They also are small in economic significance: for example, the 0.018 coefficient on current rainfall z-score in column 7 indicates that one-standard-deviation higher rainfall in a Mexican state leads to 1.8 log-points (roughly 1.8%) higher remittances. These results therefore provide little reason for concern that changes in remittances in response to Mexican rainfall constitute an alternative explanatory channel for the effects found in this paper.[25]

Robustness to restriction to non-enrolled labor force

It is also important to confirm that our results are robust to restricting attention to wages and unemployment among non-Mexicans who are not enrolled in school. In related work analyzing the extent of substitutability between natives and immigrants, Borjas et al. (2008) show that the results in Ottaviano and Peri (2008) are sensitive to whether the sample does or does not include individuals enrolled in school.

The CPS only includes information on whether an individual is enrolled in school from 1984 onwards, so our robustness check restricts the sample to the years 1984-2008. In Appendix Table 4.A.7, the columns 1-2 of the table confirm that the main results of the paper are robust to restricting the sample to the 1984-2008 period

[25]Our finding here that remittances are *pro*-cyclical vis-à-vis rainfall in Mexico stands in contrast to the findings in other countries of counter-cyclical remittance responses to home-country shocks in Yang (2007) and Yang (2008), and therefore bears further exploration. One possible explanation for remittance pro-cyclicality is that remittances respond to returns on investment in the home country. Good rainfall conditions increase the returns to productive investment in Mexico and therefore lead to greater inflows of investment funds.

(dependent variables are as in Tables 4 and 5, where non-Mexican wages and unemployment do include individuals enrolled in school). Results for wages are in Panel A, with results for unemployment in Panel B. Coefficient estimates are somewhat smaller in magnitude compared to those for the full sample period, but the original high levels of statistical significance are maintained.

The test of sensitivity to exclusion of enrolled individuals is in columns 3-4 of the table, where the dependent variables now refer only to the non-enrolled population. As it turns out, coefficient estimates are relatively unchanged (and actually slightly larger in magnitude) compared to the corresponding coefficients in columns 1-2, and statistical significance levels are maintained. Inclusion in the wage and unemployment statistics of individuals enrolled in school does not appear to have an important influence on our results.

4.6 Conclusion

In this paper, we estimate the effects of the Mexican presence in the U.S. labor force on non-Mexican labor market outcomes by using an instrument based on regional variation in rainfall shocks in Mexico and historical migration patterns between source regions in Mexico and destination regions in the U.S. We find that state-to-state migration patterns between the two countries have endured in the long run, with the 1924 migration pattern serving as a significant predictor of migration patterns 70 years later. We also find that Mexican rainfall shocks, operating through these regional migration channels, influence the Mexican labor force presence in the U.S., consistent with previous findings such as Munshi (2003). Using these rainfall shocks to instrument for the Mexican labor force share in a panel of U.S. states, we find that a higher Mexican share of the labor force leads to lower wages and higher unemployment for non-Mexicans.

The impacts we estimate are larger than other estimates in the existing literature. For example, Borjas (2003) estimates the elasticity of native wages with respect to immigration of roughly -0.6, compared to our estimate of -1.0, while "area analysis" studies typically find effects much closer to zero. Our approach likely leads to larger effects because we are able to avoid two types of biases in previous studies: positive bias that arises from immigrants being attracted to areas with improving labor market conditions, and attenuation bias due to measurement error in the immigrant share of the workforce. Another important point to note is that our identification strategy isolates variation in Mexican immigrants specifically, while the existing literature typically examines variation in the overall foreign-born population. If Mexican immigration has different effects on U.S. labor market outcomes than does immigration from other source countries, this would constitute an additional reason why our estimates differ from others'.

The analytical approach we use in this paper can be extended in a number of directions, many of which we are in the process of pursuing. We view this paper as the beginning of a research agenda that examines the impact of Mexican immigration on a variety of other U.S. outcomes, such as real estate prices, prices of services that immigrants typically provide (such as low-skilled labor and domestic service), crime rates, or utilization of public services (schooling and health care, for example). In addition, it is worthwhile extending our methodology to see whether immigration from other countries to the U.S. and other destinations can also be predicted on the basis of origin-area weather shocks, and if so to examine immigration impacts in those contexts as well.

4.7 Tables and Figures

Table 4.1: Summary statistics

	n	mean	s.d.	min.	max.
Weather shocks: Mexican states, 1979-2008					
rainfall (mL)	960	899.3	489.6	54.1	4083.1
rainfall z-score (normalized rainfall)	960	-0.05	1.04	-2.30	5.42
high rainfall (z-score$\geq$2) indicator	960	0.04	0.20	0.00	1.00
low rainfall (z-score$\leq$-2) indicator	960	0.01	0.09	0.00	1.00
high rainfall * rainfall z-score interaction	960	0.13	0.66	0.00	5.42
low rainfall * rainfall z-score interaction	960	-0.02	0.20	-2.30	0.00
Migration patterns between Mexican-US state pairs					
migrant flow (1924)	1632	75.0	504.4	0.0	10109.9
migrant flow (1924) as share of US state labor force	1632	0.0001	0.0005	0.0000	0.0117
First Stage Instruments					
weighted rainfall z-score $(t-1)$	1530	-0.000019	0.007133	-0.071163	0.115098
weighted high rainfall indicator $(t-1)$	1530	0.000191	0.001595	0.000000	0.035014
weighted low rainfall indicator $(t-1)$	1530	0.000005	0.000062	0.000000	0.001916
weighted high rainfall interaction $(t-1)$	1530	0.000538	0.004491	0.000000	0.099648
weighted low rainfall interaction $(t-1)$	1530	-0.000011	0.000128	-0.003964	0.000000
Labor market composition and outcomes (male), U.S. states 1979-2008					
Mexican labor force share					
total	1530	0.037	0.063	0.000	0.352
of which:					
high school dropouts	1530	0.017	0.030	0.000	0.154
high school graduates	1530	0.012	0.020	0.000	0.114
some college	1530	0.006	0.011	0.000	0.072
college	1530	0.003	0.005	0.000	0.031
Native male labor market outcomes					
hourly wages					
all	1530	8.25	0.96	5.92	13.58
high school dropouts	1530	5.96	0.85	1.93	10.61
high school graduates	1530	7.38	0.77	5.56	12.93
some college	1530	7.93	0.86	5.87	13.97
college	1530	11.38	1.36	8.17	15.79
log hourly wages					
all	1530	2.10	0.12	1.78	2.61
high school dropouts	1530	1.78	0.14	0.66	2.36
high school graduates	1530	1.99	0.10	1.72	2.56
some college	1530	2.06	0.11	1.77	2.64
college	1530	2.43	0.12	2.10	2.76
unemployment rate					
all	1530	0.055	0.022	0.018	0.212
high school dropouts	1530	0.113	0.045	0.000	0.285
high school graduates	1530	0.064	0.027	0.009	0.227
some college	1530	0.047	0.021	0.004	0.176
college	1530	0.023	0.011	0.000	0.073

Notes: Rainfall data from University of Delaware Center for Climatic Research. Rainfall z-score is rainfall minus long-run (1949-2008) mean divided by long-run standard deviation for the given Mexican state. High (low) rainfall indicator is for z-score $\geq$2 ($\leq$-2). "Migrant flow (1924)" is count of Mexican migrants between pairs of Mexican and US states (51x32=1632 observations), based on author's calculations using Foerster (1971) and U.S. Census data. "Migrant flow (1924) as share of U.S. state labor force" divides previous variable by U.S. state male labor force from 1970 U.S. Census (this is w_{um} in Data section of paper). Remaining variables in table are at U.S. state-year level (51x30=1530 observations). Each instrument is weighted sum of a Mexican rainfall variable across Mexican states, where weights are "Migrant flow (1924) as share of U.S. labor force" specific to respective U.S.-Mexican state pairs. U.S. labor market data from CPS Merged Outgoing Rotation Group, U.S. states and District of Columbia, 1979-2008. Mexican share of male labor force calculated using "Mexican" self-reported ethnicity. Labor force is male civilians aged 18-64 not living in group quarters. Log hourly wage is mean within group (excluding self-employed), where hourly wage calculated as weekly earnings/usual hours per week.

Table 4.2: Recent and distant-past migration patterns between U.S.-Mexico state pairs

	Migration rate, 1993-2006	
	(1)	(2)
Migration rate, 1924	0.08	0.05
	(0.01)***	(0.01)***
Observations	1632	1632
R-squared	0.1	0.34
U.S. state fixed effects		x
Mexican state fixed effects		x

Notes: Robust standard errors in parenthesis: * significant at 10%; ** significant at 5%; *** significant at 1%. Sample is all U.S. state-Mexican state pairs, including District of Columbia (U.S.) and Federal District (Mexico). Dependent variable: Migrants from Mexican state to U.S. state (intended destination), summed over EMIF Waves 1-12 (1993-2006), divided by 1930 Mexican state population. All migrants from EMIF between ages 18-64. Independent variable is 1924 migration rate between state pair, constructed as [(share of 1930 U.S. Mexican population by U.S. state within U.S. rail region) x (annualized flow of Mexican migrants from Mexican state destined for U.S. rail region, 1924)]/(1930 Mexican state population). 1924 migration data from Foerster (1971).

Table 4.3: Effect of Mexican rainfall on overall, agricultural, and non-agricultural Mexican GDP

	(1) Log GDP	(2) Log GDP	(3) Log Ag. GDP	(4) Log Ag. GDP	(5) Log Non-Ag. GDP	(6) Log Non-Ag. GDP
rainfall z-score	0.004*	0.005*	0.005	0.006	0.003	0.004*
	(0.002)	(0.003)	(0.007)	(0.008)	(0.002)	(0.002)
low rainfall indicator	-0.008	-0.461*	-0.033**	0.102	-0.007	-0.571**
	(0.012)	(0.231)	(0.016)	(0.454)	(0.014)	(0.243)
high rainfall indicator	-0.003	0.017	-0.010	-0.001	-0.001	0.015
	(0.008)	(0.020)	(0.027)	(0.067)	(0.009)	(0.019)
low rainfall * rainfall z-score		-0.211*		0.062		-0.263**
		(0.109)		(0.213)		(0.115)
high rainfall * rainfall z-score		-0.007		v-0.003		-0.006
		(0.007)		(0.022)		(0.006)
Observations	448	448	448	448	448	448
R-squared	0.999	0.999	0.993	0.993	0.999	0.999
p-Value on Test of Joint Sig.	0.202	0.069	0.005	0.012	0.319	0.057

Notes: Robust standard errors in parentheses, clustered by Mexican state: * significant at 10%; ** significant at 5%; *** significant at 1%. Data are at Mexican state-year level from 1993-2006. All regressions include state and year fixed effects. Rainfall variables defined as in note of Table 4.1.

Table 4.4: Stage 1 regressions of Mexican share of male labor force on weighted rainfall variables

	Mexican male labor force share			
Instruments	(1)	(2)	(3)	(4)
weighted rainfall z-score $(t-1)$	0.024	0.004	-1.323**	-0.252
	(0.037)	(0.020)	(0.526)	(0.169)
weighted rainfall z-score $(t-2)$	-0.091**	-0.104***	-2.110**	-0.050
	(0.041)	(0.028)	(1.016)	(0.085)
weighted rainfall z-score $(t-3)$	-0.231***	-0.230***	-2.200***	0.065
	(0.056)	(0.060)	(0.623)	(0.212)
weighted rainfall z-score $(t-4)$	-0.147***	-0.156***	-1.650***	-0.211
	(0.035)	(0.035)	(0.450)	(0.233)
weighted rainfall z-score $(t-5)$	-0.022	-0.043*	-1.628***	-0.028
	(0.032)	(0.024)	(0.388)	(0.375)
weighted rainfall z-score $(t-6)$	-0.119	-0.112***	-1.358***	-0.028
	(0.112)	(0.037)	(0.403)	(0.238)
weighted rainfall z-score $(t-7)$	-0.087	-0.074	-0.848***	-0.252
	(0.115)	(0.068)	(0.298)	(0.210)
weighted low rain ind. $(t-1)$			3,408.473*	3,311.118**
			(1,914.350)	(1,304.548)
weighted low rain ind. $(t-2)$			1,361.980	1,503.675
			(4,251.086)	(1,925.204)
weighted low rain ind. $(t-3)$			-5,803.938	-4,786.448
			(9,031.071)	(4,458.301)
weighted low rain ind. $(t-4)$			-2,535.616	159.418
			(5,106.514)	(1,164.581)
weighted low rain ind. $(t-5)$			-20,014.040***	-5,026.705**
			(5,788.451)	(2,092.155)
weighted low rain ind. $(t-6)$			-12,953.760***	143.070
			(2,643.822)	(1,763.922)
weighted low rain ind. $(t-7)$			-5,494.070**	3,315.958
			(2,695.819)	(2,480.073)
weighted high rain ind. $(t-1)$			-148.348***	-57.280***
			(30.485)	(10.483)
weighted high rain ind. $(t-2)$			-60.081*	1.019
			(34.487)	(18.903)
weighted high rain ind. $(t-3)$			-9.436	23.679**
			(35.449)	(9.637)
weighted high rain ind. $(t-4)$			27.770***	5.103
			(8.573)	(6.120)
weighted high rain ind. $(t-5)$			22.594***	9.031***
			(5.369)	(3.127)
weighted high rain ind. $(t-6)$			-31.728	12.047
			(20.101)	(9.081)
weighted high rain ind. $(t-7)$			-84.675***	1.253
			(24.950)	(10.723)
weighted low rain interaction $(t-1)$			1,648.608*	1,595.127**
			(929.506)	(634.127)
weighted low rain interaction $(t-2)$			706.427	723.153
			(2,040.046)	(930.216)
weighted low rain interaction $(t-3)$			-2,745.245	-2,322.940
			(4,345.960)	(2,153.718)
weighted low rain interaction $(t-4)$			-1,170.711	64.383
			(2,447.600)	(563.088)
weighted low rain interaction $(t-5)$			-9,418.882***	-2,360.735**
			(2,744.549)	(999.301)
weighted low rain interaction $(t-6)$			-6,084.428***	81.217
			(1,302.869)	(838.139)
weighted low rain interaction $(t-7)$			-2,585.067**	1,574.344
			(1,265.871)	(1,188.366)
weighted high rain interaction $(t-1)$			58.973***	22.238***
			(12.115)	(4.073)
weighted high rain interaction $(t-2)$			26.711**	-0.532
			(11.699)	(7.093)
weighted high rain interaction $(t-3)$			6.589	-9.619**
			(12.921)	(3.993)
weighted high rain interaction $(t-4)$			-8.030**	-1.652
			(3.238)	(2.291)
weighted high rain interaction $(t-5)$			-5.691***	-2.808***
			(2.001)	(0.928)
weighted high rain interaction $(t-6)$			13.121	-4.849
			(8.563)	(3.479)
weighted high rain interaction $(t-7)$			30.925***	-0.984
			(8.351)	(3.582)
Observations	1,530	1,530	1,530	1,530
R-squared	0.937	0.986	0.943	0.988
F-stat, all instruments	83.972	80.450	379.367	343.170
State time trends	No	Yes	No	Yes

Notes: Robust standard errors in parentheses, clustered by U.S. state: * significant at 10%; ** significant at 5%; *** significant at 1%. All regressions include state and year fixed effects. State-specific linear time trends included where indicated. Sample is panel of U.S. states from CPS MORG, 1979-2008. See Table 4.1 for variable definitions.

Table 4.5: Stage 2 regressions of non-Mexican log hourly wages on Mexican labor force share

Panel A: OLS										
Skill group	all		high school dropout		high school graduate		some college		college	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Mexican labor force share	-0.105	0.055	-0.359	-0.495*	-0.174	0.188	-0.157	0.208	0.028	0.080
	(0.150)	(0.226)	(0.221)	(0.261)	(0.154)	(0.249)	(0.127)	(0.247)	(0.140)	(0.233)
Observations	1,530	1,530	1,530	1,530	1,530	1,530	1,530	1,530	1,530	1,530
R-squared	0.875	0.931	0.692	0.767	0.824	0.882	0.782	0.826	0.858	0.896
State time trends	no	yes	no	yes	no	yes	no	yes	no	yes
Panel B: IV										
Skill group	all		high school dropout		high school graduate		some college		college	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Mexican labor force share	-0.659**	-1.080***	-0.513	-1.492*	-0.784***	1.025***	-0.675**	-1.103***	-0.656***	-1.005***
	(0.318)	(0.262)	(0.438)	(0.825)	(0.216)	(0.332)	(0.299)	(0.336)	(0.214)	(0.274)
Observations	1,530	1,530	1,530	1,530	1,530	1,530	1,530	1,530	1,530	1,530
R-squared	0.869	0.925	0.692	0.763	0.815	0.873	0.776	0.817	0.850	0.891
First stage F-stat	379.367	343.170	379.367	343.170	379.367	343.170	379.367	343.170	379.367	343.170
Implied Elasticity	-0.614	-1.006	-0.478	-1.390	-0.730	-0.955	-0.629	-1.028	-0.611	-0.937
State time trends	no	yes	no	yes	no	yes	no	yes	no	yes

Notes: Robust standard errors in parentheses, clustered at state level: * significant at 10%; ** significant at 5%; *** significant at 1%. Sample is U.S. state-years (including D.C.), CPS Merged Outgoing Rotation Group 1979-2008. All regressions include state and year fixed effects. State-specific linear time trends where indicated. See Table 4.1 for variable definitions. Dependent variables are calculated for total male labor force or by skill group as indicated. Mexican labor force share variable is for total male labor force (not by skill group) in all regressions. First stages for IV regressions are columns 3 and 4 of Table 4.4 for specifications without and with state-specific time trends, respectively.

Table 4.6: Stage 2 regressions of non-Mexican unemployment rate on Mexican labor force share

Panel A: OLS										
Skill group	all		high school dropout		high school graduate		some college		college	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Mexican labor force share	0.064	0.143***	-0.003	0.210	0.054	0.246***	0.045	0.097**	0.025	0.017
	(0.044)	(0.045)	(0.072)	(0.179)	(0.054)	(0.058)	(0.039)	(0.045)	(0.018)	(0.028)
Observations	1,530	1,530	1,530	1,530	1,530	1,530	1,530	1,530	1,530	1,530
R-squared	0.689	0.744	0.469	0.505	0.610	0.666	0.523	0.566	0.362	0.399
State time trends	no	yes	no	yes	no	yes	no	yes	no	yes
Panel B: IV										
Skill group	all		high school dropout		high school graduate		some college		college	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Mexican labor force share	0.215***	0.349**	0.237***	0.654***	0.256***	0.514***	0.192***	0.251*	0.036	0.041
	(0.051)	(0.146)	(0.086)	(0.242)	(0.059)	(0.180)	(0.037)	(0.135)	(0.067)	(0.064)
Observations	1,530	1,530	1,530	1,530	1,530	1,530	1,530	1,530	1,530	1,530
R-squared	0.677	0.739	0.462	0.499	0.596	0.660	0.510	0.562	0.362	0.399
First stage F-stat	379.367	343.170	379.367	343.170	379.367	343.170	379.367	343.170	379.367	343.170
Implied Elasticity	0.201	0.326	0.220	0.609	0.238	0.479	0.178	0.234	0.034	0.038
State time trends	no	yes	no	yes	no	yes	no	yes	no	yes

Notes: Robust standard errors in parentheses, clustered at state level: * significant at 10%; ** significant at 5%; *** significant at 1%. Sample is U.S. state-years (including D.C.), CPS Merged Outgoing Rotation Group 1979-2008. All regressions include state and year fixed effects. State-specific linear time trends where indicated. See Table 4.1 for variable definitions. Dependent variables are calculated for total male labor force or by skill group as indicated. Mexican labor force share variable is for total male labor force (not by skill group) in all regressions. First stages for IV regressions are columns 3 and 4 of Table 4.4 for specifications without and with state-specific time trends, respectively.

Table 4.A.1: Stage 1 regressions of Mexican share of male labor force on weighted rainfall, by skill group

Panel A: Without state time trends

	Mexican male labor force share				
skill group	all	HSD	HSG	SC	C
	(1)	(2)	(3)	(4)	(5)
weighted rainfall z-score $(t-1)$	0.02	0.02	0.01	-0.01	0.00
	(0.037)	(0.029)	(0.010)	(0.008)	(0.005)
weighted rainfall z-score $(t-2)$	-0.09**	-0.02	-0.06***	-0.01	-0.00
	(0.041)	(0.024)	(0.010)	(0.010)	(0.007)
weighted rainfall z-score $(t-3)$	-0.23***	-0.01	-0.08***	-0.11***	-0.02**
	(0.056)	(0.031)	(0.016)	(0.023)	(0.009)
weighted rainfall z-score $(t-4)$	-0.15***	-0.02**	-0.07***	-0.05***	-0.01***
	(0.035)	(0.009)	(0.020)	(0.013)	(0.003)
weighted rainfall z-score $(t-5)$	-0.02	0.02	-0.01	-0.02*	-0.01
	(0.032)	(0.022)	(0.010)	(0.010)	(0.006)
weighted rainfall z-score $(t-6)$	-0.12	-0.02	-0.02	-0.08***	-0.00
	(0.112)	(0.047)	(0.042)	(0.021)	(0.012)
weighted rainfall z-score $(t-7)$	-0.09	0.02	0.04	-0.13***	-0.02
	(0.115)	(0.046)	(0.043)	(0.030)	(0.015)
Observations	1,530	1,530	1,530	1,530	1,530
R-squared	0.94	0.93	0.92	0.87	0.85
F-stat on instruments	83.97	7.63	75.79	9.86	28.57

Panel B: With state time trends

	Mexican male labor force share				
skill group	all	HSD	HSG	SC	C
	(1)	(2)	(3)	(4)	(5)
weighted rainfall z-score $(t-1)$	0.00	0.01	-0.02*	0.02***	-0.01*
	(0.020)	(0.016)	(0.012)	(0.007)	(0.004)
weighted rainfall z-score $(t-2)$	-0.10***	-0.02	-0.10***	0.03***	-0.02***
	(0.028)	(0.020)	(0.022)	(0.010)	(0.004)
weighted rainfall z-score $(t-3)$	-0.23***	-0.02	-0.12***	-0.05**	-0.04***
	(0.060)	(0.019)	(0.029)	(0.021)	(0.004)
weighted rainfall z-score $(t-4)$	-0.16***	-0.02**	-0.07***	-0.05***	-0.02***
	(0.035)	(0.009)	(0.021)	(0.011)	(0.003)
weighted rainfall z-score $(t-5)$	-0.04*	0.01	-0.04***	0.00	-0.02**
	(0.024)	(0.014)	(0.010)	(0.007)	(0.006)
weighted rainfall z-score $(t-6)$	-0.11***	-0.02	-0.08***	0.01*	-0.02***
	(0.037)	(0.025)	(0.016)	(0.007)	(0.006)
weighted rainfall z-score $(t-7)$	-0.07	0.01	-0.03	-0.02	-0.03***
	(0.068)	(0.032)	(0.026)	(0.016)	(0.003)
Observations	1,530	1,530	1,530	1,530	1,530
R-squared	0.99	0.97	0.97	0.96	0.93
F-stat on instruments	80.45	6.08	62.14	8.37	71.81

Notes: Robust standard errors in parentheses, clustered by U.S. state: * significant at 10%; ** significant at 5%; *** significant at 1%. All regressions include state and year fixed effects. State-specific linear time trends included in Panel B. Sample is panel of U.S. states from CPS MORG, 1979-2008. See Table 4.1 for variable definitions. Dependent variables are Mexican male labor force share in total or by skill group as indicated. Skill groups are all, high school dropouts (HSD), high school graduates (HSG), some college (SC) and college or more (C).

Table 4.A.2: Network instrument, Stage 1 regression

	Mexican labor force share (1)	(2)
Predicted Mexican labor force share	0.47***	-0.06
	(0.082)	(0.171)
Observations	1,530	1,530
R-squared	0.95	0.98
F-stat on network instrument	33.44	0.12
State time trends	no	yes

Notes: Robust standard errors in parentheses, clustered at state level: * significant at 10%; ** significant at 5%; *** significant at 1%. Sample is U.S. state-years (including D.C.), CPS Merged Outgoing Rotation Group 1979-2008. All regressions include state and year fixed effects. Stage 1 dependent variable is Mexican share of male labor force, where "Mexican" is self-reported ethnicity. Instrument is network instrument (predicted Mexican male labor force share). Predicted Mexican male labor force share calculated as [(share of U.S. Mexican labor force in U.S. state, 1970 U.S. Census) x (Mexican labor force in U.S. at time t)]/(labor force in U.S. state at time t). I.e., predicted share is predicted Mexican labor force in U.S. state based on 1970 distribution, scaled by current U.S. state labor force.

Table 4.A.3: Non-Mexican labor market outcomes and Mexican labor force share, using network instrument (IV estimates)

Panel A: log hourly wages					
skill group	all	HSD	HSG	SC	C
	(1)	(2)	(3)	(4)	(5)
Mexican labor force share	-0.58**	-1.06**	-0.76***	-0.69***	-0.34
	(0.280)	(0.429)	(0.216)	(0.188)	(0.243)
Observations	1,530	1,530	1,530	1,530	1,530
R-squared	0.871	0.686	0.815	0.775	0.856
F-stat on network instrument	33.4	33.4	33.4	33.4	33.4
Panel B: unemployment rate					
skill group	all	HSD	HSG	SC	C
	(1)	(2)	(3)	(4)	(5)
Mexican labor force share	0.04	-0.23*	-0.00	0.08*	0.03
	(0.058)	(0.137)	(0.089)	(0.040)	(0.022)
Observations	1,530	1,530	1,530	1,530	1,530
R-squared	0.689	0.463	0.609	0.522	0.362
F-stat on network instrument	33.4	33.4	33.4	33.4	33.4

Notes: Robust standard errors in parentheses, clustered at state level: * significant at 10%; ** significant at 5%; *** significant at 1%. Sample is U.S. state-years (including D.C.), CPS Merged Outgoing Rotation Group 1979-2008. All regressions include state and year fixed effects. See Table 4.1 for variable definitions, and Table 4.A.2 for definition of network instrument. Skill groups are all skill groups, high school dropouts (HSD), high school graduates (HSG), some college (SC) and college (C).

Table 4.A.4: Second stage regressions with U.S. rainfall lags as controls

Panel A: Log Wages				
Skill group	all			
	(1)	(2)	(3)	(4)
Mexican labor force share	-0.659**	-1.080***	-0.574*	-0.956***
	(0.318)	(0.262)	(0.297)	(0.221)
Observations	1,530	1,530	1,530	1,530
R-squared	0.869	0.925	0.877	0.930
First stage F-stat	379.367	343.170	459.266	384.388
F-Stat on Joing Sig. of U.S. Rain			17.987	56.360
p-Value on U.S. Rain			0.000	0.000
State time trends	no	yes	no	yes
Controls for U.S. rainfall	no	no	yes	yes
Panel B: Unemployment				
Skill Group	all			
	(1)	(2)	(3)	(4)
Mexican labor force share	0.215***	0.349**	0.201***	0.252**
	(0.051)	(0.146)	(0.051)	(0.094)
Observations	1,530	1,530	1,530	1,530
R-squared	0.677	0.739	0.704	0.759
First stage F-stat	379.367	343.170	459.266	384.388
F-Stat on Joing Sig. of U.S. Rain			6.638	6.718
p-Value on U.S. Rain			0.000	0.000
State time trends	no	yes	no	yes
Controls for U.S. rainfall	no	no	yes	yes

Notes: Robust standard errors in parentheses, clustered at state level: * significant at 10%; ** significant at 5%; *** significant at 1%. Sample is U.S. state-years (including D.C.), CPS Merged Outgoing Rotation Group 1979-2008. All regressions include state and year fixed effects. State-specific linear time trends where indicated. See Table 4.1 for variable definitions. Dependent variables are wages and unemployment for total male labor force. Regressions in columns 1 and 2 are as in Table 4.5 and 4.6; see those tables for notes. Regressions in columns 3 and 4 include controls for U.S. rainfall in years $t-1$ to $t-7$. U.S. rainfall controls are defined analogously to Mexican rainfall variables described in Table 4.1: rainfall z-score, indicators for high and low rainfall, and interactions of z-score with high and low indicators.

Table 4.A.5: Second stage regressions with U.S. (state-level) exports to Mexico as controls, 1999-2008

Panel A: IV (Log of Wages)				
	(1)	(2)	(3)	(4)
Mexican labor force share	-0.927***	-0.877***	-0.929***	-0.871***
	(0.244)	(0.166)	(0.233)	(0.168)
ln(agricultural exports)			0.039	0.034
			(0.048)	(0.057)
ln(agricultural exports)$(t-1)$			-0.009	-0.075
			(0.083)	(0.070)
ln(manufacturing exports)			-0.385	-0.500
			(0.511)	(0.445)
ln(manufacturing exports)$(t-1)$			-0.396	-0.458
			(0.427)	(0.433)
ln(other exports)			0.103	0.075
			(0.134)	(0.102)
ln(other exports)$(t-1)$			0.085	0.058
			(0.169)	(0.165)
Observations	561	561	561	561
R-squared	0.928	0.951	0.653	0.770
First stage F-stat	1391.852	91.905	358.524	117.055
F-Stat for Joing Signif. of Exports			0.418	0.848
p-Value of F-stat.			0.863	0.539
State time trends	no	yes	no	yes
Panel B: IV (Unemployment)				
	(1)	(2)	(3)	(4)
Mexican labor force share	0.089	-0.125**	0.083	-0.143**
	(0.058)	(0.060)	(0.055)	(0.059)
ln(agricultural exports)			-0.015	-0.055*
			(0.028)	(0.030)
ln(agricultural exports)$(t-1)$			0.000	-0.008
			(0.026)	(0.024)
ln(manufacturing exports)			0.111	0.187
			(0.158)	(0.152)
ln(manufacturing exports)$(t-1)$			0.031	0.042
			(0.132)	(0.130)
ln(other exports)			0.013	0.089*
			(0.051)	(0.051)
ln(other exports)$(t-1)$			-0.009	0.023
			(0.037)	(0.040)
Observations	561	561	561	561
R-squared	0.929	0.952	0.655	0.777
First stage F-stat	1391.852	91.905	358.524	117.055
F-Stat for Joing Signif. of Exports			0.218	1.900
p-Value of F-stat.			0.969	0.099
State time trends	no	yes	no	yes

Notes: Robust standard errors in parentheses, clustered at state level: * significant at 10%; ** significant at 5%; *** significant at 1%. Sample is U.S. state-years (including D.C.), CPS Merged Outgoing Rotation Group 1999-2008. All regressions include state and year fixed effects. State-specific linear time trends where indicated. See Table 4.1 for variable definitions. Dependent variables are wages and unemployment for total male labor force. Regressions in columns 1 and 2 are as in Table 4.5 and 4.6, but restricted to years 1999-2008; see those tables for notes. U.S. trade data comes from WISERTrade.

Table 4.A.6: Effect of rainfall on remittances in Mexican states

	Log of Annual Remittances							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
rain z-score (t)	0.013					0.017*	0.018*	0.017*
	(0.009)					(0.009)	(0.009)	(0.010)
rain z-score ($t-1$)		0.013				0.018*	0.019*	0.018
		(0.009)				(0.010)	(0.011)	(0.011)
rain z-score ($t-2$)			0.004			0.009	0.011	0.010
			(0.008)			(0.009)	(0.011)	(0.011)
rain z-score ($t-3$)				-0.004			0.005	0.004
				(0.006)			(0.008)	(0.009)
rain z-score ($t-4$)					-0.012			-0.002
					(0.011)			(0.009)
Observations	192	192	192	192	192	192	192	192
R-squared	0.995	0.995	0.995	0.995	0.995	0.995	0.995	0.995
p-Value (joint signif. of rainfall vars.)	0.129	0.149	0.648	0.539	0.292	0.301	0.393	0.507

Notes: Robust standard errors in parentheses, clustered at state level: * significant at 10%; ** significant at 5%; *** significant at 1%. Remittance data from Banco de Mexico, at Mexican state-year level from 2003-2008. Rainfall data come from the University of Delaware Center for Climatic Research.

Table 4.A.7: Robustness of second stage results to exclusion of enrolled individuals, 1984-2008

Panel A: IV (Log of Wages)	(1)	(2)	(3)	(4)
Mexican labor force share	-0.439*	-0.804***	-0.522**	-0.902***
	(0.225)	(0.258)	(0.237)	(0.253)
Observations	1,275	1,275	1,275	1,275
R-squared	0.905	0.940	0.907	0.943
First stage F-stat	1345.448	2684.01	1345.448	2684.01
State time trends	no	yes	no	yes
Enrolled individuals included in dependent variable?	yes	yes	no	no

Panel B: IV (Unemployment)	(1)	(2)	(3)	(4)
Mexican labor force share	0.122**	0.185**	0.135***	0.214**
	(0.048)	(0.076)	(0.049)	(0.083)
Observations	1,275	1,275	1,275	1,275
R-squared	0.629	0.723	0.636	0.728
First stage F-stat	1345.448	2684.01	1345.448	2684.01
State time trends	no	yes	no	yes
Enrolled individuals included in dependent variable?	yes	yes	no	no

Notes: Robust standard errors in parentheses, clustered at state level: * significant at 10%; ** significant at 5%; *** significant at 1%. Sample is U.S. state-years (including D.C.), CPS Merged Outgoing Rotation Group 1984-2008. All regressions include state and year fixed effects. State-specific linear time trends where indicated. Dependent variables in columns 1 and 2 are wages or unemployment for total male labor force, as defined in Table 4.1 notes (includes individuals enrolled in school). Dependent variables in columns 3 and 4 are wages or unemployment for total male labor force, excluding individuals enrolled in school.

Figure 4.1: Migrant flows through El Paso region, 1924 and 1993-1994

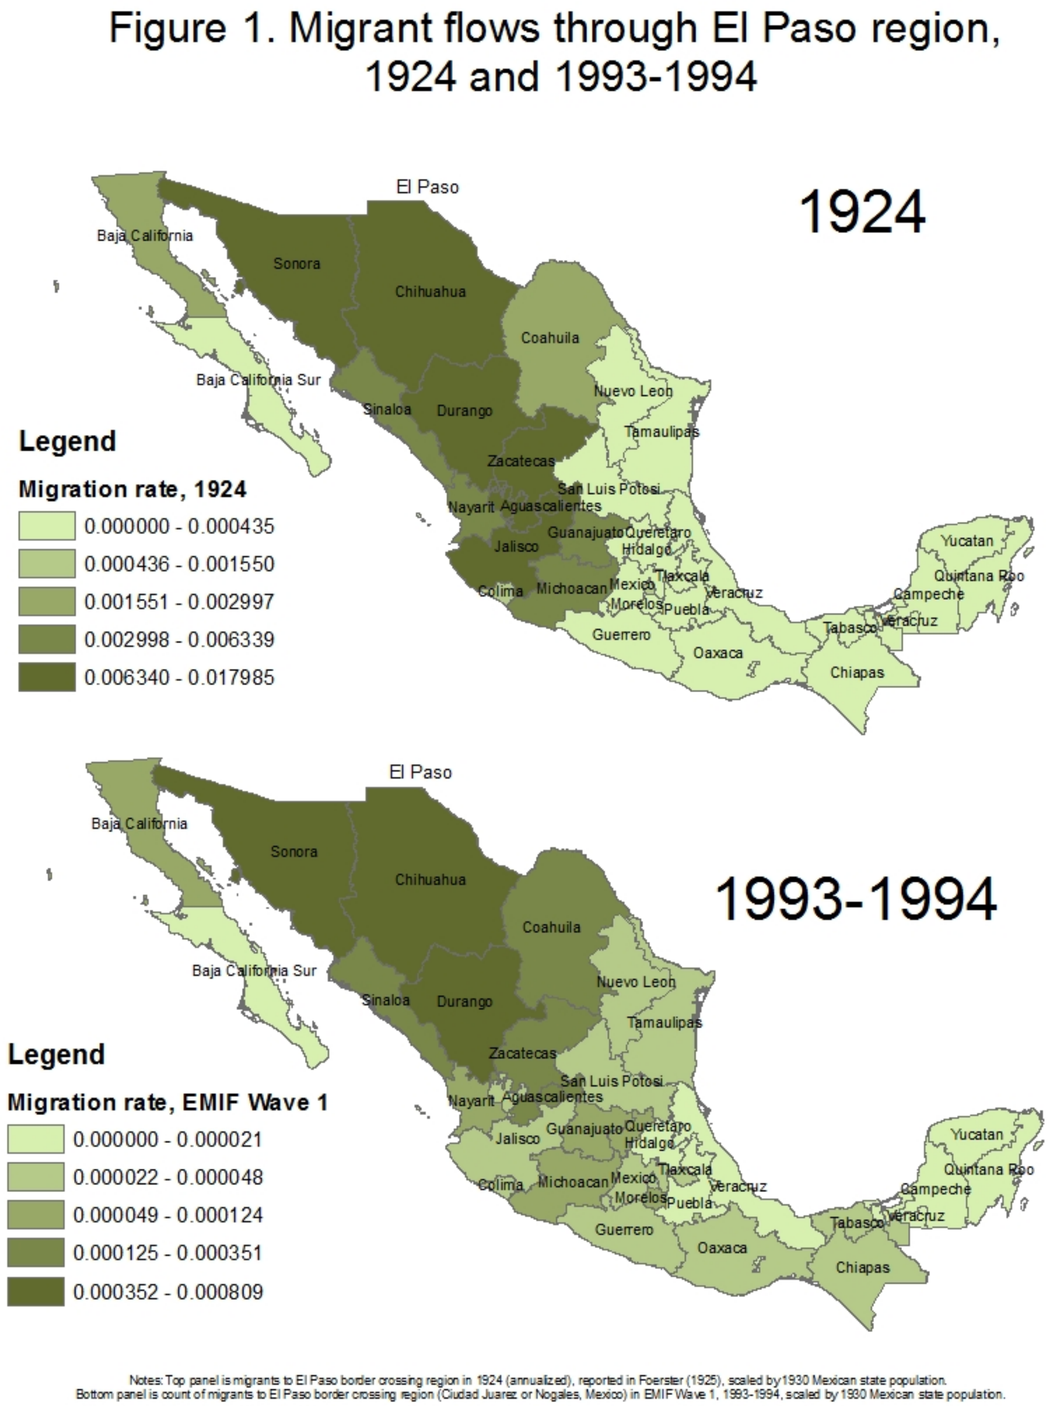

Figure 4.2: United States Railroad Regions

Figure 2. United States Railroad Regions

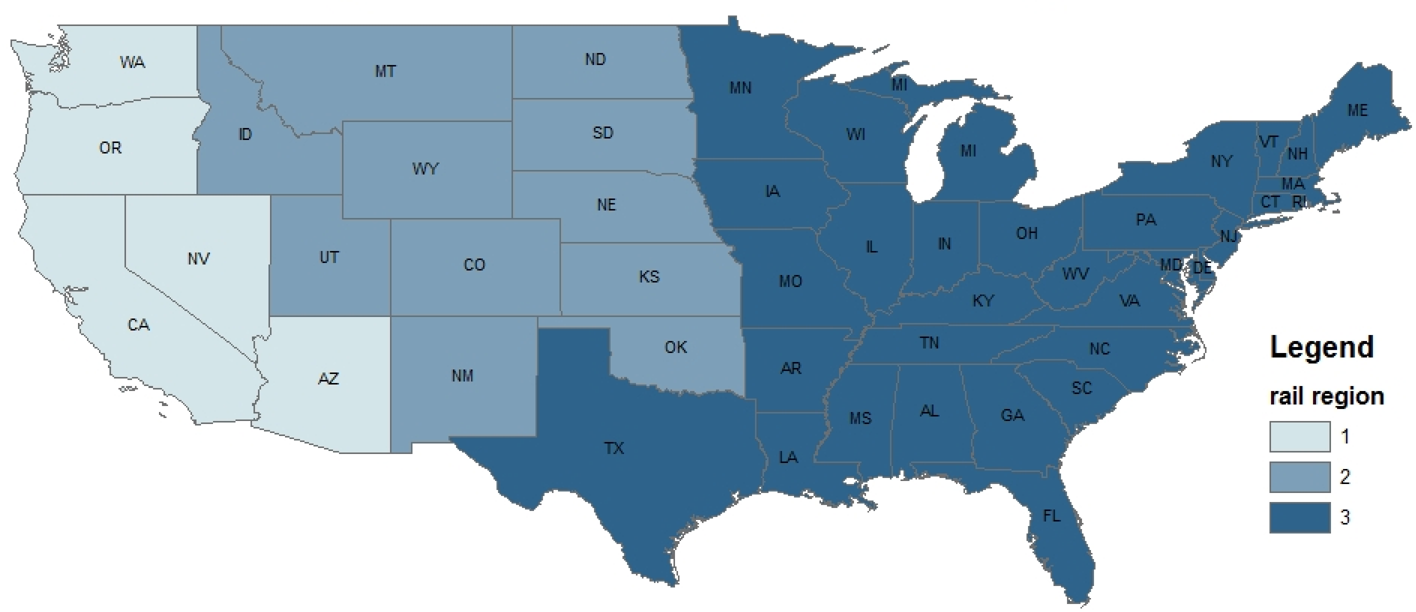

Note: Classification of rail regions based on inspection of 1927 Rand McNally World Atlas.

Figure 4.3: Migrant flows to California and Texas, 1924

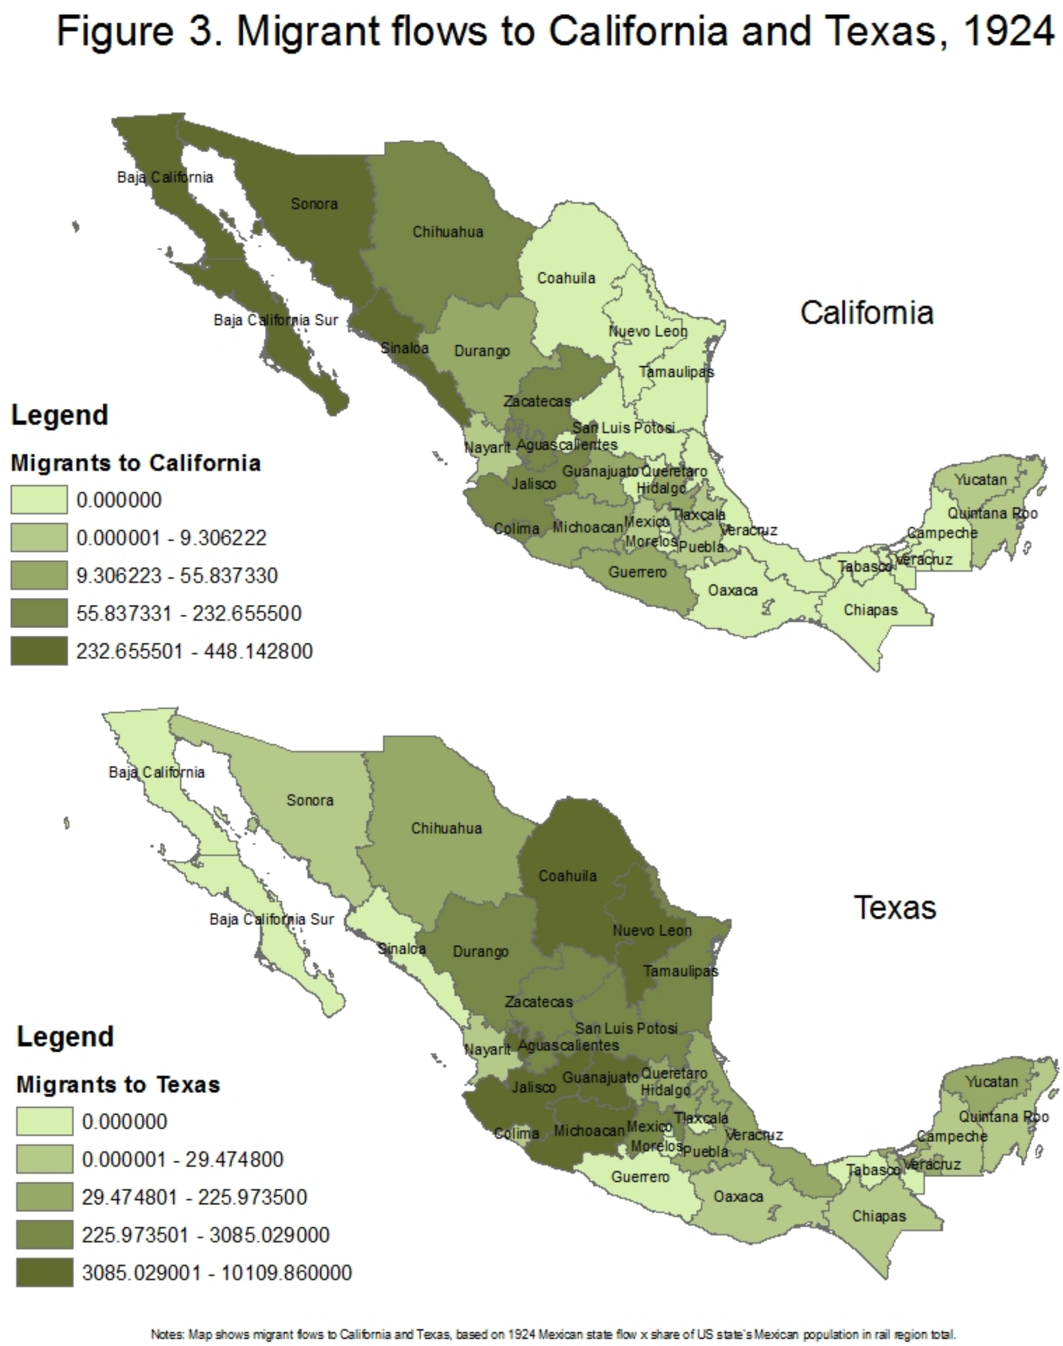

CHAPTER V

Conclusion

The first two essays of this dissertation have documented the difficult environment faced by young adults in South Africa, and have developed economic models of their transition from school to work. The first essay focused on how South African youth make school enrollment choices when their grade advancement and labor market prospects are uncertain. It also documented high rates of re-enrollment in school following periods of dropout, a behavior that is difficult to reconcile with standard models of human capital investment. Estimates of a dynamic sequential enrollment model confirmed the hypothesis that re-enrollment is driven by dynamic updating of the relative returns to schooling versus labor market participation. A simulation based on the estimated model quantified the importance of the option to re-enroll in the school to work transition in South Africa, finding that restricting re-enrollment for those who have not completed high school would increase the proportion of the sample completing 12 years of schooling by 8 percentage points. The results suggest that youth who are unable to freely re-enter school will be less willing to drop out, because they face a higher opportunity cost of school exit.

The second essay focused more narrowly on the job search process of young South Africans, looking in particular at the role of reservation wages in the search for the first job after leaving school. Using survey data on reservation wages, we find evidence that the youth labor market in South Africa is characterized by relatively high rates of job offers, but at wages that tend to be too low to be accepted. We interpret our results as consistent with a labor market that is inefficient at matching employers to employees, leading to a high "implicit refusal" rate. It is therefore unsurprising that we find simulation evidence of a potentially beneficial role for an employer wage subsidy to hire unemployed youth.

These first two essays on how South African youth transition from school to work outcomes suggest an important role for public policy in improving education and labor market outcomes: increasing the opportunity cost of dropout (through restrictions on re-enrollment) and providing incentives to hire unemployed youth could increase schooling attainment, lower unemployment, and increase wages. However, care must be taken to design policies that do not unduly reduce welfare by restricting individual choices and that pass a cost-benefit analysis. Additionally, the interaction between such policies must be considered carefully: subsidies for employers to hire unemployed youth could encourage dropout, even as other policies intend to keep youth enrolled in school. Ensuring that policies are narrowly targeted – for instance, by conditioning receipt of the wage subsidy on high school completion – would be essential to avoid conflicting effects.

The final essay of this dissertation analyzes a longstanding question in labor economics-the effect of immigration on the United State labor market-from a new angle. By making use of rainfall shocks in regions of Mexico that have historically sent migrants to particular destinations in the U.S., we are able to isolate a source of plausibly exogenous variation in the Mexican immigrant presence in U.S. states. This allows us to estimate the labor market effects of Mexican immigration in a manner that alleviates the econometric concerns plaguing much of the existing literature. We find that Mexican immigration is associated with lower wages and higher unemployment among non-Mexicans in the U.S. Our estimates are substantially larger in magnitude than previous estimates of the wage impact of immigration.

BIBLIOGRAPHY

Akee, Randall K. Q., "Who Leaves and Who Returns? Deciphering Immigrant Self-Selection from a Developing Country," 2007.

Altonji, Joseph G., "The Demand for and Return to Education When Education Outcomes are Uncertain," *Journal of Labor Economics*, January 1993, *11* (1), 48–83.

Arcidiacono, Peter, "Ability Sorting and the Returns to College Major," *Journal of Econometrics*, August 2004, *121* (1), 343–75.

Aydemir, Abdurrahman and George J Borjas, "A Comparative Analysis of the Labor Market Impact of International Migration: Canada, Mexico, and the United States," *National Bureau of Economic Research*, 2006, (12327).

____ and George J. Borjas, "Attenuation Bias in Measuring the Wage Impact of Immigration," *Journal of Labor Economics*, January 2011, *29* (1), 69–113.

Banerjee, Abhijit, Sebastian Galiani, Jim Levinsohn, Zoe McLaren, and Ingrid Woolard, "A Symposium on Fostering Growth in South Africa: Why Has Unemployment Risen in the New South Africa?," *Economics of Transition*, 2008, *16* (4), 715–40.

Bartel, Ann P., "Where Do the New U.S. Immigrants Live?," *Journal of Labor Economics*, October 1989, *7* (4), 371–391.

Becker, Gary S., *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, 3 ed., University Of Chicago Press, March 1994.

Belzil, Christian and Jörgen Hansen, "Unobserved Ability and the Return to Schooling," *Econometrica*, 2002, *70* (5), 2075–2091.

Berg, Gerard J. Van Den, "Nonstationarity in Job Search Theory," *The Review of Economic Studies*, April 1990, *57* (2), 255–277.

____ and Geert Ridder, "An Empirical Equilibrium Search Model of the Labor Market," *Econometrica*, 1998, *66* (5), 1183–1221.

Borjas, George J., "The Economics of Immigration," *Journal of Economic Literature*, December 1994, *32* (4), 1667–1717.

____ , "Does Immigration Grease the Wheels of the Labor Market?," *Brookings Papers on Economic Activity*, January 2001, *2001* (1), 69–119.

____ , "The Labor Demand Curve Is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market," *The Quarterly Journal of Economics*, November 2003, *118* (4), 1335–1374.

____ and Lawrence F. Katz, "The Evolution of the Mexican-Born Workforce in the United States," *Mexican Immigration to the United States. 2007*, 2007, pp. 13–55.

____ , Jeffrey Grogger, and Gordon H. Hanson, "Imperfect Substitution between Immigrants and Natives: A Reappraisal," *National Bureau of Economic Research Working Paper Series*, March 2008, *No. 13887.*

____ , Richard B. Freeman, and Lawrence F. Katz, "Searching for the Effect of Immigration on the Labor Market," *The American Economic Review*, May 1996, *86* (2), 246–251.

Bound, John and Alan B. Krueger, "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?," *Journal of Labor Economics*, January 1991, *9* (1), 1–24.

Burns, Justine, *Wage subsidies to combat unemployment and poverty : assessing South Africa's options*, Cape Town: Southern Africa Labour and Development Research Unit, 2010.

Cadena, Brian C., "How Do Immigration Flows Respond to Labor Market Competition from Similarly-Skilled Natives?," 2007.

Cameron, A. Colin and Pravin K. Trivedi, *Microeconometrics: Methods and Applications*, Cambridge University Press, May 2005.

Cameron, Stephen V. and Christopher Taber, "Estimation of Educational Borrowing Constraints Using Returns to Schooling," *The Journal of Political Economy*, February 2004, *112* (1), 132–182.

Cameron, Stephen V and James J Heckman, "The Dynamics of Educational Attainment for Black, Hispanic, and White Males," *Journal of Political Economy*, 2001, *109* (2001), 455–99.

Card, David, "The Impact of the Mariel Boatlift on the Miami Labor Market," *Industrial and Labor Relations Review*, January 1990, *43* (2), 245–257.

____ , "The Causal Effect of Education on Earnings," *Handbook of labor economics. Volume 3A. 1999*, 1999, pp. 1801–63.

____ , "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration," *Journal of Labor Economics*, January 2001, *19* (1), 22–64.

____ , "Immigration and Inequality," *National Bureau of Economic Research*, 2009, (14683).

____ and Ethan G. Lewis, "The Diffusion of Mexican Immigrants during the 1990s: Explanations and Impacts," *Mexican Immigration to the United States*, 2007, pp. 193–227.

____ and John DiNardo, "Do Immigrant Inflows Lead to Native Outflows?," *The American Economic Review*, May 2000, *90* (2), 360–367.

Cardoso, Lawrence A., *Mexican emigration to the United States, 1897-1931: Socio-economic patterns*, University of Arizona Press, 1980.

Carneiro, Pedro and James J. Heckman, "The Evidence on Credit Constraints in Post-Secondary Schooling," *The Economic Journal*, October 2002, *112* (482), 705–734.

Case, Anne, "Does Money Protect Health Status? Evidence from South African Pensions," *Perspectives on the economics of aging. 2004*, 2004, pp. 287–305.

____ and Angus Deaton, "School Inputs and Educational Outcomes in South Africa," *The Quarterly Journal of Economics*, 1999, *114* (3), 1047–1084.

____ and Motohiro Yogo, "Does School Quality Matter? Returns to Education and the Characteristics of Schools in South Africa," *National Bureau of Economic Research*, 1999, (7399).

Christensen, Bent Jesper and Nicholas M. Kiefer, "The Exact Likelihood Function for an Empirical Job Search Model," *Econometric Theory*, December 1991, *7* (4), 464–486.

____ and ____ , "Measurement Error in the Prototypal Job-Search Model," *Journal of Labor Economics*, October 1994, *12* (4), 618–639.

Cortes, Patricia, "The Effect of Low-Skilled Immigration on U.S. Prices: Evidence from CPI Data," *Journal of Political Economy*, 2008, *116* (2008), 381–422.

Dinkelman, Taryn and Vimal Ranchhod, "Evidence on the impact of minimum wage laws in an informal sector: Domestic workers in South Africa," 2010.

Eckstein, Zvi and Gerard J. Van Den Berg, "Empirical Labor Search: A Survey," *Journal of Econometrics*, February 2007, *136* (2), 531–64.

____ and Kenneth I. Wolpin, "Duration to First Job and the Return to Schooling: Estimates from a Search-Matching Model," *The Review of Economic Studies*, April 1995, *62* (2), 263–286.

____ and ____ , "Why Youths Drop Out of High School: The Impact of Preferences, Opportunities, and Abilities," *Econometrica*, November 1999, *67* (6), 1295–1339.

Farre, Lidia, Libertad Gonzalez Luna, and Francesc Ortega, "Immigration, Family Responsibilities and the Labor Supply of Skilled Native Women," *Department of Economics and Business*, 2009.

Fiske, Edward B. and Helen F. Ladd, *Elusive Equity: Education Reform in Post Apartheid South Africa*, Brookings Institution Press, July 2004.

Flinn, C. and J. Heckman, "New Methods for Analyzing Structural Models of Labor Force Dynamics," *Journal of Econometrics*, 1982, *18* (1982), 115–68.

Foerster, Robert Franz, *The racial problems involved in immigration from Latin America and the West Indies to the United States; a report submitted to the Secretary of Labor.*, San Francisco: R and E Research Associates, 1971.

Friedberg, Rachel M. and Jennifer Hunt, "The Impact of Immigrants on Host Country Wages, Employment and Growth," *The Journal of Economic Perspectives*, April 1995, *9* (2), 23–44.

Go, Delfin S, Marna Kearney, Vijdan Korman, Sherman Robinson, and Karen Thierfelder, "Wage Subsidy and Labour Market Flexibility in South Africa," *Journal of Development Studies*, October 2010, *46* (9), 1481–1502.

Hanson, Gordon H., "Illegal Migration from Mexico to the United States," *Journal of Economic Literature*, December 2006, *44* (4), 869–924.

____ and Antonio Spilimbergo, "Illegal Immigration, Border Enforcement, and Relative Wages: Evidence from Apprehensions at the U.S.-Mexico Border," *The American Economic Review*, December 1999, *89* (5), 1337–1357.

____ and Craig McIntosh, "The Great Mexican Emigration," *National Bureau of Economic Research Working Paper Series*, December 2007, *No. 13675.*

Heckman, J. and B. Singer, "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data," *Econometrica*, March 1984, *52* (2), 271–320.

Heckman, J. J and S. Urzua, "The option value of educational choices and the rate of return to educational choices," in "Unpublished manuscript, University of Chicago. Presented at the Cowles Foundation Structural Conference, Yale University" 2008.

Heckman, James J. and Salvador Navarro, "Dynamic Discrete Choice and Dynamic Treatment Effects," *Journal of Econometrics*, February 2007, *136* (2), 341–96.

Heckman, James J, Paul A LaFontaine, and Pedro L Rodriguez, "Taking the Easy Way Out: How the GED Testing Program Induces Students to Drop Out," *National Bureau of Economic Research*, 2008, (14044).

Hertz, Tom, "Refundable The Effect of Minimum Wages on the Employment and Earnings of South Africa's Domestic Service Workers," *W.E. Upjohn Institute for Employment Research*, 2005, (05).

Joensen, J. S., "Academic and Labor Market Success: The Impact of Student Employment, Abilities, and Preferences," *Manuscript, Stockholm School of Economics*, 2008.

Keane, Michael P. and Kenneth I. Wolpin, "The Career Decisions of Young Men," *The Journal of Political Economy*, June 1997, *105* (3), 473–522.

Keane, Michael P and Kenneth I Wolpin, "The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment," *International Economic Review*, 2001, *2001* (4), 1051–1103.

Kugler, Adriana and Mutlu Yuksel, "Effects of Low-Skilled Immigration on U.S. Natives: Evidence from Hurricane Mitch," *National Bureau of Economic Research Working Paper Series*, August 2008, *No. 14293.*

Lam, D., C. Ardington, N. Branson, K. Goostrey, and M. Leibbrandt, "Credit Constraints and the Racial Gap in Post-Secondary Education in South Africa," 2010.

——, Cally Ardington, Nicola Branson, Anne Case, Murray Leibbrandt, Alicia Menendez, Jeremy Seekings, and Meredith Sparks, "The Cape Area Panel Study: A Very Short Introduction to the Integrated Waves 1-2-3-4 Data," October 2008.

——, M. Leibbrandt, and C. Mlatsheni, "Education and youth unemployment in South Africa," *Labour markets and economic development*, 2009, p. 90.

Lam, David, Cally Ardington, and Murray Leibbrandt, "Schooling as a lottery: Racial differences in school advancement in urban South Africa," *Journal of Development Economics*, July 2011, *95* (2), 121–136.

Lancaster, Tony and Andrew Chesher, "An Econometric Analysis of Reservation Wages," *Econometrica*, November 1983, *51* (6), 1661–1676.

Leibbrandt, Murray, James Levinsohn, and Justin McCrary, "Incomes in South Africa after the Fall of Apartheid," *Journal of Globalization and Development*, 2010, *1* (1).

Levinsohn, James, "Two Policies to Alleviate Unemployment in South Africa," *Harvard Center for International Development Working Paper Series*, May 2008, *No. 166.*

—— and Todd Pugatch, "The Role of Reservation Wages in Youth Unemployment in South Africa: A Structural Approach," *Unpublished paper, University of Michigan*, 2011.

Lewis, Ethan G., "Immigration, skill mix, and capital-skill complementarity," *Quarterly Journal of Economics*, forthcoming.

Light, Audrey, "Hazard Model Estimates of the Decision to Reenroll in School," *Labour Economics*, December 1995, *2* (4), 381–406.

Luna, Libertad Gonzalez and Francesc Ortega, "Immigration and Housing Booms: Evidence from Spain," *Department of Economics and Business*, 2009.

Lynch, Lisa M., "Job Search and Youth Unemployment," *Oxford Economic Papers*, November 1983, *35*, 271–282.

Massey, Douglas S., Jorge. Durand, and Nolan J. Malone, *Beyond smoke and mirrors: Mexican immigration in an era of economic integration*, New York: Russell Sage Foundation, 2002.

Munshi, Kaivan, "Networks in the Modern Economy: Mexican Migrants in the U. S. Labor Market," *The Quarterly Journal of Economics*, May 2003, *118* (2), 549–599.

Nattrass, Nicoli and Richard Walker, "Unemployment and Reservation Wages in Working-Class Cape Town," *South African Journal of Economics*, September 2005, *73* (3), 498–509.

Orrenius, Pia M. and Madeline Zavodny, "Self-Selection among Undocumented Immigrants from Mexico," *Journal of Development Economics*, October 2005, *78* (1), 215–40.

Ottaviano, Gianmarco I.P. and Giovanni Peri, "Immigration and National Wages: Clarifying the Theory and the Empirics," *National Bureau of Economic Research Working Paper Series*, July 2008, *No. 14188.*

Pauw, Kalie and Lawrence Edwards, "Evaluating the General Equilibrium Effects of a Wage Subsidy Scheme for South Africa," *South African Journal of Economics*, September 2006, *74* (3), 442–62.

Peri, Giovanni, "Rethinking the Area Approach: Immigrants and the Labor Market in California, 1960-2005.," *National Bureau of Economic Research Working Paper Series*, July 2010, *No. 16217.*

___ and Chad Sparber, "Task Specialization, Immigration, and Wages," *American Economic Journal: Applied Economics*, July 2009, *1* (3), 135–69.

Pugatch, Todd, "Bumpy Rides: School to Work Transitions in South Africa," *Unpublished paper, University of Michigan*, 2011.

Rand McNally World Atlas. Premier Edition., New York: Rand McNally & company, 1927.

Rust, John, "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 1987, *55* (5), 999–1033.

Saiz, Albert, "Room in the Kitchen for the Melting Pot: Immigration and Rental Prices," *The Review of Economics and Statistics*, 2003, *85* (3), 502–521.

___ , "Immigration and Housing Rents in American Cities," *Journal of Urban Economics*, March 2007, *61* (2), 345–71.

Schoer, Volker and Murray Leibbrandt, "Determinants of Job Search Strategies: Evidence from the Khayelitsha/Mitchell's Plain Survey," *South African Journal of Economics*, December 2006, *74* (4), 702–24.

Stange, Kevin, "An Empirical Examination of the Option Value of College Enrollment," *Unpublished paper, University of Michigan*, 2009.

Stock, James H. and Motohiro Yogo, "Testing for Weak Instruments in Linear IV Regression," *National Bureau of Economic Research Technical Working Paper Series*, November 2002, *No. 284.*

Wolpin, Kenneth I., "Estimating a Structural Search Model: The Transition from School to Work," *Econometrica*, July 1987, *55* (4), 801–817.

Woodruff, Christopher and Rene Zenteno, "Migration Networks and Microenterprises in Mexico," *Journal of Development Economics*, March 2007, *82* (2), 509–28.

Yamada, H., "The Impact of the Introduction of Sectoral Minimum Wages on Low Wage Markets in a Low Income Country: Evidence from South Africa," 2007.

Yamauchi, Futoshi, "Race, Equity, and Public Schools in Post-apartheid South Africa: Equal Opportunity for All Kids," *Economics of Education Review*, April 2005, *24* (2), 213–33.

Yang, Dean, "Coping with Disaster: The Impact of Hurricanes on International Financial Flows, 1970-2002," *B.E. Journal of Economic Analysis and Policy: Advances in Economic Analysis and Policy*, 2008, *8* (1).

___ and HwaJung Choi, "Are Remittances Insurance? Evidence from Rainfall Shocks in the Philippines," *World Bank Economic Review*, 2007, *21* (2), 219–48.