Essays on Local Labor Markets

by

Emek Basker

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Economics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2002

Signature of Author ....
Department of Economics
August 15, 2002

Certified by ....
Daron Acemoglu
Professor of Economics
Thesis Supervisor

Certified by ....
Olivier Blanchard
Professor of Economics
Thesis Supervisor

Accepted by ....
Peter Temin
Elisha Gray II Professor of Economics
Chairman, Departmental Committee on Graduate Studies

Essays on Local Labor Markets

by

Emek Basker

Submitted to the Department of Economics
on August 15, 2002, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This dissertation consists of three chapters.

The first chapter explores the effect of Wal-Mart expansion on local retail employment. The phenomenal expansion of Wal-Mart provides a clean case for studying the labor-market effects of increased efficiency. I estimate the effect of Wal-Mart entry on retail employment at the county level. Using an instrumental-variables approach to correct for both measurement error in entry dates and possible endogeneity of the timing of entry, I find that Wal-Mart entry increases retail employment by 100 jobs in the year of entry. Half of this gain disappears over the next five years, leaving a statistically significant net gain of 50 jobs at the five-year horizon. The decline in retail employment in the years immediately following entry is associated with the closing of both small and large retail establishments. At the same time, retail employment in neighboring counties declines by approximately 30 jobs, and wholesale employment in the entered county declines by a similar number.

In the second chapter, I explore several differences between the migration and job-search behavior of workers with different levels of education, both theoretically and empirically. I start with two stylized facts. First, the propensity to migrate increases with education. Second, conditional on migration, the probability that a worker moves with a job in hand (rather than moving to search for a job in the new location) also increases with education. I present a simple consumer-choice model that captures these facts and generates a number of predictions about differential sensitivity of migration to observed variables by education. These predictions are verified using CPS data.

The third chapter documents a small permanent effect of idiosyncratic shocks to agricultural revenues – due to variation in crop yields and crop prices – on the number of farms in the United States. Using county-level data on the number of farms by size and ownership structure (family-owned vs. corporate-owned), I show that following negative deviations from expected revenue, the number of farms declines; this decline is disproportionaly due to a decline in the numbers of small and/or family-owned farms.

Thesis Supervisor: Daron Acemoglu
Title: Professor of Economics

Thesis Supervisor: Olivier Blanchard
Title: Professor of Economics

Acknowledgments

My years at MIT have been shaped by the people around me: my teachers, my classmates, and my students. Each of them contributed enormously to my experience (almost all positively).

Most importantly, I thank my two main advisors, Daron Acemoglu and Olivier Blanchard, for their patience, support, guidance and good humor. Working with them for the past three years has been both fun and rewarding. They have each taught me a great deal – in no small part by their disagreements with one another.... I thank them tremendously for giving me the skills, and the confidence, to make my own choices.

A number of other faculty members have helped me along the way. Sendhil Mullainathan provided many insightful comments and a fresh perspective – when I could find him. Josh Angrist has been helpful and supportive throughout my years at MIT. And Frank Fisher has been a friend and a mentor whose presence has greatly enriched my graduate-school experience.

My classmates deserve many thanks for tolerating my sarcasm and helping out in many small and big ways. Saku Aura was always available to discuss my research and provided many thoughtful comments and regular doses of encouragement. Amy Finkelstein has been my inspiration, both for the joy she takes in her research and for the honesty with which she conducts it. And so many others – too numerous to name – provided advice and support along the way. I will miss our coffee breaks, our walks, and the fascinating conversations we inevitably had over lunch – about politics, people, and banana-peeling techniques.

Finally, I wish to thank my family and friends outside MIT whose support and encouragement have kept me (relatively) sane throughout graduate school. Dorothy deserves special credit for providing a firm anchor when I drifted and for forgiving my moodiness almost daily. My father and sisters have been supportive and generous over my 22(!) years of schooling, never questioning whether I have overshot the optimal level of education. My mother, to whom I dedicate this dissertation, is not around to see me complete my education, but I know she would have been proud.

Formal Acknowledgements

I wrote these three papers under the tutelage of Daron Acemoglu and Olivier Blanchard. Their comments and suggestions have improved all three papers beyond measure. Despite their best efforts some errors may remain; they are my own.

The first chapter could not have been written without the help of Dorothy Carpenter, Kathy Cosgrove, Rich Lindrooth, Erich Muehlegger, Mike Noel, Jon Zinman, and especially Maurice Drew, Catherine Friedman and Steven Sadoway, all of whom helped me gather Wal-Mart data. The chapter benefitted substantially from the comments of Josh Angrist, Saku Aura, David Autor, Glenn Ellison, Amy Finkelstein, Guido Kuersteiner, Jeffrey Miron, Sendhil Mullainathan, Whitney Newey, Marko Terviö, Ken Troske and seminar participants in numerous institutions. That chapter was completed while visiting the Federal Reserve Bank of St. Louis, which I thank for its hospitality. Needless to say, the views expressed are not necessarily those of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

The second chapter was much improved thanks to comments from Saku Aura, Hoyt Bleakley, Al Nucci and others. Pegah Ebrahimi and Amy Mok helped with typing.

The third chapter was completed in record time with tremendous help from Andrew Hertzberg and Barrett Kirwan; Josh Angrist, Saku Aura, Bengte Evenson, Guido Kuersteiner, Robin McKnight and Steve Taff all provided helpful insights.

I thank the MIT Department of Economics and the MIT Schultz Fund for financial support.

For my mother

Chapter 1

Competition, Efficiency and Employment: Labor-Market Effects of Wal-Mart Expansion

> "By contributing overwhelmingly to the productivity growth jump in general merchandise retail, Wal-Mart demonstrates the impact that managerial innovation and effective use of IT can have on market structure, conduct, and performance."
> – McKinsey Global Institute, 2001

1.1 Introduction

A recent study by McKinsey Global Institute (2001) attributes the increase in the productivity growth rate in the 1990s to only six industries: retail, wholesale, securities, telecommunications, semiconductors, and computer manufacturing. New research has also concluded that nearly all of the gains in productivity within the retail industry have occurred through entry (of firms throughout the productivity distribution) and exit (of the least-efficient firms) (Foster, Haltiwanger and Krizan 2001). Combined, these two facts suggest that we can look at entry of efficient firms into the retail industry to understand the effects of increased productivity on employment.[1]

[1] To see why the sign of the effect is ambiguous, note that firms with higher productivity need to hire fewer workers to produce a constant level of goods or services, but as they lower prices, quantity demanded increases as

In this paper, I use a case-study approach to track the effect of one specific type of productivity shock – entry of Wal-Mart – on sectoral and total county-level employment. Wal-Mart, the largest as well as the most efficient retailer in the world, entered local markets gradually beginning in 1962, creating a natural experiment for studying the effect of the introduction of new technology on labor markets. One industry study estimates that Wal-Mart requires only 75% of other retailers' employment levels to achieve the same revenue stream (Muller 1999). Analysts at Lehman Brothers have noted Wal-Mart's "operational and technological superiority" over all its competitors (Feiner 2001). And a Wal-Mart competitor was quoted in the Washington Post saying, "The real problem [with Wal-Mart] is that they're so good at what they do" (1990).

Using an instrumental-variables specification that exploits the variable lag between store-planning dates and store-opening dates, I am able to correct for both measurement error in reported entry dates and endogeneity of the entry decision. I find a positive effect on retail employment in the entered county even at the five-year horizon. Retail employment in neighboring counties declines, as does wholesale employment in the entered county. The combined effect on the retail and wholesale sectors and neighboring counties yields a net nil effect on employment in the distributive trades in the entered region. I find no effect on aggregate county-level employment.

This paper contributes to a related discussion on the effect of changes in industry structure and industry-level technology on employment. This area has not received much attention by economists, but a couple of recent papers have begun to explore the relationship between competition, ownership structure, and labor-market outcomes. In one recent paper, Bertrand and Kramarz (forthcoming) examine entry regulation in the French retail industry over the past three decades. The authors find that regulation limiting entry of large retailers slowed employment growth in the French retail industry. Their finding is consistent with my finding that Wal-Mart entry increases local retail employment.[2]

well. The latter effect is due to a combination of demand substitution between firms in the industry and demand substitution across industries. Which effect dominates depends on the elasticity of demand for the industry's product. This intuition is formalized in a simple model in Appendix 1.A.

[2] Several studies have aimed to address the effect of Wal-Mart entry on retail employment, and on the number of small retail businesses, directly. Of these, the most cited are a series of case studies by Kenneth Stone (1989, 1991, 1997) arguing that Wal-Mart entry has had a negative impact on small retailers in rural Iowa. Others have

The remainder of the chapter is organized as follows. Section 1.2 provides background information on the retail industry in general and Wal-Mart in particular. Section 1.3 describes the data. My empirical strategy is explained in Section 1.4, and results are presented and discussed in Section 1.5. Section 1.6 concludes.

1.2 Background

1.2.1 The Retail Industry

In a recent study of churning in the retail industry over the period 1987-1997, Foster, Haltiwanger and Krizan (2001) find a very large dispersion in the distribution of productivity (measured as the difference between log real output and log labor input), relative to the dispersion found in the manufacturing industries. They also find extremely high rates of job reallocation, due mostly to entry and exit: 70% of gross job and output creation (destruction) in the retail industries is accounted for by firm entry (exit). While new firms enter throughout the productivity distribution, exit is concentrated at the lower tail of the productivity distribution; this fact drives productivity growth in the retail industry. Unlike the manufacturing industries, Foster, Haltiwanger and Krizan claim, the retail industry would have experienced no productivity growth over the period studied without entry and exit.

Mass merchandisers, such as Wal-Mart and its main competitors K-Mart and Target, differ from traditional retailers in that they sell a large variety of products at low prices. Data from the 1997 Census of Retail Trade show that about 7.5% of retail workers were employed by "discount or mass merchandising department stores" the week of March 12, 1997. A recent study of international productivity differences across industries found that traditional stores in the United States are only 60% as productive as U.S. mass merchandisers (Baily and Solow 2001).

found more positive effects on other areas (Franz and Robb 1990, Barnes et al. 1996).

1.2.2 Wal-Mart

> "The purchasing power of a chain unquestionably gives it certain advantages. But I believe that it owes much, if not most, of its success to the intelligence with which it is operated."
> – W.D. Darby, *The Story of The Chain Store*, 1928.

The first Wal-Mart store opened in Benton County, Arkansas in 1962. By the time the company went public in 1969, it had 18 stores throughout Arkansas, Missouri, and Oklahoma. The company slowly expanded its geographical reach, building new stores and accompanying distribution centers further and further away from its original location, and continued, at the same time, to build new stores in areas already serviced. Figure 1-1 shows maps of the 48 contiguous states with approximate locations of Wal-Mart stores over time to illustrate this point. By 1998 Wal-Mart had approximately 2400 stores in all 50 states and about 800,000 employees in the United States. At the end of 2001 Wal-Mart had 1.2 million employees worldwide, of which about 962,000 (77%) were employed in the United States.

Wal-Mart is the largest retailer in both the United States and the world. The company operates Wal-Mart discount stores as well as Wal-Mart "supercenters" which include grocery departments and constitute approximately one third of all current Wal-Mart stores. The typical Wal-Mart store spans 100,000-150,000 square feet and employs 150-350 people, many of them in part-time jobs. By 1998, one quarter of the 1614 counties entered by Wal-Mart had more than one store; of these, 234 counties had two stores, and 151 counties had 3 or more stores (among them Harris County, Texas, with 19 stores in 1998).

Wal-Mart is extremely efficient even compared with other "big-box" retailers. It has been cited for its technological advantage by many industry analysts. Lehman Brothers analysts have noted Wal-Mart's "leading logistics and information competencies" and compared it favorably to the world's second-largest retailer, Carrefour, saying that Wal-Mart's "operational and technological superiority has allowed Wal-Mart to gain a comparative advantage over every competitor it has faced, including Carrefour" (Feiner 2001). The *Financial Times* is more expressive, calling Wal-Mart "an operation whose efficiency is the envy of the world's storekeepers" (Edgecliffe-Johnson 1999).

Wal-Mart's competitive edge is driven by a combination of conventional cost-cutting and

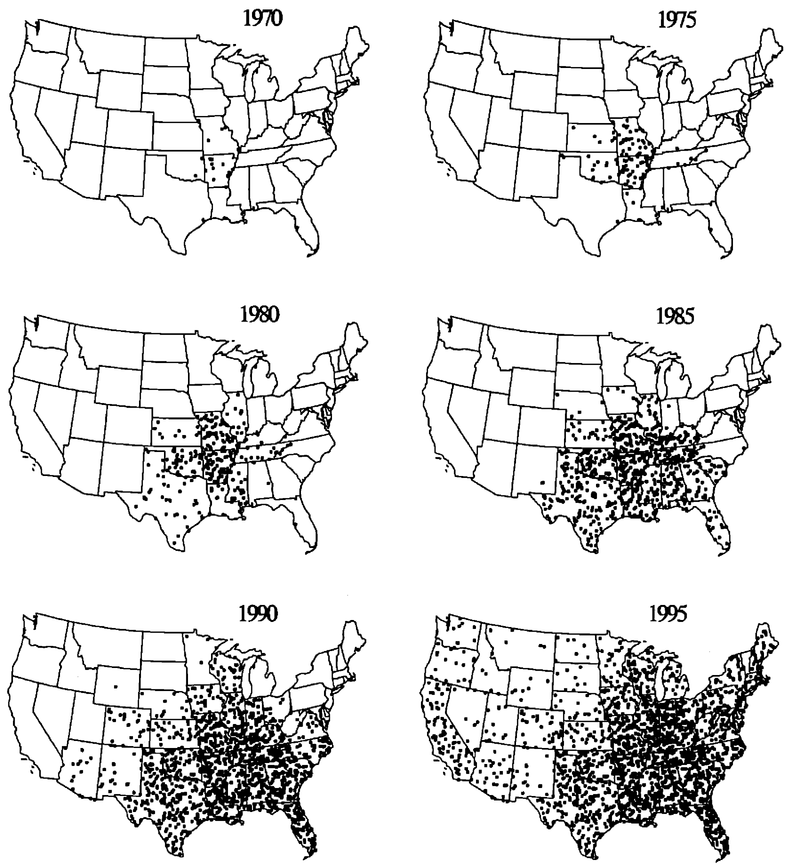

Figure 1-1: Location of Wal-Mart Stores, Various Years

sensitivity to demand conditions, and by superior technology: the company uses software-based logistics and distribution systems, and its divisions are well integrated. Wal-Mart's most cited advantage over small retailers is probably economies of scale and access to capital markets, whereas against other large retailers, such as K-Mart and Target, the most cited factors are the following:[3]

- Superior logistics, distribution, and inventory control: Wal-Mart's proprietary software, Retail Link, links stores directly to Wal-Mart's distribution centers, and links those directly with suppliers (like GE and Proctor & Gamble) who get daily sales data and are able to plan their own inventories accordingly. This system has reduced Wal-Mart's inventory costs to levels substantially below its competitors' (Stalk and Hout 1990).

- Size: This distribution network is made even more efficient by the geographic proximity of its many stores; Wal-Mart's size also gives it market power in some goods as well as input markets.

- Cost-conscious "corporate culture".

- Demand-sensitivity: Inventories and prices differ from store to store based on climate and consumer demographics. Reorders are made based on actual store needs (communicated to the nearest distribution center) rather than centralized forecasting, and pricing is competitive given market conditions.

There is no single best measure of productivity in the retail industry. One commonly used measure is sales per square foot. Figure 1-2 shows sales per square foot at K-Mart and Wal-Mart stores (in nominal dollars) for selected years. By way of comparison, a series of studies by the Urban Land Institute put average sales per square foot for mall stores slightly below K-Mart's sales over the period 1978-1997 (Urban Land Institute, various years). Wal-Mart does well also by other measures of productivity. Figures for sales for employee, cited in Johnson (2002), show Wal-Mart consistently ahead of other firms by a large margin; the effect of Wal-Mart on county-wide sales per employee are investigated in Appendix 1.C.

[3]Many of the details cited here on Wal-Mart's operations are from Harvard Business School's three Case Studies about Wal-Mart (Ghemawat 1989, Foley and Mahmood 1996, and Ghemawat and Friedman 1999). Similar points are also made by McKinsey Global Institute (2001).

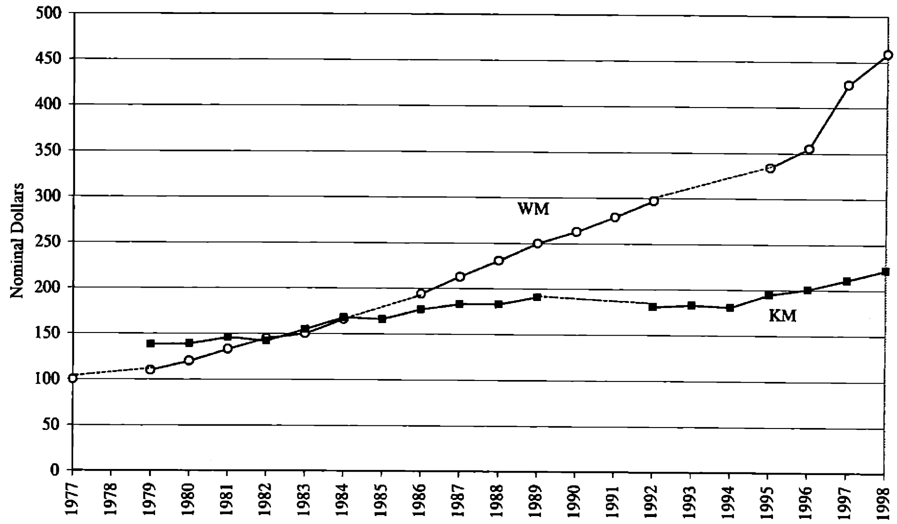

Figure 1-2: Annual Sales per Square Foot, Wal-Mart vs. K-Mart

1.3 Data

1.3.1 Wal-Mart Stores

I use data on the locations and opening dates of 2,382 Wal-Mart stores in the United States, collected primarily from Wal-Mart annual reports, Wal-Mart editions of Rand McNally Road Atlases and annual editions of the Directory of Discount Department Stores. The available data include store location (by town) and store number.

The following data sources, which I refer to collectively as "directories", provide one measure of opening dates: Vance and Scott (1994) list store entries to 1969, the year the company become publicly traded. Annual reports between 1970 and 1978 include lists of current stores. After 1978 annual reports became largely uninformative, listing only the current number of stores per state. The annual *Directory of Discount Department Stores* provides store lists between 1979 and 1993. The directory is published in the beginning of each calendar year, and contains the store list for the end of the previous calendar year. Finally, for recent years I use a special edition of the popular Rand McNally road atlas, sold only at Wal-Mart stores, which contains a list of store locations, and includes each store's company-assigned number. The variable $\mathbf{WMlist_{jt}}$

gives the number of new stores to open in county j in year t based on these directories and store lists.

I also construct an alternate set of Wal-Mart entry identifiers using a combination of company-assigned store numbers (from the Rand McNally atlases) and the net change in the number of stores each year (from company annual reports). This alternate set of entry dates is then used in an instrumental-variables specification to correct for measurement error in, and potential endogeneity of, the timing of entry. Wal-Mart assigns store numbers roughly in sequential order, with store #1 opening first, followed by store #2, and so on. I therefore assign entry dates to stores sequentially, based on their store numbers. This assignment method provides a very good approximation to the true distribution of entry dates. Aggregating these store-level entry dates to the county-year level, I construct $\mathbf{WMnum_{jt}}$: the number of Wal-Mart stores in county j whose store numbers correspond to those opened in year t.

For more details on the construction of the Wal-Mart variables, see Appendix 1.B.

1.3.2 Labor Market Data

My unit of observation is a county-year. Although there are currently 3111 counties in the contiguous 48 states, some counties have been created (usually by splitting one county in two) and others have merged over the period studied; in those cases, I have merged the observations into one observation for the entire study period.[4] I limit the data set to the 1777 counties with 1964 employment above 1500, positive employment growth between 1964 and 1977, and no Wal-Mart entry prior to 1977.

Annual county-level employment by SIC (or NAICS) for 1977-1999 comes from the Census Bureau's County Business Patterns (CBP) serial. The panel contains 40,871 observations (1777 counties * 23 years).[5] Unfortunately, no wage data are available from CBP.

Table 1.1 lists some summary statistics for labor-market data. More details are available in Appendix 1.B.

[4] For details on these newly-created and merged counties, see Appendix 1.B.

[5] The relevant SIC (NAICS) codes are:
Retail: SIC 52-- except 5800, NAICS 44
Wholesale: SIC 50--, NAICS 42
Restaurants: SIC 5800, NAICS 721.

Table 1.1: Summary Statistics

	Sample Counties	Excluded Counties
Total Employment (Mean)	42,000	6,000
Total Employment (Median)	11,000	1,500
Fraction Retail Employment	15.4%	16.8%
Fraction Wholesale Employment	5.0%	5.8%
Fraction with Wal-Mart	75%	13%
Median Number of Small Establishments[a]	172	37
Median Number of Medium Establishments[a]	13	1
Median Number of Large Establishments[a]	1	0

[a] Small establishments: 1-19 employees; medium: 20-99; large: 100+

1.4 Methodology

1.4.1 OLS Regressions

Because the data do not appear to contain unit roots, the analysis is done using employment levels (see Appendix 1.B for details on unit root tests). The regressions are:

$$\frac{\text{retail}_{jt}}{\text{pop}_{jt}} = \alpha + \sum_k \sum_t \delta_{tk}\, \text{urban}_{jk}\, \text{year}_t + \sum_j \psi_j \text{county}_j + \theta(L)\frac{\text{WalMart}_{jt}}{\text{pop}_{jt}} + \text{u}_{jt} \quad (1.1)$$

where retail_{jt} is retail employment in county j in year t; pop_{jt} is the population of county j in year t; year_t is a year dummy; $\text{urban}_{jk} \in \{\text{urban, suburban, rural}\}$ is an urbanization dummy allowing for different year fixed effects for urban, suburban, and rural counties;[6] WalMart_{jt} is the number of new Wal-Mart stores built that year in county j; county_j is a county dummy; and $\theta(L)$ is a lag polynomial with six lags and five leads (the sixth lag represents the collective period "six or more years after year t"; the omitted category (reference period) is six or more years before a given store was opened.[7] Note that employment 6 or more years before entry is

[6] Urban_{jk} = urban if county j was inside an MSA (metropolitan statistical area) in 1960; suburban if it was ≤25 miles from the nearest MSA in 1960; and rural otherwise.

[7] In other words,
$\theta(L) = \theta_1 F^5 + \theta_2 F^4 + \theta_3 F^3 + \theta_4 F^2 + \theta_5 F + \theta_6 + \theta_7 L + \theta_8 L^2 + \theta_9 L^3 + \theta_{10} L^4 + \theta_{11} L^5 + \theta_{12} \sum_{\tau \geq 6} L^\tau$ where L is the lag operator and F is the lead operator.

normalized to be zero for all counties. The error term u_{jt} is clustered at the county level.[8]

Both employment and the number of Wal-Mart stores are divided by the current county population, so the coefficients $\theta(L)$ can be interpreted as the effect of one additional Wal-Mart store per-capita on retail employment per capita.[9] Plots of the coefficients $\theta(L)$ are therefore used to show the evolution of employment over a 10-year period, starting five years before and ending five years after Wal-Mart entry into a county. The coefficient θ_{12}, intended to capture the permanent effect of Wal-Mart entry on employment six or more years after entry, is omitted from the graphs because it is identified using a relatively small number of observations.

The OLS estimates are valid if Wal-Mart entry is correctly measured and exogenous to employment changes. Unfortunately, neither measure of Wal-Mart entry – WMlist_{jt}, which uses directory data, and WMnum_{jt}, which uses store numbers to impute opening dates – is measured without error. Concerns about endogeneity in the timing of entry offer further complication. An instrumental-variables specification is therefore used to correct these problems.

1.4.2 Measurement Error

Measurement error in the Wal-Mart entry variables – WMlist_{jt} and WMnum_{jt} – takes a particular form: while the entered counties are correctly identified, the *timing* of entry may be incorrectly measured for a variety of reasons.

The timing error in WMlist_{jt} is due to errors in the directories. A particular egregious example of such errors is the lack of updating of the *Directory of Discount Department Stores* between 1990 and 1993. In addition, stores may appear in directories a year or two late, planned stores may appear before they open, and typos can cause a single store to appear multiple times in one year. The error in WMnum_{jt}, constructed using store numbers, is due to the variable lag between store planning and store entry, as well as to random assignment of store numbers

[8] There has been some confusion in the literature about the use of clustered standard errors with fixed-effect models (see `www.stata.com/support/faqs/stat/aregclust.html`). Kézdi (2001) and Bertrand, Duflo, and Mullainathan (2001) show that clustering with fixed-effect models is not a problem, and is, in fact, generally recommended when the autocorrelated process is not well understood.

[9] The use of per-capita terms on both the left- and right-hand sides of Equation (1.1) could in principle cause a spurious correlation between the variables that would bias the estimated coefficients. In practice, however, the year-to-year variation in county population is small enough that it is not driving the results presented here; similar results arise when retail employment and WMlist are normalized by a constant such as the 1990 population of the county.

to approximately 40 stores whose numbers are not known.

An instrumental-variables approach, in which one variable is used to instrument for the other, can be used to correct for this measurement error if the measurement errors in the two variables is classical and uncorrelated. That the measurement error across the two variables is uncorrelated seems plausible.[10] Because WMlist_{jt} and WMnum_{jt} are discrete, however, their measurement error is not classical, as noted by Kane, Rouse, and Staiger (1999). This induces bias in the instrumental-variables results reported here.[11]

The 12 first-stage regressions are

$$\frac{\text{WMlist}_{j,t-s}}{\text{pop}_{jt}} = \widetilde{\alpha} + \sum_k \sum_t \widetilde{\delta}_{tk}\, \text{urban}_{jk}\, \text{year}_t + \sum_j \widetilde{\psi}_j \text{county}_j + \widetilde{\theta}(L) \frac{\text{WMnum}_{jt}}{\text{pop}_{jt}} + \widetilde{\text{u}}_{jt} \quad (1.2)$$

where $s = -6^+, -5, ..., 4, 5$. Predicted values are denoted by $\frac{\widehat{\text{WMlist}_{jt}}}{\text{pop}_{jt}}$, and the second stage is

$$\frac{\text{retail}_{jt}}{\text{pop}_{jt}} = \alpha + \sum_k \sum_t \delta_{tk}\, \text{urban}_{jk}\, \text{year}_t + \sum_j \psi_j \text{county}_j + \theta(L) \frac{\widehat{\text{WMlist}_{jt}}}{\text{pop}_{jt}} + \text{u}_{jt}. \quad (1.3)$$

1.4.3 Endogeneity

Another difficulty in assessing the impact of Wal-Mart entry on the level and composition of county employment is the possible endogeneity of Wal-Mart's entry decision. This endogeneity has two dimensions: Wal-Mart chooses both the locations to enter (entry dimension) and the timing of entry into those counties (timing dimension).

If Wal-Mart selects counties whose growth rates exceed those of non-entered counties, a spurious positive effect will be registered by the estimated coefficients $\widehat{\theta}(L)$. To address this concern, I limit the analysis to counties that constitute a good control group for entered coun-

[10]This assumption would be violated if some stores, for example in metropolitan areas, experience shorter planning phases – for example due to quicker zoning changes – and were also more likely to appear in the directories sooner, due to better directory coverage. This does not appear to be the case.

[11]Kane, Rouse and Staiger (1999) suggest a GMM estimator to address this problem. Unfortunately, due to the size of the panel and the hundreds of covariates, their solution is not computationally feasible in this setting.

ties: counties with a 1964 population above 1500 and a positive average growth rate of total employment between 1964 and 1977. Finally, I remove counties entered by Wal-Mart before 1977 to eliminate concerns about the endogeneity of employment growth. Wal-Mart entered 75% of the remaining 1777 counties between 1977 and 1998, compared with only 13% of the excluded counties.[12]

Moreover, the timing of entry may be endogenous to employment outcomes if Wal-Mart enters counties during growth spurts (or, what is less likely, during temporarily slowdowns). If entry is timed to coincide with growth spurts, estimated coefficients would reveal a spurious positive relationship between Wal-Mart entry and employment growth.

The instrumental-variables strategy described also corrects for this endogeneity concern. The correction is valid if store numbers represent store planning dates, plans are made well in advance of entry, and, as above, the lag between the two measures of entry is independent of employment outcomes. The first condition appears to hold. Consider, for example, stores 762, 763, and 764. All three are located in Jefferson County, Alabama, and their sequential numbering suggests they were probably planned together. Two of them (763 and 764) opened in 1984, while the third opened in 1990. If this difference in entry dates is random, the IV strategy should be valid.[13]

The second condition is that planning must be done sufficiently in advance of entry that spurts in employment growth cannot be reasonably forecasted for the year of entry. Ideally, the order of store entry would have been planned before any stores opened (circa 1960). A more likely scenario is that Wal-Mart determines entry into blocks of counties at regular intervals, which would reduce the endogeneity of the timing of entry.

Finally, lags in entry must be uncorrelated with employment outcomes. This assumption seems reasonable in general, but it would not hold if, for example, towns that resisted and delayed Wal-Mart entry had a disproportionate number of inefficient incumbents that closed after Wal-Mart's entry, or if the residents of such towns were more (or less) avid shoppers than residents of other towns.

[12] Indistinguishable results are obtained if the sample is limited instead to entered counties.

[13] One way to check the validity of this assumption is to regress the difference in assigned entry dates based on the two sources on county characteristics. Such regressions consistently yield no relationship between county characteristics (such as size and urbanization) and the difference between the two entry dates.

If the instrumental-variables strategy outlined above does not correct for endogeneity, the lead coefficients would most likely betray this fact. In other words, if Wal-Mart times entry to take advantage of temporary retail growth spurts, then unless it times its entry perfectly, we would expect to see some increase in retail employment in the years before Wal-Mart entry. As the results below show, for the most part this is not the case; in other words, the IV strategy appears to correct for endogeneity as well as measurement error.

1.5 Results

1.5.1 Retail Employment

To begin, I present OLS results using the two alternative measures of Wal-Mart entry dates. Figure 1-3 shows the OLS estimates using the RHS variable WMlist_{jt}; Figure 1-4 shows OLS estimates for the same regressions, using WMnum_{jt} instead.[14] In both cases, retail employment increases by an estimated 40 jobs in the year of entry, up to half of which are eliminated within five years. In both cases as well, 20 jobs are estimated to have been created in the year *before* Wal-Mart entry. While this number is small in absolute magnitude, it is disconcertingly large relative to the estimated post-entry effect.

The IV results are shown in Figure 1-5. The effect of entry is estimated much more cleanly at approximately 100 jobs. In the years immediately following entry, there is a loss of 40-70 additional jobs. The net effect at the five-year horizon, however, is positive and significant (p-value 0.0003).[15]

Recall that the typical Wal-Mart store employs 150-350 workers. These results suggest that employment increases by less than the full amount of Wal-Mart's hiring, even before allowing other firms time to fully adjust to Wal-Mart's entry. Part of this discrepancy can be explained by buyouts of existing chain stores by Wal-Mart Corporation, and prompt exit and cutbacks by other retailers.[16] Another (albeit unlikely) possibility is that Wal-Mart replaces existing

[14]Like all regression results presented here, unless otherwise noted, the 95% confidence intervals shown use asymptotic standard errors and allow for any intertemporal correlation of errors for a given county.

[15]The long-run effect, six or more years after entry (not shown in Figure 1-5) is a net increase of 15 jobs over employment in year 0; this long-run increase is not statistically significant (p-value 0.3749).

[16]This possibility is explored in more detail in Section 1.5.2.

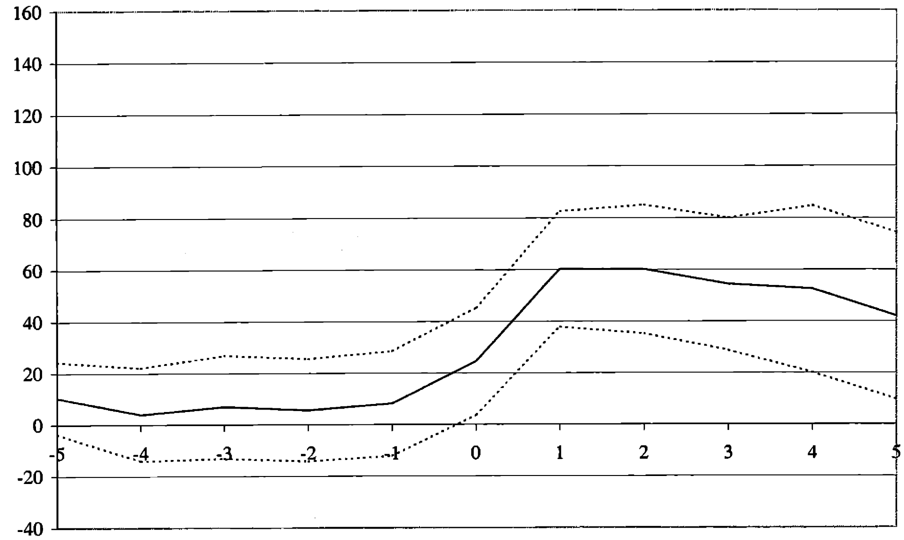

Figure 1-3: OLS Retail Employment Results (WMlist)

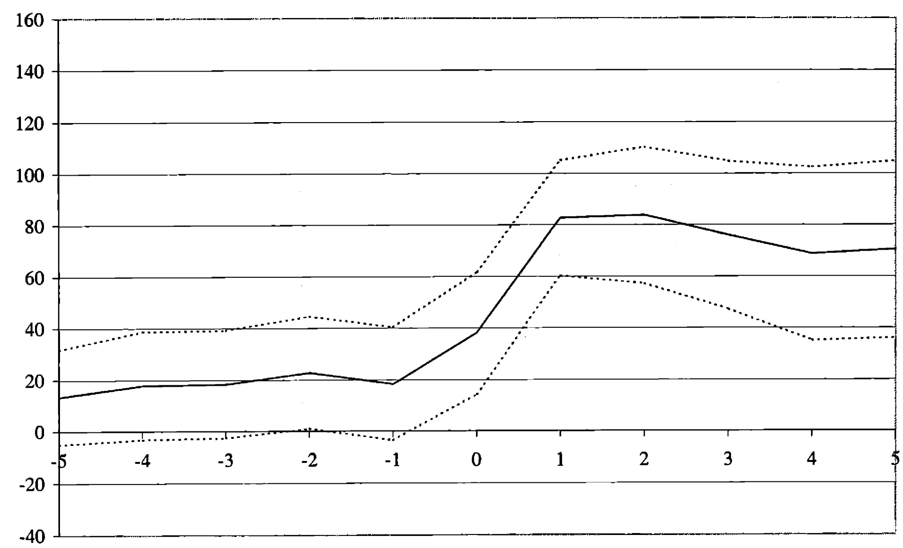

Figure 1-4: OLS Retail Employment Results (WMnum)

part-time jobs with full-time jobs. CBP employment figures do not control for hours worked, so full-time and part-time employees are weighted equally.

It is difficult to compare hours of work for the typical Wal-Mart employee with hours of work at other retailers, because very little is known about employment conditions at Wal-Mart. A reasonable prior is that Wal-Mart employees work *fewer*, not more, hours than other retail workers, based on the finding by Bertrand and Kramarz (forthcoming) that as more entry is allowed into the French retail industry, part-time employment increases relative to all retail employment. Wal-Mart claims that 70% of Wal-Mart employees work 28 hours a week or more (Wal-Mart 2001a). This figure appears to be within the norm for workers in the discount retail industry.[17] It is also in keeping with the rest of the retail industry: the 30th percentile of hours worked by retail employees, obtained from the March Current Population Survey (CPS) for 1978-1999, is 28 hours across employer size, state, and year.

As a specification check, Figure 1-6 shows IV results from a regression that allows year fixed effects to vary not by urbanization but by Census region. This second-stage regression is

$$\frac{\text{retail}_{jt}}{\text{pop}_{jt}} = \alpha + \sum_k \sum_t \lambda_{tk}\,\text{region}_{jk}\,\text{year}_t + \sum_j \psi_j \text{county}_j + \theta(L)\,\frac{\widehat{\text{WMlist}}_{jt}}{\text{pop}_{jt}} + \text{u}_{jt}. \quad (1.4)$$

(with appropriate modification to the first stage regression). The results are extremely similar to those results presented in Figure 1-5, where year fixed effects are allowed to differ by 1960 urbanization status. The instantaneous effect of entry is estimated at 100 jobs, with a decline of 20-50 jobs in the years immediately following entry.[18,19]

As noted in Section 1.4.3, if the timing of entry were endogenous, we would expect to see deviations from the county's long-run level of per-capita retail employment, relative to other counties, prior to entry. No such effect is evident in the leading coefficients. Given that construction alone can take several months (Murzell 1993), it is unlikely that Wal-Mart could

[17]See `http://www.pbs.org/storewars/stores3.html`.

[18]Ideally, we would like to control for state*year fixed effects or state*year*urban status fixed effects. Unfortunately, this is computationally infeasible given the simultaneous inclusion of county fixed effects in all models.

[19]The number of retail workers in Benton County, Arkansas, where Wal-Mart headquarters are located, has increased by 5-6 workers for each new Wal-Mart store over the last 20 years.

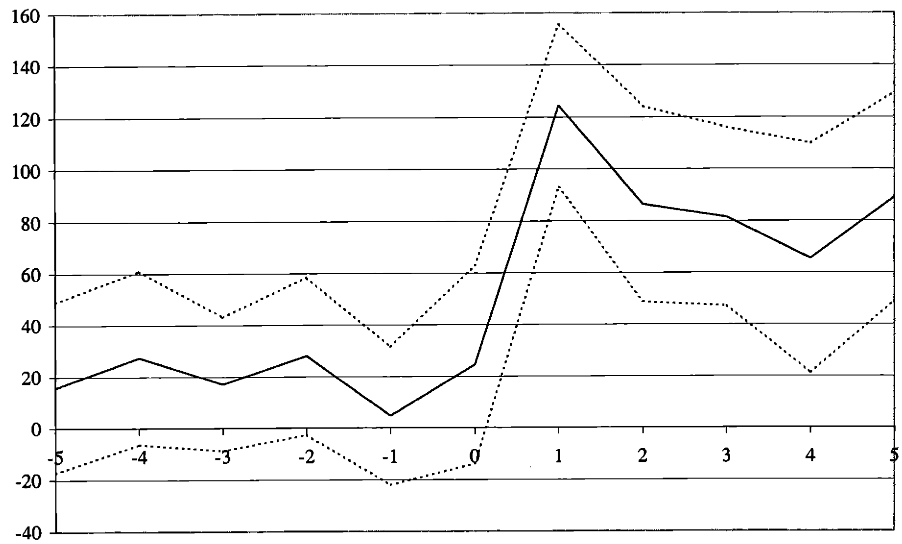

Figure 1-5: IV Retail Employment Results

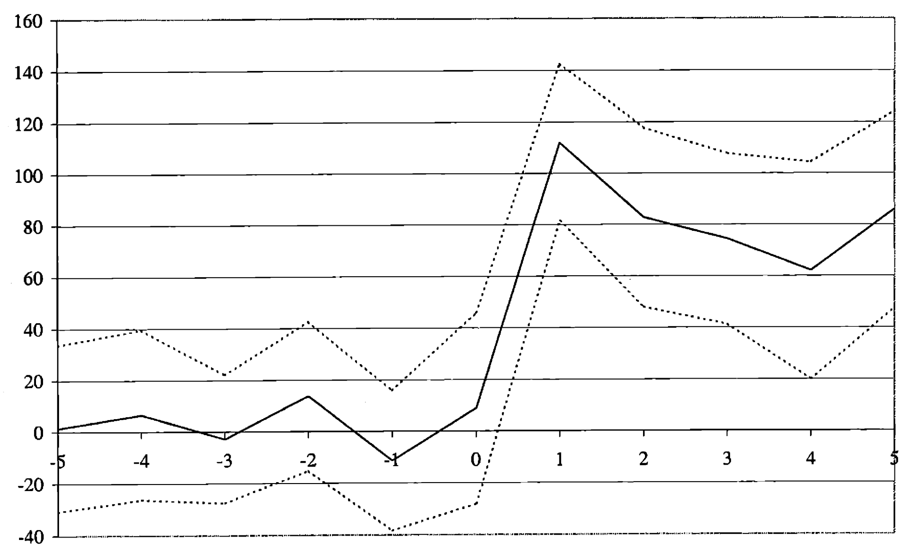

Figure 1-6: IV Retail Employment Results, Region*Year Fixed Effects

introduce a store – with its requisite rezoning, planning, construction, and set-up – in less than a year.

All regression results reported for the remainder of the paper will be IV results with WMnum_{jt} instrumenting for WMlist_{jt}, and year fixed effects allowed to differ across urbanization categories.

1.5.2 Distribution of Retailer Size

Because Wal-Mart competes with retailers across categories – not only with general-merchandise stores, but with apparel stores, drug stores, etc. – it is interesting to look at changes in the overall distribution of retailer size in the years following Wal-Mart entry. We expect a decline in the number of competing retailers following Wal-Mart entry, with less-efficient (and disproportionately small) retailers most likely to exit. As noted by Stone (1987 and elsewhere), some retailers selling complementary products may be positively affected by Wal-Mart entry. In addition, retailers located near Wal-Mart may benefit from the externality of increased customer traffic, if Wal-Mart behaves like an anchor store in a traditional mall (see Pashigian and Gould 1998, Gould, Pashigian, and Prendergast 2002).

To capture the effect of Wal-Mart on the number of retail firms in each size category, I run instrumental-variables regressions with second stage

$$\frac{\text{estab}_{jt}}{\text{pop}_{jt}} = \alpha + \sum_k \sum_t \delta_{tk} \text{urban}_k \text{year}_t + \sum_j \psi_j \text{county}_j + \theta(L) \frac{\widehat{\text{WMlist}}_{jt}}{\text{pop}_{jt}} + \text{u}_{jt} \quad (1.5)$$

where estab_{jt} are, respectively, the number of small retail establishments (under 20 employees) in county j at year t; the number of medium-sized establishments (20-99 employees); and the number of large establishments (100+ employees). IV results are presented below.

Figure 1-7 shows the effect of Wal-Mart on the number of small establishments, defined as having fewer than 20 employees. There is a significant decline of 3 small retail establishments in the years after Wal-Mart entry (p-value 0.0009).

Figure 1-8 shows a decline also in the number of medium-sized establishments (with 20-99 employees) following Wal-Mart entry. Note that the scaling is not as in Figure 1-7, because there is less fluctuation in the number of medium-sized establishments. There is a slight, marginally significant, decline in the number of medium-sized establishments (p-value 0.0337).

Finally, Figure 1-9 shows the effect of Wal-Mart on the number of large establishments (with 100 or more employees). Note that the estimated coefficients mirror those on retail employment shown in Figure 1-5. The increase in the number of large retail establishments, of approximately 0.7, reinforces the interpretation that Wal-Mart's entry coincides with exit or contraction of other large retailers.[20] Additional firms exit, or shrink in size below 100 employees, in the years following Wal-Mart entry.

1.5.3 Retail Employment in Neighboring Counties

Shopping in neighboring counties is generally an imperfect substitute for shopping in one's county of residence, since additional travel is involved. Nevertheless, due to their low prices and large selection, Wal-Mart stores in rural areas typically draw customers from a wide radius which may include several neighboring counties. In a series of studies, Kenneth Stone argues that much of the negative effect of Wal-Mart on retail employment occurs not in the communities in which Wal-Mart located, but in nearby communities (see, e.g., Stone 1997).

At the same time, competition in the labor market may drive workers from retail establishments in their own counties to neighboring counties. (County Business Patterns data attribute jobs to the county in which the employer is located, rather than the workers' counties of residence.) Both effects are expected to work in the same direction, so the expected effect on neighboring counties' retail employment is unambiguously negative.

I define counties as "neighbors" if the distance between their geographic centers is under 10 miles.[21] Formally, let $J = \{\text{neighbor}\,(j)\}$ be the set of county j's neighbors, and define for any variable X,

$$\mathrm{X}_{Jt} = \sum_{k \in J} \mathrm{X}_{kt}. \tag{1.6}$$

[20] In some cases, Wal-Mart acquired a large number of stores from a competitor; in those cases Wal-Mart entry was not associated with a net increase in the number of large retail establishments. Examples include the 1977 purchase by Wal-Mart of 16 Mohr Value Discount Department Stores in Missouri and Illinois, and the 1981 purchase of 106 stores in nine states from Kuhn's-Big K Stores Corp.

[21] Other definitions of "neighbors" were also tried as robustness checks; the results were not sensitive to the exact definition.

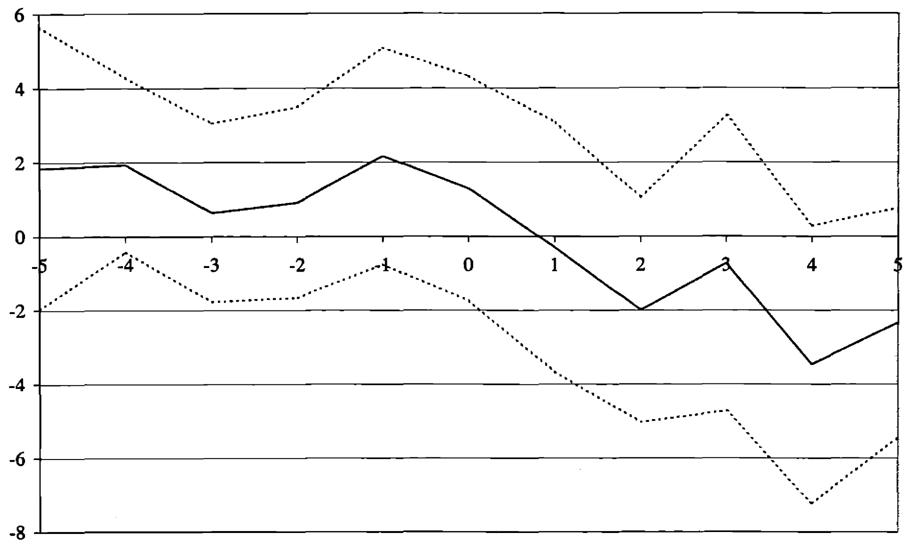

Figure 1-7: Small Retail Establishments (1-19 Employees)

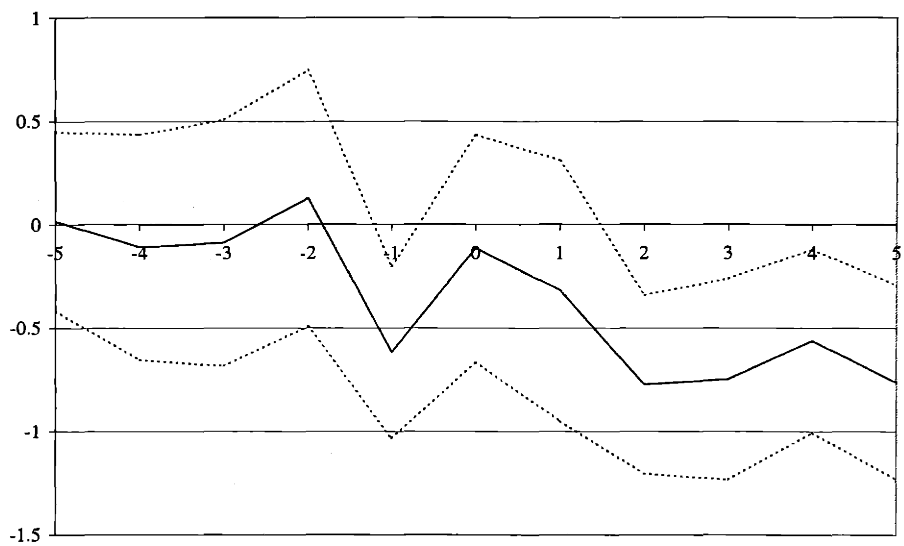

Figure 1-8: Medium-Sized Retail Establishments (20-99 Employees)

To estimate the effect of Wal-Mart entry in county j on retail employment in the surrounding area J, I estimate an IV regression with second stage

$$\frac{\text{retail}_{Jt}}{\text{pop}_{Jt}} = \alpha + \sum_t \delta_t \text{year}_t + \sum_j \psi_j \text{county}_j + \phi \frac{\sum_{s\leq t} \widehat{\text{WMlist}}_{Js}}{\text{pop}_{Jt}} + \theta(L) \frac{\widehat{\text{WMlist}}_{jt}}{\text{pop}_{Jt}} + \text{u}_{jt} \quad (1.7)$$

(to economize on estimated parameters, year fixed effects are assumed constant across urbanization categories). $\frac{\sum_{s\leq t} \widehat{\text{WMlist}}_{Js}}{\text{pop}_{Jt}}$ and $\frac{\widehat{\text{WMlist}}_{jt}}{\text{pop}_{Jt}}$ are predicted values from the appropriate first-stage regressions, with instruments $\frac{\sum_{s\leq t} \widehat{\text{WMnum}}_{Js}}{\text{pop}_{Jt}}$ and appropriate leads and lags of $\frac{\widehat{\text{WMnum}}_{jt}}{\text{pop}_{Jt}}$. The variable $\frac{\sum_{s\leq t} \widehat{\text{WMnum}}_{Js}}{\text{pop}_{Jt}}$ is the number of existing stores, per-capita, in counties J at year t (in other words, the cumulative number of new stores opened no later than year t). It is included in the regression to avoid confounding the effect of Wal-Mart entry in county j with the effect of entry in county j's neighbors, and is assumed to have a once-and-for-all effect on employment in those counties. Note that as the estimated effect is on retail employment per capita in neighboring counties, the appropriate normalization of the number of Wal-Mart stores also uses neighboring counties' population.

IV results are shown in Figure 1-10. No significant effect of Wal-Mart entry on retail employment in neighboring counties can be detected, although the mean of the pre-entry coefficients is roughly 50 jobs above the post-entry mean. To interpret the coefficients, note that the regression includes retail employment in *all* neighboring counties on the LHS. The average county in the sample has about 5 neighbors, so the annual fluctuations shown are on the order of 10-20 jobs per neighboring county.[22]

As an alternative specification with more power, I estimate a simple once-and-for-all IV

[22]The estimated coefficient ϕ is approximately 47 with 95% confidence-interval $[7, 86]$. This number is in line with the estimated five-year effects in the own-county IV regressions. When own-county effects are constained to be once-and-for-all, with second-stage regression

$$\frac{\text{emp}_{jt}}{\text{pop}_{jt}} = \alpha + \sum_t \delta_t \text{year}_t + \sum_j \psi_j \text{county}_j + \phi \frac{\widehat{\text{WMlist}}_{jt}}{\text{pop}_{jt}} + \text{u}_{jt}$$

the estimated coefficient $\widehat{\phi}$ is approximately 62, with confidence interval $[11, 113]$.

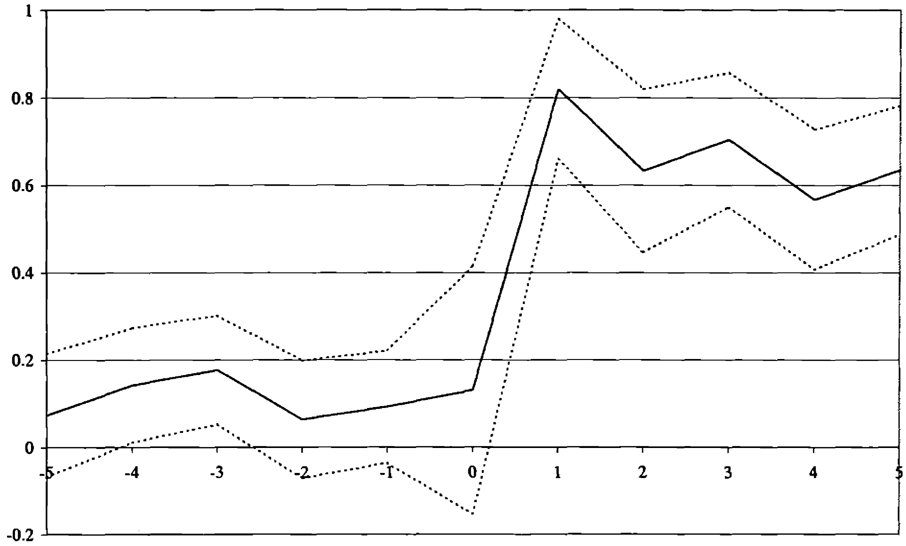

Figure 1-9: Large Retail Establishments (100+ Employees)

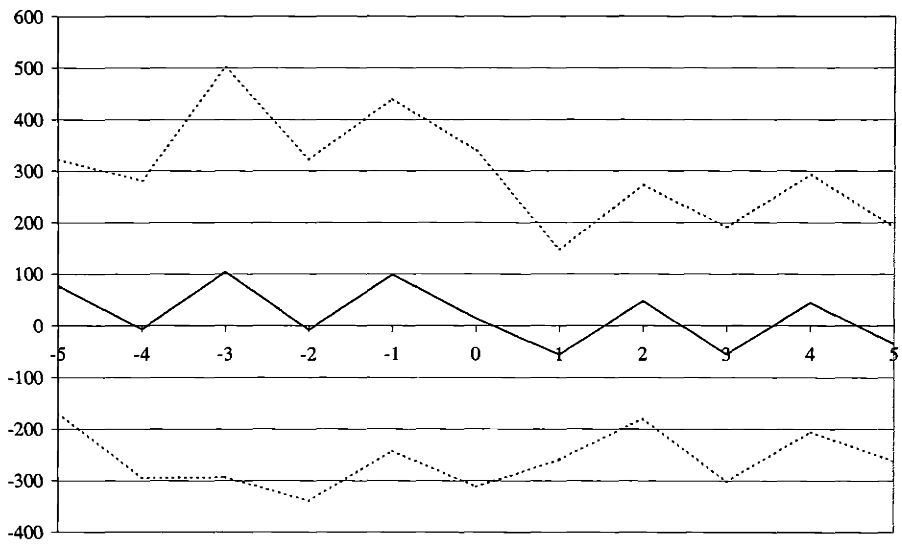

Figure 1-10: Retail Employment in Neighboring Counties

regression with second-stage regression

$$\frac{\text{retail}_{Jt}}{\text{pop}_{Jt}} = \alpha + \sum_t \delta_t \text{year}_t + \sum_j \psi_j \text{county}_j + \phi \frac{\sum_{s \le t} \widehat{\text{WMlist}}_{Js}}{\text{pop}_{Jt}} + \theta \frac{\sum_{s \le t} \widehat{\text{WMlist}}_{js}}{\text{pop}_{Jt}} + \text{u}_{jt} \quad (1.8)$$

where $\frac{\sum_{s \le t} \widehat{\text{WMlist}}_{js}}{\text{pop}_{Jt}}$ is the (predicted) number of existing Wal-Mart stores county j per J capita. The estimated coefficient θ is -29.119 and is significant at the 1% level. In other words, while Wal-Mart entry into county j permanently increases retail employment in county j by approximately 50 jobs, retail employment in neighboring counties decreases by approximately 30 jobs.

1.5.4 Other Sectors

Wholesale

Because Wal-Mart is vertically integrated, Wal-Mart entry is unlikely to complement wholesale employment even though it is associated with an increase in measured retail employment. Moreover, anecdotal evidence suggests that some competing retailers find that Wal-Mart's prices are better than their wholesalers. *Inc.* magazine, in an article about the effect of Wal-Mart on small businesses, interviewed a retailer who noted that his local Sam's Club, a membership-only retail chain owned by Wal-Mart, carried some items at a lower price than his distributor (Welles 1993). Thus, though a retail store, Wal-Mart may be a substitute for wholesalers. In addition to this direct competition from Wal-Mart, wholesalers may also be affected indirectly as small retailers, who traditionally buy from regional wholesalers, shut down (Shills 1997).

The estimated effect of Wal-Mart entry on wholesale employment is shown in Figure 1-11. The observed decline of 25 wholesale jobs following Wal-Mart entry is statistically significant (p-value 0.0334).[23]

[23] In long-run – six or more years after entry – a further 10 jobs are lost in the wholesale sector. The long-run effect is therefore a loss of 35 wholesale jobs (p-value 0.0029).

Restaurants

Restaurant employment is used as a control for retail employment, because the two are likely to be highly correlated, but restaurant employment is not expected to be substantially affected by Wal-Mart entry. There are several caveats to this claim: retail and restaurant employees may be drawn from the same labor market (though restaurant employees are on average younger and less-educated than retail workers); and there is some anecdotal evidence to suggest that restaurants, at least fast-food restaurants (which cannot be separated from other eating establishments in County Business Patterns data), may complement shopping at Wal-Mart.

There is no perceptible impact of Wal-Mart on restaurant employment, as Figure 1-12 shows: neither a discontinuity as with retail employment (see Figure 1-5 above), nor a change in the pattern of growth as with wholesale employment (Figure 1-11). The observed trend in restaurant employment is most likely due to other factors not captured by the regression, and suggests that the instrumental-variables specification addresses some, but not all, concerns about endogeneity in the timing of entry.[24]

Total Employment

As noted above, the typical Wal-Mart store has 150-350 employees, less than 2% of total employment in the average county at the time of the Wal-Mart entry. The chances of finding a statistically-significant effect on total county employment are therefore slim, and in fact, Figure 1-13 shows the estimated effect is statistically zero. Neither the five-year effect nor the long-run effect is statistically different from the coefficient in the entry year. The increasing trend in the years before Wal-Mart entry, however, again suggests that some endogeneity in the timing of entry remains.

1.6 Conclusion

The effect of productivity increases on employment has been analyzed at both the macroeconomic level and the plant level, but the intermediate level has been neglected by researchers. This level of analysis is important, because much of the productivity growth we have observed

[24]The long-run coefficient (not depicted) is not significantly different from the coefficient for year 0.

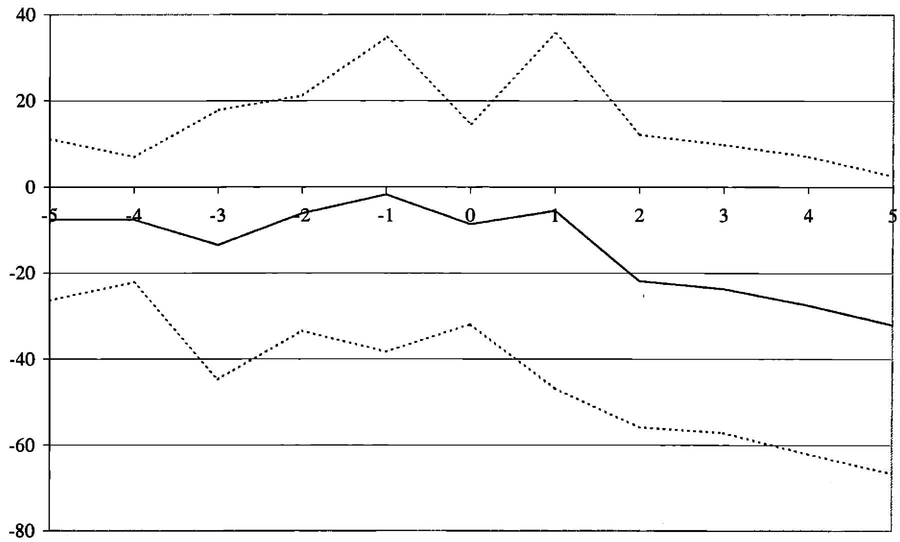

Figure 1-11: Wholesale Employment

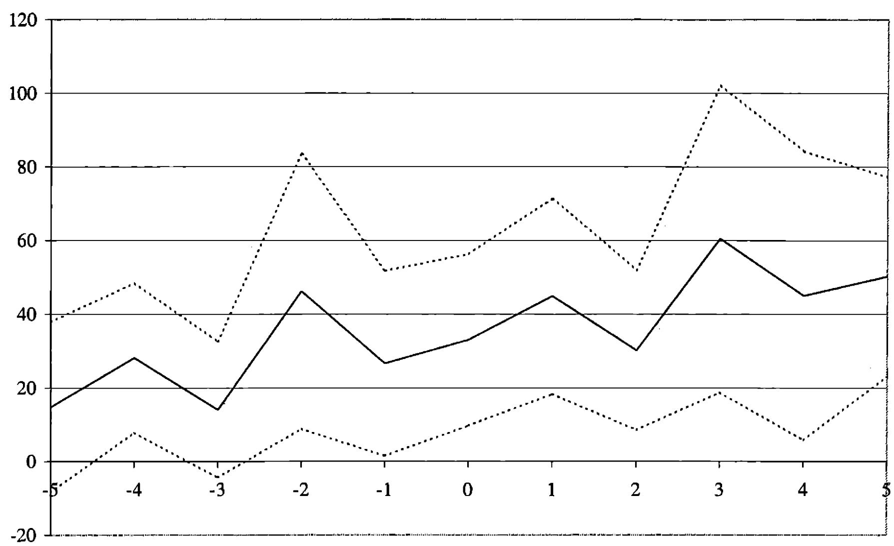

Figure 1-12: Restaurant Employment

in the past decade was associated with entry, rather than technology adoption by existing firms and establishments. This paper takes a first stab at estimating the effect of entry of a more productive firm on sector-level employment and the reallocation of jobs across firms.

The experiment is a clean one, because I am able to identify the date of entry precisely, using an instrumental-variables specification. The effect I estimate is a flexible reduced-form effect, allowing both Wal-Mart and other firms in the county of entry as well as in surrounding counties to adjust to the shock over a period of several years. Finally, because I use a large panel of nearly 1800 counties over 23 years, and because Wal-Mart entry is a "large" shock relative to the size of the local retail market in most counties – median retail employment in 1990 was only 850, while the average Wal-Mart store had approximately 200 employees – the effect can be estimated with relative precision.

I find an increase of 100 retail jobs in the county at entry; half of that increase remains five years after entry. This effect is substantially mitigated when neighboring counties are also considered, where there is a decline of approximately 30 retail jobs. There is also a negative effect on county-level wholesale employment. Combined, these negative effects are large enough to fully offset the gains to retail employment in the entered county.

In closing, it should be emphasized that this paper does not attempt to answer the question whether entry of Wal-Mart has a positive or negative net impact on a local economy. The answer to that question depends on many other factors, which are beyond the scope of this paper; these include concerns about market concentration, income effects, distribution of rents, and more.

Appendices to Chapter 1

Appendix 1.A: Theoretical Framework

This appendix presents a simple model to highlight the ambiguity of the effect of a positive technology shock on sectoral employment. I focus on the retail sector, abstracting away from other sectors in the economy. This approach is legitimate assuming the retail sector is sufficiently small and insulated, so that wage and price changes in the retail sector do not have widespread consequences for prices and production levels in other sectors (and by implication, on aggregate price and production levels).

Consider a perfectly competitive retail sector with an upward-sloping labor supply function

$$L = \eta w$$

where $\eta > 0$. The upward-sloping labor-supply function should be viewed as a reduced-form representation; it could be due, most plausibly, to workers' heterogeneous outside options in other sectors.

Two intermediate retail services are produced competitively with a single input (labor) and CRS production functions

$$\begin{aligned} R_1 &= AL_1 \\ R_2 &= L_2 \end{aligned}$$

where $A \geq 0$ parametrizes firm 1's technology.[25]

A final retail service, with price p, is produced competitively from the intermediate services with CES production function

$$R = \left[\gamma\left(R_1\right)^{\rho} + \left(1-\gamma\right)\left(R_2\right)^{\rho}\right]^{\frac{1}{\rho}} \tag{1.9}$$

[25] A more complicated model with multiple inputs would be able to capture an added dimension of uncertainty due to the elasticity of substitution between labor and capital, which may vary between the old and new technologies.

where $\gamma \in (0,1)$ is the relative importance of R_1 in the final good production, $\frac{1}{1-\rho}$ is the elasticity of substitution in demand for retail services, and $\rho \in (-\infty, 1)$. The two intermediate services are gross substitutes if $\rho > 0$, and they are gross complements if $\rho < 0$. Demand for R is given by the invertible function $R(p)$.

Total retail employment is just the sum of employment in the two intermediate sectors:

$$L_1 + L_2 = L.$$

First-order conditions in the final good sector are

$$\begin{aligned} p_1 &= \gamma p R^{1-\rho} (R_1)^{\rho-1} = \gamma p R^{1-\rho} (AL_1)^{\rho-1} \\ p_2 &= (1-\gamma) p R^{1-\rho} (R_2)^{\rho-1} = (1-\gamma) p R^{1-\rho} (L_2)^{\rho-1} \end{aligned} \tag{1.10}$$

and in the two intermediate sectors

$$\begin{aligned} p_1 &= \frac{w}{A} \\ p_2 &= w. \end{aligned} \tag{1.11}$$

Equating $Ap_1 = p_2$ in the final good FOCs yields

$$L_2 = \left(\frac{\gamma A^\rho}{1-\gamma}\right)^{\frac{1}{\rho-1}} L_1.$$

From the labor supply function,

$$L_1 + L_2 = \eta w = \eta (1-\gamma) p \left[\gamma A^\rho (L_1)^\rho + (1-\gamma)(L_2)^\rho\right]^{\frac{1-\rho}{\rho}} (L_2)^{\rho-1}$$

where the second equality comes from recognizing that $w = p_2$ (from Equation (1.11)) and evaluating p_2, from Equation (1.10), using Equation (1.9).

Solving simultaneously yields employment equations

$$L_1 = \eta p \left(\gamma A^{\rho}\right)^{\frac{1}{1-\rho}} \left[\left(\gamma A^{\rho}\right)^{\frac{1}{1-\rho}} + (1-\gamma)^{\frac{1}{1-\rho}}\right]^{\frac{1-2\rho}{\rho}} \tag{1.12}$$

$$L_2 = \eta p (1-\gamma)^{\frac{1}{1-\rho}} \left[\left(\gamma A^{\rho}\right)^{\frac{1}{1-\rho}} + (1-\gamma)^{\frac{1}{1-\rho}}\right]^{\frac{1-2\rho}{\rho}} \tag{1.13}$$

$$L = \eta p \left[\left(\gamma A^{\rho}\right)^{\frac{1}{1-\rho}} + (1-\gamma)^{\frac{1}{1-\rho}}\right]^{\frac{1-\rho}{\rho}} \tag{1.14}$$

and price equations

$$w = p_2 = p \left[\left(\gamma A^{\rho}\right)^{\frac{1}{1-\rho}} + (1-\gamma)^{\frac{1}{1-\rho}}\right]^{\frac{1-\rho}{\rho}} \tag{1.15}$$

$$p_1 = p \left[\gamma^{\frac{1}{1-\rho}} + \left(\frac{1-\gamma}{A^{\rho}}\right)^{\frac{1}{1-\rho}}\right]^{\frac{1-\rho}{\rho}} \tag{1.16}$$

where p is defined implicitly by the function

$$p = R^{-1}\left(\eta p \left[\left(\gamma A^{\rho}\right)^{\frac{1}{1-\rho}} + (1-\gamma)^{\frac{1}{1-\rho}}\right]^{\frac{2-2\rho}{\rho}}\right). \tag{1.17}$$

Comparative statics cannot be analyzed without putting some structure on the demand function $R(p)$. Assume it is given by

$$R(p) = e^{1-\alpha \ln(p)}$$

with constant elasticity of demand $|\alpha| > 0$. Solving Equation (1.17), we get

$$p = \exp\left[\frac{1 - \ln\left(\eta \left[\left(\gamma A^{\rho}\right)^{\frac{1}{1-\rho}} + (1-\gamma)^{\frac{1}{1-\rho}}\right]^{\frac{2-2\rho}{\rho}}\right)}{\alpha + 1}\right].$$

The main result of this model highlights the ambiguity of the sign of the effect of increased productivity in one firm on overall sectoral employment.

Result 1 *If demand for retail services is inelastic ($\alpha < 1$), total retail employment decreases with A. If demand for retail services is elastic ($\alpha > 1$), total retail employment increases with A.*

Proof. Follows from differentiating Equation (1.14):

$$\frac{\partial L}{\partial A} = \frac{\eta}{A} \exp \left[\frac{1 - \ln \left(\eta \left[(\gamma A^{\rho})^{\frac{1}{1-\rho}} + (1-\gamma)^{\frac{1}{1-\rho}} \right]^{\frac{2-2\rho}{\rho}} \right)}{\alpha + 1} \right] \bullet$$

$$\left[(\gamma A^{\rho})^{\frac{1}{1-\rho}} + (1-\gamma)^{\frac{1}{1-\rho}} \right]^{\frac{1-2\rho}{\rho}} (\gamma A^{\rho})^{\frac{1}{1-\rho}} \left(\frac{\alpha - 1}{\alpha + 1} \right)$$

and noting that all but the last term are strictly positive. ∎

To see the intuition for this result, note that the productivity increase in the retail sector has two direct effects. Given a fixed quantity demanded, fewer workers are needed to supply it. At the same time, however, lower prices will increase quantity demanded. Which effect dominates depends on the price elasticity of demand. Note that this result does not depend on the particular functional forms used here (for final-good production and demand), but is much more general.

Appendix 1.B: Data and Empirical Issues

Wal-Mart Data

Table 1.2 shows the sources from which store opening dates, used in the construction of the variable WMlist_{jt}, were drawn. Chain Store Guides' *Directories of Discount Department Stores* from 1990-1993 are available, but are largely uninformative; the directories do not appear to have been updated in those years. For stores that do not appear in the 1989 directory, but do appear in the 1995 Rand McNally road atlas (i.e., exist in 1994), opening dates are assigned according to the following algorithm. From the annual reports, I obtain the net increase (rarely, a decrease) in the number of Wal-Mart stores in each state each year. Since there are very few store closures, I use the net increase to proxy for the gross increase, i.e., the number of new stores to open each year in each state. For example, in Arizona, 5 new stores opened in 1990, 7 in 1991, and one each in 1992 and 1993. Using the list of stores that existed in 1994 but not in 1989, I assign entry dates randomly, in proportion to their probability of opening in each year. Therefore each store that opened in Arizona during this period has a probability $\frac{5}{14}$ of being assigned to 1990, probability $\frac{1}{2}$ of being assigned to 1991, and probability of $\frac{1}{14}$ each of being assigned to 1992 and 1993. In all, 680 stores' opening dates are assigned in this way, as follows: 203 in 1990, 145 in 1991, 181 in 1992, and 151 in 1993.[26]

Table 1.2: Directory Sources for Wal-Mart Opening Dates

Years	Source
1962-1969	Vance and Scott (1994)
1970-1978	Wal-Mart Annual Reports
1979-1982	Directory of Discount Department Stores
1983-1986	Directory of Discount Stores
1987-1989	Directory of Discount Department Stores
1990-1993	See text
1994-1997	Rand McNally Road Atlas

Table 1.3 shows the assignment of store opening dates by store numbers. Store openings in county j in year t implied by these assigned entry dates are aggregated to the county-year level and form the variable WMnum_{jt}. The accuracy of this method depends critically on Wal-Mart assigning store numbers in a roughly sequential order, and not reassigning numbers in the event of store closure. Only 40 stores closed over the entire period 1964-1999, so the latter condition appears to be satisfied; this also implies that reassignment of store numbers, if it takes place at all, cannot be common.[27]

Figure 1-14 shows the distribution of the difference between the two measures of opening dates, at the store level. Over 40% of stores are assigned the same opening year with both measures, approximately 80% of stores are assigned two opening dates within one year of one another, and 90% are within two years.

County Merges and Splits

The following is a complete list of county merges and splits in the contiguous 48 states since 1960.[28]

[26] Entry dates assigned in this way clearly suffer from measurement error, but they are unbiased. This method is therefore preferred to the naïve alternative of assigning all stores that opened in those years the date they first appear in the data (generally, 1993); entry dates assigned this way would be biased as well as measured with error.

[27] Relocation of stores within a community, which is much more common than store closure, does not pose a problem.

[28] For more information about county definitions and changes see

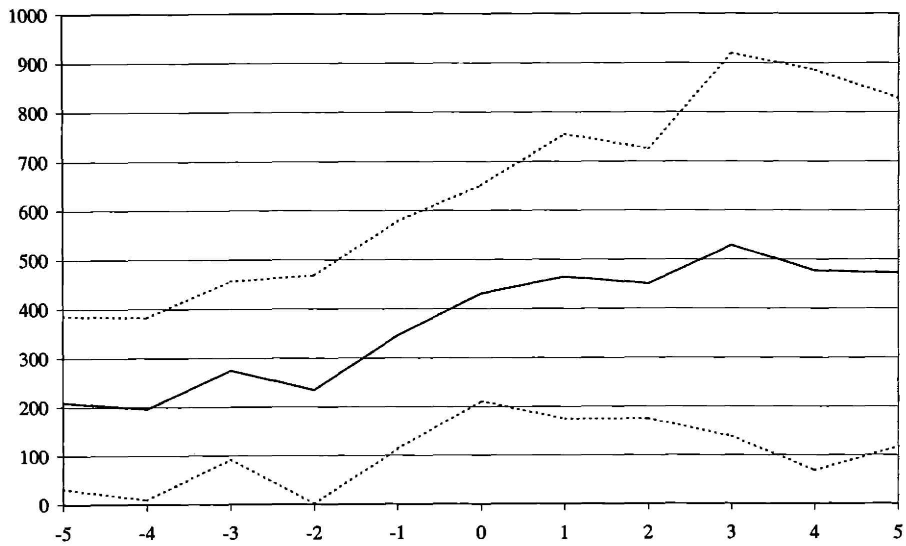

Figure 1-13: Total Employment

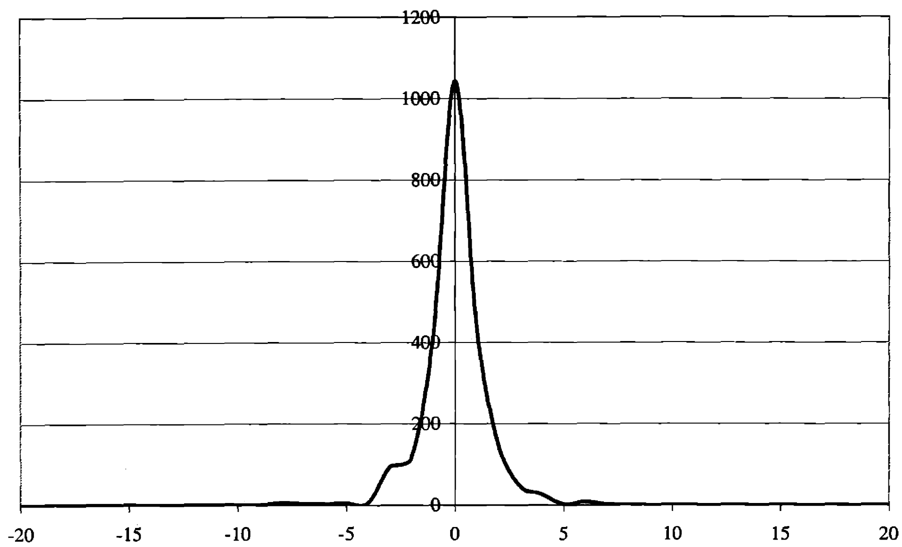

Figure 1-14: Differences between the Two Measures of Opening Dates (Years)

Table 1.3: Store Opening Year Assignment by Store Number

Store Numbers	Assigned Year	Store Numbers	Assigned Year
1-2	1964	330-489	1981
3	1965	490-549	1982
4-5	1966	550-640	1983
6-7	1967	641-743	1984
8-12	1968	744-857	1985
13-17	1969	858-977	1986
18-37	1970	978-1111	1987
38-51	1971	1112-1256	1988
52-62	1972	1257-1399	1989
63-77	1973	1400-1569	1990
78-103	1974	1570-1701	1991
104-124	1975	1702-1874	1992
125-152	1976	1875-2013	1993
153-193	1977	2014-2123	1994
194-225	1978	2124-2231	1995
226-274	1979	2232-2999	1996
275-329	1980	2300+	1997

Arizona: Yuma County split into two counties, La Paz and Yuma, in 1987.

Florida: Dade County changed its name to Miami-Dade County in 1997.

Georgia: Muscogee County became Columbus-Muscogee county, in 1971; reflected in Census data beginning 1974.

Nevada: Carson City and Ormsby County were consolidated into one community in 1969.

New Mexico: Valencia County split in two in 1983 to form Cibola County and Valencia County.

Virginia has seen the largest number of changes, many of them involving the creation of "independent cities" that are not part of any county. Bedford city split from Bedford County in 1968; Emporia city split from Greensville County in 1967; Lexington city split from Rockbridge County in 1966; Manassas city and Manassas Park city split from Prince William County in 1975; Nansemond County became Nansemond city in 1972 and then was annexed to Suffolk city in 1974; Poquoson city split from York County in 1975; Salem city split from Roanoke County in 1968; South Boston city merged with Halifax County in 1995.

http://www.census.gov/population/cencounts/00-90doc.txt.

Wisconsin: Menominee County was formed from parts of Shawano and Oconto Counties in 1961.

Wyoming: Since 1970, employment in Yellowstone National Park has been reported with Teton and Park Counties; I have combined these three counties and treat them as one throughout.

Employment Data Accuracy

Occasionally, in counties with a small number of employers, data on the total number of employees in a sector, or even in the entire county, is omitted from County Business Patterns to avoid disclosure of the number of employees in individual firms. Only the number of firms in each of eight employment-size categories (1-19, 20-99, ... 50,000-99,999, 100,000+ employees) is given in these cases. In those instances, I assume that the actual number of employees of a firm of size X is a weighted mean of the lower and upper bounds on its employment-size class (with weight $\frac{2}{3}$ on the lower bound and $\frac{1}{3}$ on the upper bound); the exception is the class-size of 100,000+, to which I assign 150% of the lower bound, or 150,000. For example, a firm with 1-19 employees is assigned a value of 7.[29]

Unit Roots

There is a high degree of persistence in county-level employment. To test whether the employment series used admit unit roots, I run a Dickey-Fuller (DF) test on each county series separately, after removing year fixed effects interacted with 1960 urbanization status (urban, suburban, rural). By construction, a 5% rejection rate is to be expected at the 95% confidence level if the series have unit roots. The actual rejection rates vary by series from 6%-14%, and are shown in Table 1.4.

Panel-data unit root tests provide a powerful alternative to county-by-county testing. I apply two such tests, one by Maddala and Wu (1999) and another by Levin and Lin (1993). Maddala and Wu use a variant of a test by Fisher which uses a combination of the p-values

[29] I chose to weight the lower and upper bounds of each interval by $(\frac{2}{3}, \frac{1}{3})$, respectively, rather than $(\frac{1}{2}, \frac{1}{2})$, because counties small enough to elicit concerns about disclosure of information on individual firms in aggregate data seem likely to have a disproportionate number of small employers. The results are robust to this specification.

from the county-by-county DF tests. The Maddala-Wu test statistic is

$$-2\sum_{i=1}^{N}\ln(\pi_i) \sim \chi^2_{(2N)}$$

where π_i is the p-value from the Dickey-Fuller test for county i, and the null hypothesis is that all series share a unit root. The Levin-Lin test is a panel variant of the Dickey-Fuller test with the same null hypothesis.

The panel tests are subject to two caveats. First, the null hypothesis will be violated if even one of the series is stationary. Rejection of the null hypothesis may therefore be interpreted to imply that some series have unit roots and others do not, or, if we believe the series constitute realizations of a single process, rejection implies that the common process is stationary. The tests are therefore meaningless under the assumption that each county employment series may be the realization of a unique process. Second, both tests assume that the observations are independent; this assumption may be violated by spatial correlation.

Table 1.4 reports the test results. The first column shows the fraction of counties for which county-by-county Dickey-Fuller tests rejected the presence of unit roots at 95% significance. The rejection rates of 6%-14% for these series are higher than the expected 5% under the null hypothesis of unit roots. The second and third columns report p-values from Maddala-Wu and Levin-Lin tests, respectively.

Table 1.4: Unit Root Tests

Employment	Dickey-Fuller % Rejected	Maddala-Wu p-Value	Levin-Lin p-Value
Total	5.74	0.0016	0.000
Retail	8.50	0.0000	0.000
Wholesale	13.67	0.0000	0.000
Restaurant	13.67	0.0000	0.000

Although the tests reject unit roots for all employment series, the caveats mentioned above render these rejections less than perfectly informative. The decision to use levels or first differences is therefore somewhat arbitrary. I report the results in levels because they are somewhat easier to interpret.

Appendix 1.C: Sales per Worker

In this section I attempt to quantify the productivity differences between Wal-Mart and other retailers, subject to some important caveats. I use a common proxy for productivity in the retail sector, sales per worker, available for selected years at the county level from the Census of Retail Trade (CRT).

The CRT is conducted every five years, in years ending in 2 or 7 (1972, 1977, etc.), and provides county-level sales volume and total retail employment. To compute sales per worker, I use these variables: total sales revenues for all establishments (available for 1972-1992), total sales revenues for establishments with paid employees (available only for 1977-1992), and total number of paid employees (available for 1972-1992).[30] I use these data to compute two measures of sales per worker. The first is the ratio of sales in establishments with paid employees to the number of paid employees. This measure is available for 1977, 1982, 1987, and 1992. The second, noisier, measure is the ratio of total sales (in establishments with *and without* paid employees) to the number of paid employees. The second measure is available back to 1972, but has the disadvantage that sales in establishments with no employees (e.g., only a proprietor) are attributed to employees in other establishments. In practice, sales in establishments without employees are a very small fraction of total sales, so the figures are extremely similar.

I estimate the simple regression equation

$$\frac{\text{sales}_{jt}}{\text{retail}_{jt}} = \alpha + \sum_t \delta_t \text{year}_t + \sum_j \psi_j \text{county}_j + \theta \sum_{s \leq t} \text{WalMart}_{js} + \text{u}_{jt} \tag{1.18}$$

where sales_{jt} is sales revenue (in real 1982-1984 dollars) of all retail establishments in county j in year t, retail_{jt} is retail employment in county j in year t, $\sum_{s \leq t} \text{WalMart}_{js}$ is (as in Section 1.5.3) the number of Wal-Mart stores in existence in county j in year t, and the other variables are as defined above. t runs from 1972 to 1992. The error term is clustered at the county level to allow for arbitrary autocorrelation at the county level.

Table 1.5 presents the OLS and IV results from these regressions. The OLS results are shown separately for $\sum_{s \leq t} \text{WMlist}_{js}$ and $\sum_{s \leq t} \text{WMnum}_{js}$. The IV results use $\sum_{s \leq t} \text{WMnum}_{js}$ to in-

[30] County-level data from the 1997 CRT have unfortunately not yet been released.

strument for $\sum_{s\leq t} \text{WMlist}_{js}$. "All sales" indicates that sales from all establishments are used in the computation of the LHS variable (allowing inclusion of the 1972 data); "Sales by employees" indicates that sales data used refer only to sales in establishments with employees. The IV estimates suggest that every Wal-Mart store increases sales per worker in the county by \$550-\$750 per year (in constant 1982-1984 dollars). Mean and median sales per worker for the period studied are about \$75,000, so this increase, while highly significant, is smaller than 1%.

Table 1.5: Estimated Effect of WalMart on Sales per Worker

	OLS		
	WMlist	WMnum	IV
All sales[a]	\$602.9	\$720.8	\$765.2
	(160.8)	(161.6)	(156.8)
Sales by employees[b]	442.3	579.0	578.0
	(159.1)	(159.9)	(143.9)

[a] Sales in all establishments per paid employee

[b] Sales in establishments with paid employees per paid employee

These figures should be interpreted with caution. While they show that Wal-Mart entry coincides with an increase in county-level sales per employee, this coincidence is neither necessary nor sufficient to prove that Wal-Mart is more productive than its competitors. To see why it is not sufficient, note that Wal-Mart entry may increase sales per employee by substituting capital for labor; measures of capital are unfortunately not available for the retail sector. At the same time, Wal-Mart prices tend to be lower than its competitors', so sales figures may decline even as productivity increases, implying that the coincidence is not even necessary.

References

Baily, Martin Neil and Robert M. Solow (2001). "International Productivity Comparisons Built from the Firm Level." *Journal of Economic Perspectives* 15.

Banerjee, Anindya (1999). "Panel Data Unit Roots and Cointegration: An Overview." *Oxford Bulletin of Economics and Statistics* 61.

Barnes, Norma Ganim, Allison Connell, Lisa Hermenegildo, and Lucinda Mattson. "Regional Differences in the Economic Impact of Wal-Mart." *Business Horizons* 39.

Bertrand, Marianne, and Francis Kramarz (forthcoming). "Does Entry Regulation Hinder Job Creation? Evidence from the French Retail Industry." *Quarterly Journal of Economics.*

Bertrand, Marianne, Esther Duflo and Sendhil Mullainathan (2001). "How Much Should We Trust Difference-in-Difference Estimates?" MIT mimeograph.

Chain Store Guide (1979-1982, 1987-1993). *Directory of Discount Department Stores.* New York: Business Guides, Inc.

Chain Store Guide (1983-1986). *Directory of Discount Stores.* New York: Business Guides, Inc.

Darby, W. D. (1928). *The Story of the Chain Store: A Study of Chain Store Policies and Methods, Particularly as They Affect the Independent Merchant in the Dry Goods Field, Together With a General Survey of Chain Store Developments.* New York: Dry Goods Economist.

Dobbs, I. M., M. B. Hill, and M. Waterson (1987). "Industrial Structure and the Employment Consequences of Technical Change." *Oxford Economic Papers* 39.

Edgecliffe-Johnson, Andrew (June 19, 1999). "A Friendly Store from Arkansas: Wal-Mart's Devotion to Customer Service and Cost-Cutting Technology is Starting to Reshape Retailing Around the World." *Financial Times.*

Entorf, Horst, Michel Gollac, and Francis Kramarz (1999). "New Technologies, Wages, and Worker Selection." *Journal of Labor Economics* 17.

Feiner, Jeffrey M., et al. (1993). *The Wal-Mart Encyclopedia: A Portfolio Manager's Guide to Growth.* New York: Salomon Brothers.

Feiner, Jeffrey M., et al. (2001). *Wal-Mart Encyclopedia X: Building a Global Brand.* New York: Lehman Brothers.

Foley, Sharon, and Takia Mahmood (1996). "Wal*Mart Stores, Inc." Harvard Business School Case Study 9-794-024.

Franz, Lori, and Edward Robb (1990). "Effect of Wal-Mart Stores on Economic Environment of Rural Communities." Business and Public Administration Research Center mimeograph.

Foster, Lucia, John Haltiwanger, and C. J. Krizan (2001). "The Link between Aggregate and Micro Productivity Growth: Evidence from Retail Trade." Presented at NBER Summer Institute 2001.

Ghemawat, Pankaj (1989). "Wal-Mart Stores' Discount Operations." Harvard Business School Case Study 9-387-018.

Ghemawat, Pankaj and Gregg Friedman (1999). "Wal-Mart in 1999." Harvard Business School Case Study N9-799-118.

Pashigian, B. Peter and Eric D. Gould (1998). "Internalizing Externalities: The Pricing of Space in Shopping Malls." *Journal of Law and Economics* 41.

Gould, Eric D., B. Peter Pashigian and Canice Prendergast (2002). "Contracts, Externalities, and Incentives in Shopping Malls." Mimeograph.

Im, Kyung So, M. Hashem Pesaran, and Yongcheol Shin (1997). "Testing for Unit Roots in Heterogeneous Panels." Mimeograph, University of Cambridge.

Johnson, Bradford C. (2002). "Retail: The Wal-Mart Effect." *The McKinsey Quarterly.*

Kane, Thomas J., Cecilia Elena Rouse, and Douglas Staiger (1999). "Estimating Returns to Schooling When Schooling is Misreported." Princeton Industrial Relations Section Working Paper 419.

Kézdi, Gábor (2001). "Robust Standard Error Estimation in Fixed-Effects Panel Models." University of Michigan mimeograph.

KMart Corporation (1997-1999). *KMart Fact Book.*

Levin, Andrew, and Chien-Fu Lin (1993). "Unit Root Tests in Panel Data: New Results." University of California San Diego Discussion Paper 93-56.

Maddala, G.S., and Shaowen Wu (1999). "A Comparative Study of Unit Root Tests with Panel Data and a New Simple Test." *Oxford Bulletin of Economics and Statistics* 61.

McKinsey Global Institute (2001). "US Productivity Growth 1995-2000: Understanding the contribution of Information Technology relative to other factors." Washington DC: MGI.

Muller, Thomas (1999). "Impact of Anticipated Wal-Mart on the Northern Neck (including Kilmarnock and nearby communities)." Fairfax, Virginia: mimeograph.

Muzeroll, Phillis A. (1993). "Upper Valley Economic Report: Ascutney, Wal-Mart Follow Growth." *Vermont Business Magazine.*

Shils, Edward B. (1997). "The Shils Report: Measuring the Economic and Sociological Impact of the Mega-Retail Discount Chains on Small Enterprise in Urban, Suburban and Rural Communities." `http://www.lawmall.com/rpa/rpashils.htm`

Shnuch, Mary T. (June 3, 1990). "Business Isn't the Same When Wal-Mart Comes to a Small Town." *The Washington Post.*

Stalk, George, Jr., and Thomas M. Hout (1990). *Competing Against Time: How Time-based Competition Is Reshaping Global Markets.* New York: The Free Press (Macmillan).

Stone, Kenneth E. (1991). "Competing with the Mass Merchandisers." *Small Business Forum*, Spring 1991.

Stone, Kenneth E. (1997). "Impact of the Wal-Mart Phenomenon on Rural Communities." *Increasing Understanding of Public Problems and Policies.* Chicago: Farm Foundation.

Stone, Kenneth E. (1989). "The Impact of Wal-Mart Stores on Other Businesses in Iowa." Iowa State University mimeograph.

Urban Land Institute (various years). *Dollars and Cents of Shopping Centers: A Study in Receipts and Expenses in Shopping Center Operations.* Washington, DC: Urban Land Institute.

US Bureau of the Census (1964-1999). *County Business Patterns.* Washington, DC: US Government Printing Office.

Vance, Sandra S., and Roy V. Scott (1994). *Wal-Mart: A History of Sam Walton's Retail Phenomenon.* New York: Twayne Publishers.

Wal-Mart Stores, Inc. (1971-2001). *Annual Report.*

Wal-Mart Stores, Inc. (2001a). *Wal-Mart Associate Handbook.*

Welles, Edward O. (1993). "When Wal-Mart Comes to Town." *Inc. Magazine*, July 1, 1993.

Chapter 2

Education, Job Search and Migration

> "If you don't already have a job in your new location, your number one priority probably will be to find one. Looking for a job long distance can be more fun than looking in your existing location.... No matter how thrilling the prospects, however, you have to be certain you can afford a long-distance search."
> – *Steiner's Complete How-to-Move Handbook*

2.1 Introduction

Between two and three percent of Americans move their residence across state lines every year, many of them for employment-related reasons. Some move to search for work in a new location, though most move to take jobs they have already secured. Not much is known about this process of job search and migration: Why do some people move first, before they have found a job? How do employment outcomes vary by the type of move? How sensitive are movers to local and overall economic conditions? The purpose of this paper is, therefore, to explore the interaction between job search and migration, both theoretically and empirically. I focus on the ways in which education changes workers' incentives, and therefore the migration process.[1]

This much is known: The propensity to migrate decreases with age and increases with

[1] I use the terms "migrant" and "mover" interchangeably throughout the paper. I also use the terms "high-skill", "high-education" and "high-wage" interchangeably.

education (see Greenwood 1975 and 1993). There is also evidence that the unemployed are more likely to migrate than the employed (Schlottmann and Herzog 1981), and that the unemployed are more sensitive than the employed to the overall unemployment rate in their migration decision (DaVanzo 1978; Bartel 1979). Since the incidence of unemployment is higher among less-educated workers, this effect may mitigate the direct positive effect of education on migration. Table 2.1 shows some summary statistics on the differential rates of migration across education categories, computed from the March Current Population Survey (CPS) from 1981-2000.[2] While the overall rates of migration are small, the differences between groups are striking: college educated workers are 82% more likely to migrate in any given year than are high-school dropouts.[3]

In recent years the CPS questionnaire has solicited information about movers' main reason for their move. Approximately half of all migrants over the period 1997-2000 moved for job-related reasons. Of these, 90% moved to take a new job or for a job transfer, and the remaining 10% moved to search for work.[4,5] This fraction varies substantially by education level, as shown in Table 2.2. The probability that a migrant is moving to take a job increases monotonically with her education level; the probability that she moves to look for work or for non-job related reasons decreases monotonically with her education. Strikingly, of the high-school dropouts who moved either to take a job or to look for a job, nearly a third moved to search for a job. Fewer than 3% of college graduates who moved for one of these two reasons (to take a new job or search for a job) moved for the purpose of searching.[6]

The focus of this paper is on the interaction of the migration decision with job-search, and the ways in which this interaction depends on a worker's level of education. I start with a

[2] See Section 2.3 for a description of the data.

[3] Mauro and Spilimbergo (1999) redo the classic analysis of Blanchard and Katz (1992) using Spanish data, but are able to obtain separate estimates by education level of the population. They find that, following an adverse regional employment shock, adjustment for highly-educated workers occurs quickly via out-migration, whereas adjustment for low-education workers is much slower and involves high unemployment and low participation rates for a prolonged period. The implication of these findings is that highly-educated workers migrate away in response to a negative idiosyncratic regional shock, whereas low-education workers do not.

[4] Unfortunately, the *ex-ante* labor-force status (employed, unemployed, not in labor force) of these workers is not known.

[5] Since this question was only asked in the last four years of a long economic expansion, results may not generalize. For example, a larger fraction of moves may be for the purpose of looking for a job in leaner years.

[6] The order of job search and migration has implications for employment outcomes: workers who move to take a job they have already found are up to 13% more likely to be employed the following March than workers who move first and search later.

Table 2.1: Education and Migration Statistics

	Fraction of Population	Propensity to Migrate	Fraction of Migrants[a]
All	100%	2.69%	100%
HS Dropouts	12.52%	2.03%	9.45%
HS Graduates	35.06%	2.16%	28.23%
Some College	25.18%	2.70%	25.27%
College Grads +	27.24%	3.66%	37.06%

[a] May not add to 100% due to rounding

Source: Author's calculations from CPS, 1981-2000

Table 2.2: Main Reason for Migration by Education

Main Reason for Move[a]	All Movers[b]	HS Dropouts	HS Grads	Some College	College Grads +
	Panel A: Full Sample				
New job / job transfer	46.43%	26.76%	32.30%	39.29%	60.46%
Looking for work / lost job	5.47%	12.81%	9.28%	6.79%	1.78%
Other job-related reason[c]	8.13%	7.71%	7.99%	8.84%	7.84%
Non-job related reason[d]	39.97%	52.72%	50.44%	45.09%	29.92%
	Panel B: Men Only				
New job / job transfer	49.70%	31.58%	33.29%	38.92%	66.02%
Looking for work / lost job	6.26%	15.48%	11.21%	8.79%	1.21%
Other job-related reason[c]	8.81%	7.73%	9.41%	9.81%	8.15%
Non-job related reason[d]	35.22%	45.21%	46.10%	42.49%	24.62%

[a] May not add to 100% due to rounding

[b] Includes only movers whose moving status and reason for moving are not imputed

[c] Includes retirement, easier commute, and miscellaneous job-related reasons

[d] Includes family reasons (e.g., move for spouse), health reasons, etc.

Source: Author's calculations from CPS, 1997-2000 simple consumer-choice model in which workers have the choice of searching for a job locally, searching for a job globally (in multiple regions simultaneously), or moving to another region and searching for a job there. I derive conditions under which each strategy dominates the others; these depend on the worker's expected wage (a proxy for her education or skill) as well as on economic conditions, both aggregate and local. I find that, as expected, high-skilled workers are more likely to search globally (and therefore to migrate for job-related reasons) than are low-skilled workers; high-skilled workers are also least likely to engage in labor-market arbitrage in the sense of moving from high-unemployment states to low-unemployment states. I also find that, while migration is expected to be pro-cyclical, the cyclicality of migration should be greatest for workers with intermediate skill (education) level.

I next turn to testing these hypotheses using US data. Using pooled cross-section individual-level data, spanning two decades, from the CPS, I estimate individual migration equations. I find that, as predicted, migration is pro-cyclical in the aggregate and counter-cyclical with respect to state-level conditions: workers are most likely to move out of high-unemployment states when the economy as a whole is strong. When I allow the effect of economic variables to differ across education categories, I find that high-school graduates are more sensitive to aggregate business-cycle conditions than are workers with both higher and lower education levels. Workers with the lowest education level are, however, the most sensitive to arbitrage opportunities in unemployment rates across states.

Existing models of job search assume either that migration must precede search or that search must precede migration. McCall and McCall (1987) assume that migration (between cities) must precede job search. Coulson, Laing, and Wang (2001) model search in a single metropolitan area, and allow agents to search in either the central business district or the suburbs, but not both simultaneously (though they may search – and work – in either market with a commuting cost); they argue that global search will never occur. Spilimbergo and Ubeda (2001) develop a model of migration in which they focus on double matching in the labor market and social setting. They assume that unemployed workers can find a job with certainly upon migration (and with probability less than 1 locally), so there is no job search following a move. In contrast to these papers, I derive the conditions under which an unemployed worker will search locally and globally, and correspondingly conditions under which a worker will move to search for a job or to take a job.

The remainder of the chapter is organized as follows. Section 2.2 presents the model. Section 2.3 describes the data used in the analysis, and Section 2.4 discusses the empirical strategy. Results are presented in Section 2.5. Section 2.6 concludes.[7]

2.2 Model

2.2.1 Setup

Consider an unemployed worker, who, if employed, would produce expected output y and earn an expected wage w.

There is one period. At the beginning of the period, the worker searches for a job. The worker is risk-neutral and maximizes expected income less search costs. She has several alternative technologies for job search. Local search in the worker's region of residence (WLOG Region 1) is costless, but yields a job with lower probability than global search. Global search, in all regions simultaneously, yields a job with higher probability, but at a cost (c); if the job is found in the other region (Region 2), moving costs need also to be incurred. Finally, the worker may choose to move preemptively to Region 2, incurring the moving cost with certainty, but searching only locally once she arrives. A worker who moves – to search for a job or to take a job found in a global search – incurs a cost of moving m.

Let h be the probability that the worker finds a job anywhere; h proxies for global business-cycle conditions. Let ph be the probability that a worker finds a job in Region 1, and let the probability of finding a job in Region 2 be $(1-p)\,h$. Normalize the cost of searching locally to zero.

The choice for a worker in Region 1 is therefore among the following three options. Searching locally is costless and yields a job with probability ph. Searching globally costs c and yields a job with probability h, with an additional cost of moving m with probability $(1-p)\,h$. Moving to Region 2 and searching there involves a moving cost m with certainty, and a probability of (costlessly) finding a job $(1-p)\,h$. The worker's expected utility, conditional on each of these

[7] Throughout the paper, I use the terms "mover" and "migrant" interchangeably. I use the term "job migrant" (or job mover) to refer to a person who moved with a job in hand, and "search mover" to refer to a person who moved in order to search for work.

three actions, are:

$$\begin{aligned} U^L &= phw && (2.1) \\ U^G &= hw - c - (1-p)\,hm && (2.2) \\ U^M &= (1-p)\,hw - m. && (2.3) \end{aligned}$$

where U^L is her expected utility from conducting a local search, U^G is her expected utility from conducting a global search, and U^M is her expected utility from moving to Region 2 and conducting a local search there.

For simplicity, we start by focusing on the case where $h = 1$ (i.e., global search yields a job with certainty, and local search yields a job with probability p). In that case

$$\begin{aligned} U^L &= pw && (2.4) \\ U^G &= w - c - (1-p)\,m && (2.5) \\ U^M &= (1-p)\,w - m. && (2.6) \end{aligned}$$

Given these expected utility functions, workers in Region 1 will choose

$$A = \begin{cases} G & \text{if} \quad w > \overline{w} = \max\left\{\frac{1}{1-p}c + m, \frac{1}{p}c - m\right\} \\ L & \text{if} \quad w < \underline{w} = \min\left\{\frac{1}{1-p}c + m, \max\left\{\frac{1}{1-2p}m, 0\right\}\right\} \\ M & \text{if} \quad w \in [\underline{w}, \overline{w}] \end{cases} \tag{2.7}$$

The probability that a worker moves from Region 1 is therefore

$$P\,(\text{move}) = \begin{cases} 0 & \text{if} \quad w < \underline{w} \\ 1 & \text{if} \quad w \in [\underline{w}, \overline{w}] \\ 1-p & \text{if} \quad w > \overline{w} \end{cases} . \tag{2.8}$$

Figure 2-1 shows the decision space for the worker in (p, w) space if $m = c = 2$ and $h = 1$. For low w, searching locally dominates for sufficiently low p. Searching globally dominates for sufficiently high wages regardless of p. For an intermediate set of wages and sufficiently low probability of finding a job locally, moving to Region 2 to search there dominates.

Extrapolating to a population of workers facing similar problems (perhaps with idiosyncratic values of w, corresponding to ability, and m, corresponding to local attachment or the "psychic cost" of moving), we can make the following predictions. When local conditions are very bad – p is very low – all but the lowest-skilled workers migrate to search for work elsewhere. As p increases, and local conditions improve, high-wage workers turn to global search, which decreases the probability that they will migrate, and low-wage workers turn to local search. At very high levels of p, when local conditions are very favorable, high-wage workers too eventually switch to local search. As p increases, therefore, the probability of migration decreases for two reasons. First, for high wage workers, the nature of the search changes discretely: from certain or possible migration (if M or G dominate) to certain non-migration. Second, in the region where G dominates, the probability of migration decreases as p increases since the probability that a global search will result in a job outside the region declines with p.

Note that the above results do not depend on the one-period specification, but would carry through (with obvious modifications) to a dynamic setting in which jobs may be lost and the decision to move may be revisited. The key assumption driving the results is that the costs of search and moving (c and m, respectively) are fixed, whereas the wage increases with ability. A high-skilled worker has a high opportunity cost associated with unemployment, and is therefore willing to spend more resources – in the form of c and m – to increase the probability of finding a job. In contrast, a worker whose wage, and therefore opportunity cost of unemployment, is low, will not spend as many resources on job search.

2.2.2 Comparative Statics

As m decreases (i.e., the workers becomes less migration-averse), all thresholds shift down, so that both M and G dominate on larger regions. Figure 2-2 shows decision space for the case of $m = h = 1$ and $c = 2$; as $m \to 0$ local search exists only when $p \geq \frac{1}{2}$.

As the cost of global search, c, decreases, global search becomes relatively more attractive. Figure 2-3 shows decision for $m = 2$ and $c = h = 1$; at the limit as $c \to 0$, global search always dominates moving to search, so the relevant choice margin is between local search and global search.

Finally, consider the case where the probability of finding a job anywhere is $h < 1$. For

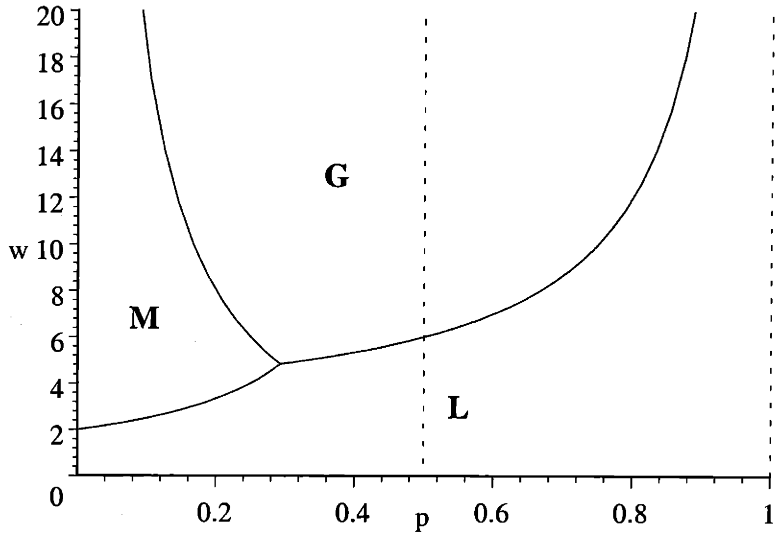

Figure 2-1: Decision Space for $m = c = 2, h = 1$

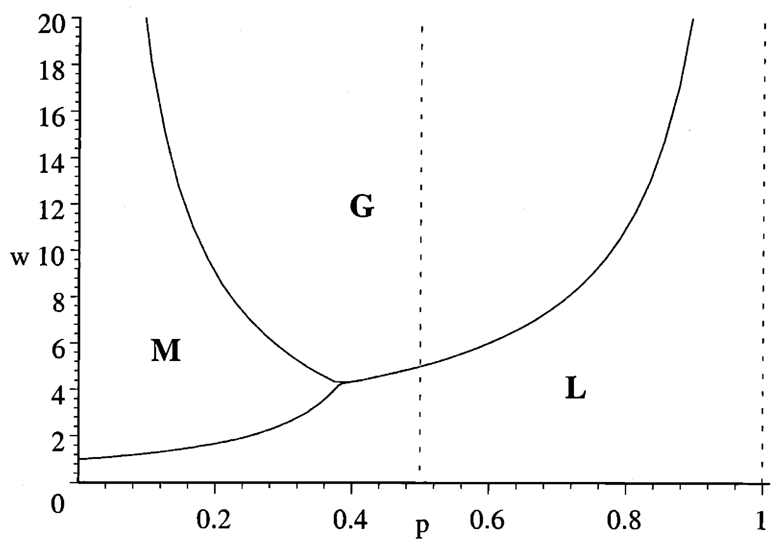

Figure 2-2: Decision Space for $m = 1, c = 2, h = 1$

$m = c = 2$ and $h = \frac{1}{2}$, the regions are then as shown in Figure 2-4. As $h \to 0$, L dominates in an ever-increasing region; at the limit, all workers search locally since there is no expected return to a global search or move.

2.2.3 Multiple Regions

Though they were derived in a 2-region setting, these results hold for the case in which there are $n > 2$ regions as well. Continue to assume that the worker's home (or origin) region is Region 1, where the probability of finding a local job is $p_1 h$. WLOG let $p_2 \geq p_3 \geq \cdots \geq p_n$ and consider the case of $h = 1$ as a baseline. By construction, the only region to which the worker would move in order to search is Region 2 (since it is a more attractive destination *ex ante* than any other region). The decision problem of the worker is therefore:

$$\begin{aligned} U^L &= p_1 w && (2.9) \\ U^G &= w - c - (1 - p_1)\, m && (2.10) \\ U^{M2} &= p_2 w - m && (2.11) \end{aligned}$$

where U^{M2} is her utility of moving to Region 2 and searching locally from there.

Area "A" in Figure 2-5 shows the domain (p_1, p_2) for the case where $n = 3$: given $p_1 \in [0, 1]$, we restrict $p_2 \in \left[\frac{1-p_1}{2}, 1 - p_1\right]$.[8] As the number of regions increases, the domain of p_2 increases to $\left[\frac{1-p_1}{n-1}, 1 - p_1\right]$; at the limit, the domain includes area "B" as well as area "A". Figure 2-6 shows the decision space on the domain ("A"+"B") in (p_1, p_2, w) space. It is easy to see that the comparative static results derived above for the 2 region case remain intact in the more general case, with slight notational modifications.

2.2.4 Implications

The model yields the following testable empirical implications:

1. High-skilled workers are more likely to migrate than are low-skilled workers. Although the propensity to migrate is, in general, not monotonic in skill – at low p workers with

[8] The upper bound on p_2 is to ensure that the 3 probabilities do not exceed 1. The lower bound is necessary and sufficient to ensure that the assumption $p_2 \geq p_3$ holds.

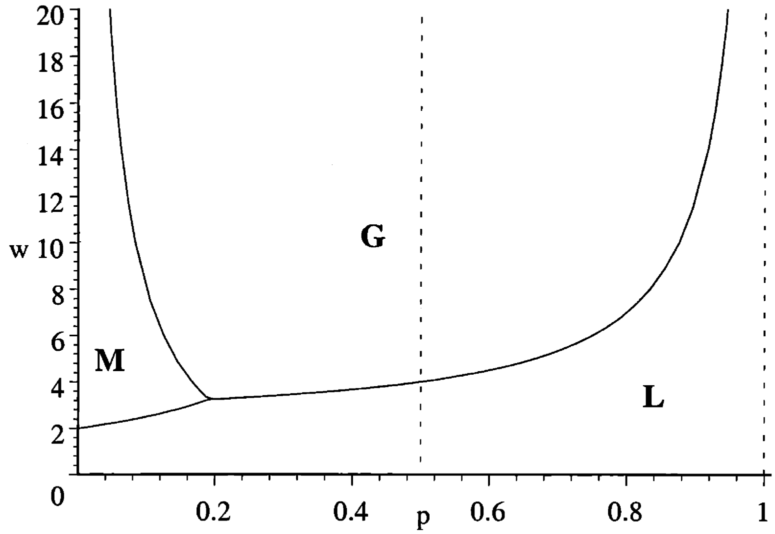

Figure 2-3: Decision Space for $m = 2, c = 1, h = 1$

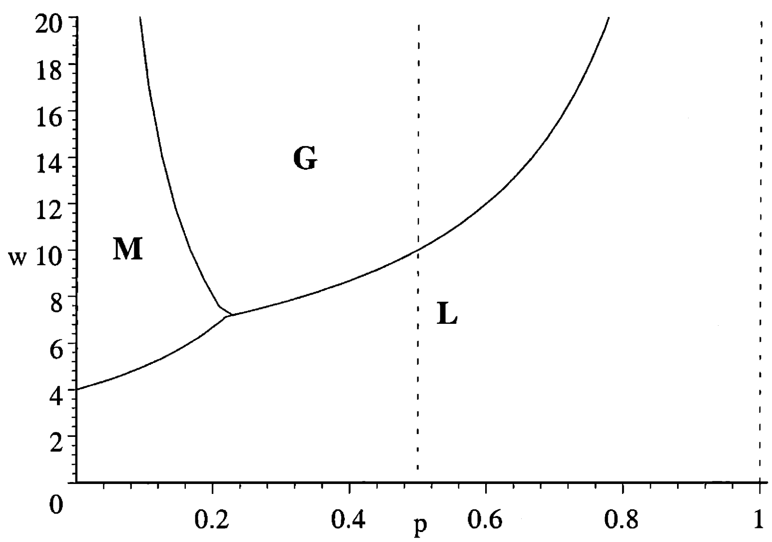

Figure 2-4: Decision Space for $m = c = 2, h = \frac{1}{2}$

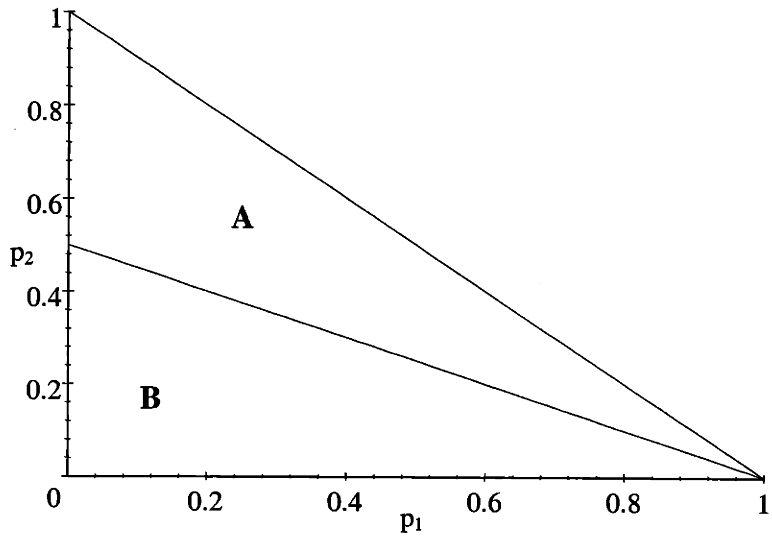

Figure 2-5: Decision Space Domain for $n \geq 3$

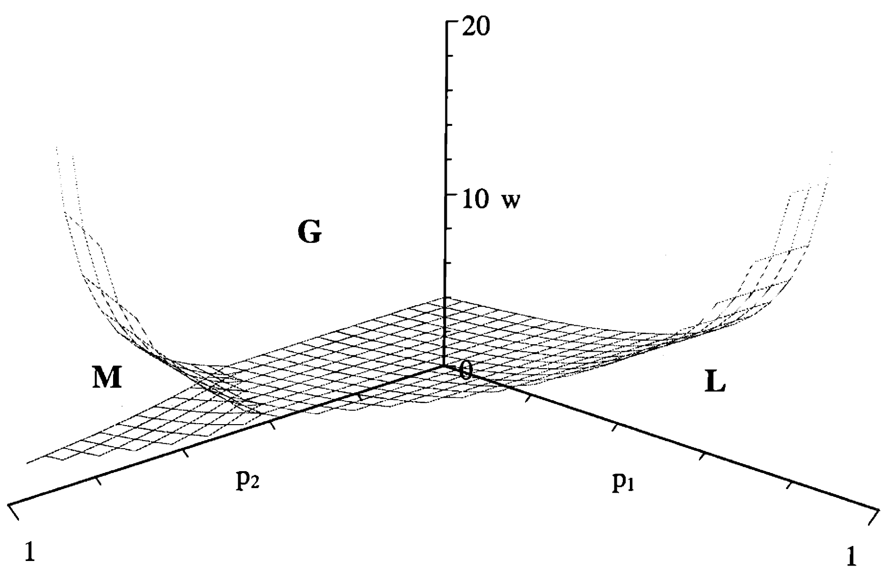

Figure 2-6: Decision Space for n-Region Case given $m = c = 2, h = 1$

intermediate wages may migrate with higher probability than high-wage workers (because of the discreteness of the M vs G choice) – it is monotonic when evaluated at the average region's conditions (i.e., at $p = \frac{1}{2}$).

2. The propensity to migrate may be non-monotonic in skill for sufficiently depressed regions, where intermediate-skilled workers will migrate at higher rates than high-skilled workers. This situation corresponds to the case where p very close to zero in the model.

3. Migration will tend to arbitrage unemployment-rate differentials:

   (a) Workers in states with bad economic conditions will be more likely to move than those living in states with good economic conditions.

   (b) Destinations will have better economic conditions than origin states.

4. Low-skilled workers' destinations will, on average, represent a larger improvement over their origins than is the case for high-skilled workers. This is because destination economic conditions figure directly in the selection of a location for low-skilled workers who move first and search later ("search migrants"), but only indirectly – by affecting the probability that a job is found – in the location choice of workers who search first and move later ("job migrants").

5. Migration is pro-cyclical: as h falls, migration decreases.

6. The effect of fluctuations in aggregate economic conditions will not be uniform across skill groups. Low-skilled workers, never very likely to move, and high-skilled workers, who search globally for a wide range of local conditions, will change their behavior only slightly. The marginal workers will be intermediate-skilled; they are likely to be most sensitive to such business-cycle fluctuations. In other words, we expect to find a non-monotonic relationship between skill and elasticity of migration with respect to business-cycle conditions.

7. For given state-level economic conditions (p and h), migrants who move before finding a job should be less skilled than those who move after finding a job. This implication, of

course, is part of the motivation for the model (see Table 2.2), but needs to be confirmed in a regression with controls for other observable characteristics.

8. Among migrants, even after controlling for skill, the probability of being employed is higher for those who searched globally and moved only after they found a job than for "searching migrants" who move first and search for work later. This is a direct consequence of the fact that global searchers move only if they find a job in the destination state, whereas searching migrants move in order to search.

These hypotheses are tested in Section 2.5 below.

2.3 Data

I use March Current Population Survey (CPS) data from 1982-2001 (excluding 1985 and 1995). For many variables, including migration, I attribute the variable values to the previous calendar year: for example, the 1982 survey supplies 1981 data. I therefore distinguish between the *survey year* (the calendar year in which the survey was administered) and the *reference year* (the year preceding the survey year). Respondents have been asked whether they moved in the last year (and where from), almost every year since 1982; exceptions are the 1985 and 1995 surveys. Since 1997 they have also been asked for the main reason for their move.

The Census Bureau selects residential addresses (dwelling units), not their occupants, for inclusion in the CPS; each dwelling unit is included in the March CPS twice, one year apart. By construction, then, non-movers are interviewed twice, whereas movers are interviewed only once – the address is visited twice but two different households respond to the survey. While the CPS questionnaire is quite thorough, most variables – labor force status (employed, unemployed, or not in the labor force), marital status, student status (full- or part-time student), and homeownership status – are available only on a current basis (survey year) but on a lagged basis (reference year).[9]

To eliminate as many non-labor-market-related moves as possible, I limited my sample to civilian adults ages 25-60 in the reference year who were not students during the survey year.[10]

[9] For further information about CPS design and methodology, see Current Population Survey (2000).
[10] In 1981-1983, I could not eliminate students due to incomplete data.

This restriction provides me with approximately 60,000 observations per year.[11]

As with many surveys, data accuracy is a concern. Questions that are not answered during a survey are replaced by imputed ("allocated") values, which are generated from other ("donor") records. Allocations can be common for some variables and can have a large effect on mean values of some variables, notably migration status. Unfortunately, records with altered or imputed data were not properly flagged by the Census Bureau until the 1996 survey (referring to 1995 migration data). Since the time series for which allocated values are properly flagged is very short (and migration is a rare event), the analysis presented here cannot be repeated using only allocated values. More details on CPS allocation procedures are presented in Appendix 2.A.[12]

2.4 Empirical Methodology

2.4.1 Migration Regressions

Because migration is a relatively rare event, even a large sample such as the CPS contains only a small number of migrant observations in any given year. I therefore use pooled cross-section data to estimate the individual-level migration equation

$$\mathbb{P}\left(\text{migrate}_{it}\right) = \Phi\left(\alpha + \beta x_{it} + \sum_{s} \sigma_s \text{state}_{ist} + \sum_{t} \delta_t \text{year}_t\right) + \mathrm{u}_{it} \tag{2.12}$$

where $\mathbb{P}(\mathrm{A})$ is the probability that event A occurs, migrate_{it} is an indicator equal to 1 if individual i moved between years t and $t+1$ (and 0 otherwise), $\Phi(\cdot)$ is the standard normal CDF, state_{ist} is an indicator equal to 1 if individual i lived in state s in year t, year_t is a year indicator, and x_{it} is a vector of additional explanatory variables. For regressions with only individual-level demographic variables, the error term u_{it} are clustered at the household level, allowing correlation between the migration decisions of spouses, as well as a across the two interviews of each dwelling unit. For regressions where the coefficient of interest varies only

[11]Because the interpretation of interstate migration is ambiguous for individuals living in Washington D.C., I omit both current D.C. residents and individuals who moved out of D.C. in the past year. None of the results reported here are sensitive to this omission.

[12]To complicate matters further, allocation procedures changed in 1988. More details on this change are in the Appendix.

by state and year, u_{it} is clustered at the state level. By pooling the data, I implicitly assume that the effect of individual characteristics, such as education, and economic variables on the propensity to migrate is not time-varying.[13]

Year fixed effects are included in some, but not all, of the regressions. Year-to-year fluctuations in the migration rate due to unobservable changes will be captured by these year fixed effects, when included. Year fixed effects are omitted when aggregate economic conditions are included in the regression.

Unless otherwise noted, I report the derivatives $\frac{\partial\Phi(x,z)}{\partial x}|_{(\overline{x},\overline{z})}$ (where z is the vector of all explanatory factors excluding x) rather than probit coefficients. The numbers reported may therefore be interpreted as the effect of an infinitesimal change in the variable of interest, x, on the probability that the "average" person migrates. In cases where x is a binary variable (such as an indicator for race, sex, or education), I report instead the change in the probability of migration associated with a discrete change in x:

$$\Phi\left(x=1;\overline{z}\right)-\Phi\left(x=0;\overline{z}\right).$$

Most regressions use all available observations, with the exception of a few focusing only on movers. Because women are more likely than men to move for reasons other than work (specifically, for a spouse or other family member), I have also estimated results using a male-only sample. The results tend to be qualitatively similar though most effects are magnified because of men's higher sensitivity to labor-market conditions.[14] I report the male-only results only when they are sufficiently different from the results for the full sample to be of independent interest.

2.4.2 Unemployment Rates

I compute state-level unemployment rates, as well as state unemployment rates by education category (high-school dropouts, high-school graduates, some college, and college degree or be-

[13] For completeness, I also estimate the equivalent OLS regression
$\text{migrate}_{it}=\alpha+\beta x_{it}+\sum_s\sigma_s\text{state}_{st}+\sum_t\delta_t\text{year}_t+u_{it}.$

Results from these regressions are uniformly extremely similar to the probit results, and are therefore not reported.

[14] Of course, since the sample is about half the size, some significance is lost despite the (absolutely) larger estimated coefficients

yond), using records of male non-migrants ages 25-60.[15]

The arbitrage motive for migration, by the model, depends on the difference between the unemployment rate in a worker's current region of residence and the unemployment rate in other regions to which the worker could potentially move. I compute the "target unemployment rate" for a worker in state s as

$$\text{targetue}_{st} = \sum_{m \neq s} \omega_{sm} \text{unemp}_{mt}$$

where unemp_{mt} is the unemployment rate in state m in year t, and the weights ω_{sm} represent the share of movers from state s who moved to state m over the period 1981-2000:

$$\omega_{sm} = \frac{\text{migrants}_{sm}}{\sum_{k \neq s} \text{migrants}_{sk}}. \tag{2.13}$$

The target unemployment rate can therefore be interpreted as the unemployment rate in the "average" state of destination for state-s out-migrants.

The relevant variable is then the difference between the worker's current state's unemployment rate and the unemployment rate in the target area, evaluated before migration:

$$\text{targetdiff}_{st} = \text{unemp}_{st} - \text{targetue}_{st}. \tag{2.14}$$

One issue that warrants mention is measurement error. The predications of the model in Section 2.2 have all to do with the *individual's* employment prospects, not with the average employment opportunity in her state. In the empirical analysis, the state unemployment rate is used to proxy for individual employment prospects, but as such it is measured with error. Attenuation bias in the coefficient on the differential unemployment rates is therefore to be expected, both because the unemployment rate is itself measured with error (especially when broken down by education category, where very few observations inform each calculation), and because it is an imperfect proxy.

[15]The unemployment rate by state and education is measured with more error than the average unemployment rate in the state due to the relatively small number of observations in every state*year*education cell. 258 of 3600 state*year*education cells – more than half of them for the highest education category – have no observations of unemployed workers. These cells are assigned an unemployment rate of zero.

2.5 Results

2.5.1 Baseline Regression

In this section I test hypotheses (1) and (2) from Section 2.2.4. Recall, these are:

1. Evaluated at average regional conditions, high-skilled workers are more likely to migrate than are low-skilled workers.

2. The propensity to migrate may be non-monotonic in skill for sufficiently depressed regions.

To test these hypotheses, the baseline regression includes only individual-level explanatory variables: age, sex, race, and education, as well as state and (for some regressions) year fixed effects. I report the coefficients on these demographic variables from probit regressions in Table 2.3. Coefficients on age fixed effects from Column (1) are plotted in Figure 2-7. Three education categories are reported (high school diploma, some college, college degree or beyond); the omitted education category is less than high school. As expected, migration increases monotonically with education, and decreases with age.

When the number of weeks worked last year is included in the regression (either as a continuous variable or as an indicator for 50 or more weeks worked), the estimated effect of education increases. Since more-educated workers are less likely to be unemployed, and since the unemployed may be more likely to migrate, controlling for education without controlling for labor-force status biases downward the estimated coefficient on education. Controlling for the number of weeks worked last year therefore increases the measured effect of education on migration. Since the number of weeks worked last year is itself endogenous to the migration decision, however, results with these controls should be interpreted with care. The coefficients on weeks of work should be interpreted with caution, if at all.[16]

Next, I test the prediction that propensity to migrate may not be monotonic in education for regions with sufficiently high unemployment rates. Let bad_{st} be an indicator for bad economic

[16]To see why the number of weeks worked is endogenous, note that (1) moving takes time, and (2) workers who move in order to search for work are expected to be unemployed for a time in their new location. The number of weeks worked is also a noisy proxy for the worker's labor-force status the previous year, which is the (unavailable) control variable of interest. While the endogeneity biases the coefficient upwards (in absolute terms), the measurement error biases it towards zero; whether the true effect of having been employed one year ago is larger or smaller than the one estimated is impossible to say with the available data.

Table 2.3: Probability of Migration: Baseline Regression Estimates

Variable	(1)	(2)	(3)	(4)	(5)	(6)
Male	0.174	0.483	0.499	0.175	0.489	0.505
	(0.027)	(0.027)	(0.026)	(0.027)	(0.027)	(0.026)
White	-0.055	-0.009	-0.041	-0.053	0.004	-0.026
	(0.065)	(0.061)	(0.061)	(0.064)	(0.061)	(0.060)
High School	0.001	0.273	0.293	-0.004	0.263	0.279
Diploma Exactly	(0.067)	(0.067)	(0.066)	(0.067)	(0.067)	(0.066)
Some College,	0.474	0.832	0.862	0.475	0.818	0.841
No Degree	(0.076)	(0.079)	(0.078)	(0.076)	(0.079)	(0.078)
College Degree	1.452	1.964	1.993	1.446	1.941	1.961
or Beyond	(0.086)	(0.092)	(0.091)	(0.086)	(0.091)	(0.090)
Weeks Worked		-0.075			-0.075	
Last Year (Number)		(0.001)			(0.001)	
Worked 50+ Weeks			-3.209			-3.238
Last Year (Indicator)			(0.060)			(0.061)
χ^2 test for equality of	793	1,100	1,145	790	1,085	1,123
education coefficients	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Age Fixed Effects	Y	Y	Y	Y	Y	Y
Origin State FE	Y	Y	Y	Y	Y	Y
Year Fixed Effects	N	N	N	Y	Y	Y

Notes: 944,061 observations used. Standard errors are clustered at the household level. All coefficients are multiplied by 100.

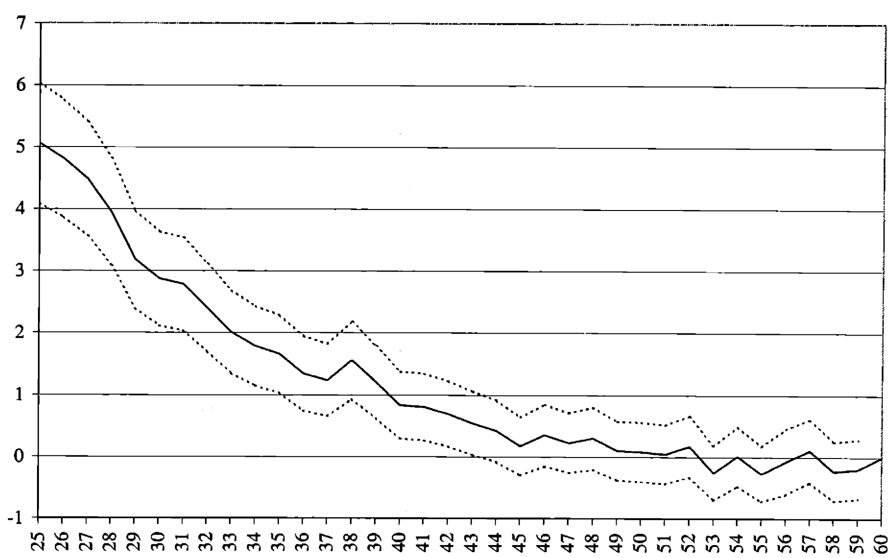

Figure 2-7: Coefficients on Age Indicators ($\beta_{60} \equiv 0$)

conditions in the state relative to the "target area" for potential migrants:

$$\text{bad}_{st} = \begin{cases} 1 & \text{if targetdiff}_{st} > 0 \\ 0 & \text{otherwise,} \end{cases}$$

where targetdiff_{st} is as defined in Equation (2.14). Table 2.4 presents coefficients on *interactions* of education and bad economic conditions, to test whether the effect of education is monotonic in the good regime (when $\text{bad}_{st} = 0$) but non-monotonic in the bad regime ($\text{bad}_{st} = 1$). Each regression is presented in two columns, the first showing coefficients on education for state-year cells with the bad regime, and the second showing the coefficients for state-year cells with the good regime.

The first two columns show coefficients from a regression with no year fixed effects; the last two columns repeat this exercise with year fixed effects. In each column, I show first the coefficients on the interaction terms, followed by a χ^2 test for equality of the coefficients by education for each regime. In every case, equality of the coefficients on education *within* the regime is clearly rejected.

Below this first χ^2 test statistic I show the χ^2 statistic of interest: testing whether changes in the probability of migration change with education *differently* across the two regimes. The prediction that intermediate-skilled workers will migrate more than both low- and high-skilled workers out of states with bad economic conditions fails: there is no statistical difference between the elasticity of migration with respect to education in the bad regime and that in the good regime.

2.5.2 Arbitrage

In this section I test hypotheses (3) and (4) from Section 2.2.4:

3. (a) Workers in states with bad economic conditions will be more likely to move than those living in states with good economic conditions.

   (b) Destinations will have better economic conditions than origin states.

4. Low-skilled workers' destinations will, on average, represent a larger improvement over their origins than is the case for high-skilled workers.

Before turning to regression analysis, I first present a simple tabulation of the average improvement in migrants' state-level unemployment rate. Table 2.5 tabulates the fraction of migrants whose destination-state unemployment rate (measured in year t, before the move) is lower than their year-t origin-state unemployment rate, by education. I show results using both average state unemployment rates (Column 1) and education-specific state unemployment rates (Column 2), both constructed from March CPS files using male non-movers ages 25-60.

The table confirms hypotheses (3b) and (4). On average, migrants move to states with lower unemployment rates than their states of origin *for their education category*, though not necessarily to states with lower overall unemployment rates. As the model predicts, the fraction of migrants whose destinations have lower unemployment rates than their origins decreases monotonically with education.

To test hypothesis (3a), I add the variable targetdiff_{st} (the difference between the state unemployment rate and the unemployment rate in the state's "target region") to the regression presented in Table 2.3. Table 2.6 shows the estimated effect of targetdiff_{st} on migration. Column (1) shows results using the average target unemployment-rate differential. Columns (2) and (3) control for weeks worked last year (as a continuous variable and as a discrete variable, respectively); Columns (4)-(6) repeat these regressions with year fixed effects. All other controls from Table 2.3 are included in all regressions.

When no controls for employment status are included, the effect of targetdiff_{st} is positive and significant: the higher is a state's unemployment rate, relative to the region to which its residents are likely to migrate, the higher is the probability that they will move. For the regressions with year fixed effects, this result continues to hold when weeks of work are included in the regression; the results are marginally-significant (significant at the 10% confidence level) when no year fixed effects are included. As above, it is important to interpret the results of regressions which control for weeks of work with care, since the number of weeks worked is endogenous to the migration decision.

2.5.3 Cyclical Patterns

I now turn to testing the hypotheses regarding the cyclical behavior of migration:

5. Migration is pro-cyclical.

Table 2.4: Baseline Regression Estimates with *bad* interaction

	(1)		(2)	
Variable	bad = 1	bad = 0	bad = 1	bad = 0
High-School Dropout	(dropped)	-0.277 (0.111)	(dropped)	-0.226 (0.122)
High School Diploma Exactly	-0.107 (0.094)	-0.181 (0.099)	-0.071 (0.093)	-0.164 (0.097)
Some College, No Degree	0.308 (0.109)	0.327 (0.114)	0.369 (0.108)	-0.342 (0.112)
College Degree or Beyond	1.358 (0.134)	1.320 (0.137)	1.439 (0.131)	1.316 (0.135)
χ^2 test for equality of education coefficients	364.55 0.0000	445.53 0.0000	398.28 0.0000	411.25 0.0000
χ^2 test for joint significance of *bad* interactions	6.54 0.1626		5.84 0.2113	
Year Fixed Effects	N		Y	

Notes: 944,061 observations used. Regressions include all controls from Table 2.3. Standard errors are clustered at the household level. All coefficients are multiplied by 100.

Table 2.5: Unemployment Rate Arbitrage by Migrants' Education

	Fraction Moving to States with Lower...	
	Unemployment Rate	Education*UE Rate
All Movers	48.98%	52.51%
HS Dropouts	53.69%	54.59%
HS Graduates	49.78%	54.20%
Some College	48.65%	52.81%
College Grads +	47.40%	50.49%

Source: Author's calculations from CPS, 1981-2000

Table 2.6: Arbitrage Regression Estimates

Variable	(1)	(2)	(3)	(4)	(5)	(6)
targetdiff$_{st}$	5.431 (2.410)	3.812 (2.297)	3.572 (2.205)	5.849 (2.010)	4.345 (1.905)	4.156 (1.839)
Weeks Worked Last Year (Number)		-0.075 (0.002)			-0.075 (0.002)	
Worked 50+ Weeks Last Year (Indicator)			-3.206 (0.096)			-3.234 (0.093)
Year FE	N	N	N	Y	Y	Y

Notes: 944,061 observations used. Regressions include all controls from Table 2.3. Standard errors are clustered at the state level. All coefficients are multiplied by 100.

6. The cyclicality of migration is strongest for workers with intermediate skills (education).

Table 2.7 shows coefficients on the U.S. unemployment rate when interacted with education variables. In Column (1), the US unemployment rate is interacted with individuals' education categories. In Columns (2) and (3), controls are added for weeks worked last year (continuously and discretely, respectively). Column (4) shows results similar to Column (1) with the difference that individuals' education categories are interacted with the unemployment rate *by education category.* Columns (5) and (6) repeat this specification again controlling for weeks worked last year. Standard errors are clustered at the education-category level. Note that education enters into these regressions directly as well as through the interaction terms.

The χ^2 statistic shows a test for equality of the effect of business cycles conditions across educational categories. In all cases a χ^2 test rejects equality of the coefficients across education groups. The estimated effect of the business cycle is nonlinear: the migration rate of high-school graduates is more sensitive to business-cycle conditions that that of high-school dropouts. This nonlinearity is strong in some specifications and almost invisible in others.

2.5.4 Reasons for Migration

In this section I test whether the reasons for migration are sensitive to observable variables. Since questions about the main reason for migration were not asked before 1997, only a small set of observations is available to answer this question, and some variables (specifically, cyclical ones) cannot be used on the right-hand side. Moreover, since workers who search globally migrate with probability $(1-p)\,h < 1$, only a fraction of global searchers will be included in the migrant sample.

With these caveats in mind, define "job movers" as movers who moved to accept a job (following a global search) and "search movers" as movers who move to search for a job. I first show that search movers are more likely to arbitrage unemployment rate differences across states than are job movers. Table 2.8 shows the average arbitrage in unemployment rates by type of migration for 1997-2000 (the years for which type of migration is solicited). Strikingly, a statistically-significant 64% of search movers move to states with lower unemployment rates, while only 52% of job movers do so (and for them the fraction is not statistically different from half); the numbers are slightly smaller, but the pattern the same, when education-specific unemployment rates are used.

I next turn to testing the model's predictions about differences between job movers and search movers:

7. For given state-level economic conditions (p and h), search migrants will be less skilled than job migrants.

8. Among migrants, even after controlling for skill, the probability of being employed is higher for job migrants than for search migrants.

To verify that the stylized fact presented in Table 2.2 – that workers with higher education are more likely to be job-movers than search-movers, with the converse being true of less-educated workers, I limit my sample to job movers and search movers (removing all other movers, as well as all non-movers), and run the regression

$$\mathbb{P}\left(\text{jobmove}_{it} \mid \text{migrate}_{it}\right) = \Phi\left(\alpha + \boldsymbol{\beta} x_{it} + \sum_t \delta_t \text{year}_t\right) + \text{u}_{it} \tag{2.15}$$

where jobmove$_{it}$ is an indicator equal to 1 if the migrant is a "job mover", and x_{it} are demographic variables – sex, race, age, and education – on the remaining subsample. The results are shown in Table 2.9. No state-of-origin fixed effects are included due to the small sample size. Column (1) shows results using a full set of age fixed effects; because of the small sample, Column (2) uses fixed effects for 6-year age groups. Columns (3) and (4) repeat Columns (1) and (2) with the addition of year fixed effects. The results show that, even after controlling for age, the probability that a worker is a job-mover increases with education.

Finally, we turn to outcomes. How important is the type of migration to employment outcomes? To answer this question, I regress current employment status on type of move, controlling for the same household characteristics:

$$\mathbb{P}\left(\text{employed}_{i,t+1} \mid \text{migrate}_{it}\right) = \Phi\left(\alpha + \boldsymbol{\beta} x_{it} + \kappa \text{jobmove}_{it} + \sum_t \delta_t \text{year}_t\right) + \text{u}_{it} \tag{2.16}$$

where employed$_{i,t+1}$ is an indicator which equals 1 if the worker is employed in year $t+1$, and zero otherwise (i.e., for both the unemployed and non-participants). The results are shown in

Table 2.7: Cyclicality of Migration by Education Category

Variable	(1)	(2)	(3)	(4)	(5)	(6)
Unemp Variable		US Average			US Average by Education	
High School	-8.889	-14.531	-15.869	-4.370	-7.750	-8.553
Dropout	(0.525)	(0.635)	(0.600)	(0.317)	(0.396)	(0.362)
High School	-10.491	-14.925	-16.104	-9.063	-12.690	-13.588
Graduate	(0.580)	(0.696)	(0.550)	(0.453)	(0.548)	(0.431)
Some College,	-1.769	-6.056	-7.409	-4.328	-9.095	-10.540
No Degree	(0.820)	(0.948)	(0.847)	(0.854)	(0.987)	(0.874)
College Degree	1.315	-1.107	-1.686	5.059	-0.062	-2.062
or Beyond	(0.786)	(0.856)	(0.830)	(1.309)	(1.467)	(1.397)
Weeks Worked		-0.076			-0.076	
Last Year (Number)		(0.003)			(0.003)	
Worked 50+ Weeks			-3.249			-3.247
Last Year (Indicator)			(0.159)			(0.161)
χ^2 test for equality	441,286	26,992	149,981	7,208	17,500	8,788
of interaction terms	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Notes: 944,061 observations used. Regressions include all controls from Table 2.3. Standard errors are clustered at the education-category level. All coefficients are multiplied by 100.

Table 2.8: Unemployment Rate Arbitrage by Migration Type

	Fraction Moving to States with Lower...	
	Unemployment Rate	Education*UE Rate
All Movers	53.15%	52.28%
Job Movers	52.34%	50.88%
Search Movers	64.26%	58.86%

Source: Author's calculations from CPS, 1997-2000

Table 2.9: Regression Results for Job-Mover Characteristics

Variable	(1)	(2)	(3)	(4)
Male	-1.144	-1.543	-1.042	-1.488
	(1.103)	(1.174)	(1.112)	(1.185)
White	5.510	5.696	5.668	5.891
	(2.911)	(3.140)	(2.953)	(3.183)
High School	4.052	3.552	4.094	3.580
Diploma Exactly	(1.812)	(2.007)	(1.798)	(2.001)
Some College,	6.602	6.116	6.655	6.179
No Degree	(1.711)	(1.873)	(1.703)	(1.867)
College Degree	23.388	22.460	23.271	22.322
or Beyond	(3.819)	(3.680)	(3.815)	(3.664)
χ^2 test for equality of	108.71	94.47	106.51	94.07
education coefficients	0.0000	0.0000	0.0000	0.0000
Age Fixed Effects	36	6	36	6
Year Fixed Effects	N	N	Y	Y
Observations	1830	1845	1830	1845

Notes: Standard errors are clustered at the household level.
All coefficients are multiplied by 100.

Table 2.10. Columns (1) and (2) show results for the full sample of movers, without and with year fixed effects, respectively; Columns (3) and (4) repeat the analysis for men only.

These results suggest that, conditional on migration, male workers who moved to take jobs are 13% more likely to be employed the following March than male workers who moved to search for work. Interestingly, controlling for type of move, education does not significantly change the probability of being employed. Unfortunately, these results are probably driven, at least in part, by an unaddressed endogeneity problem: workers who move to search for a job were probably unemployed for a while before they moved, whereas workers who move to take jobs (job movers) are more likely to have engaged in on-the-job search before the move. This implies that being a job-mover is probably correlated with other, unobserved, characteristics that make the worker more employable in any location. It is therefore hard to judge how much of the increased probability of being employed is due to the type of migration; the 13% figure (16% for both sexes) should be taken as an upper bound.

Table 2.10: Employment Probability for Migrant Sub-Sample

	All Movers		Men Only	
Variable	(1)	(2)	(3)	(4)
Job Mover	16.091	16.098	13.229	12.973
	(4.369)	(4.319)	(4.237)	(4.228)
Male	18.280	18.379		
	(1.947)	(1.942)		
White	-1.221	-1.159	-0.372	-0.210
	(2.724)	(2.792)	(2.631)	(2.649)
High School	0.652	0.467	0.243	1.647
Diploma Exactly	(3.703)	(3.752)	(2.940)	(2.915)
Some College,	1.184	1.105	1.422	1.299
No Degree	(3.749)	(3.761)	(2.672)	(2.654)
College Degree	3.482	3.306	3.239	3.053
or Beyond	(3.877)	(3.879)	(3.190)	(3.110)
χ^2 test for equality of	2.26	2.16	2.51	2.39
education coefficients	0.5203	0.5400	0.4734	0.4959
Age Fixed Effects	6	6	6	6
Year Fixed Effects	N	Y	N	Y
Sex Composition	M & F		M only	
Observations	1861	1861	1088	1088

Notes: Standard errors are clustered at the household level.
All coefficients are multiplied by 100.

2.6 Conclusion

This paper presents a very simple one-period consumer-choice model of migration and job search intended to capture two stylized facts: the positive relationship between education and propensity to migrate, and the tendency of more-educated workers to find work before they move, while less-educated workers move first and search for work second. Several predictions about aggregate migration behavior emerge from the model, and are tested empirically, with very good results.

First, I test whether the monotonic relationship between education and migration holds both in below-average and above-average states. I find that the relationship is somewhat more stable in states with worse-than-average economic conditions than in states with better-than-average conditions. The strong prediction of the model, that intermediate-skilled workers will have the highest out-migration rates from worse-than-average states fails empirically. At the same time, the difference between the out-migration rates of the highly-educated and the low-educated is smaller in these worse-than-average states, as the model predicts.

Second, I show that, *conditional on moving*, low-skilled workers are more sensitive to relative unemployment rates than are high-skilled workers. I also look at the sensitivity of different educational groups to business-cycle conditions, and find, as the model predicts, that intermediate-skilled workers are most sensitive to business-cycle conditions in their migration decision.

Finally, I show that search movers are substantially less likely to be employed following their move than are job movers.

The model presented here is not intended to capture all aspects of the migration decision; as Table 2.2 (in the Introduction) shows, nearly half of all migrants give reasons other than work for their decision to move. And while the model imposes identical preferences and identical search-and-migration technologies on all workers (allowing them to differ along a single dimension – wages, assumed to increase monotonically with skill), in reality there are many other differences between low- and high-skilled workers. On the preference dimension, workers may care differentially about their career. If skill is acquired, workers who are "career-minded" may choose to acquire skill and, concurrently, be more willing to migrate even when the expected gain is small. In the terminology of the model presented here, this would imply a correlation between the (psychic) cost of moving and skill: $\rho(m, w) < 0$. On the technology dimension, skilled workers may face lower global-search costs, so that c may decrease with skill. Such modifications to the model would increase the migration rate of the high-skilled relative to the low-skilled, and further decrease the sensitivity of high-skilled workers to cyclical patterns.

Appendix to Chapter 2: Data Issues

Allocated Values

As mentioned in Section 2.3, missing data in the CPS are replaced by allocated values, which are generated from other ("similar") records. Unfortunately, records with altered or imputed data were not properly flagged by the Census Bureau before 1995. One variable that is particularly susceptible to allocation is the migration variable. Table 2.11 lists the number of records with an allocated migration status for each year, their relative weight in the sample, and the propensity of allocated records to be coded as migrations. Beginning with the 1996 survey, over a thousand observations annually are allocated, and the fraction of these observations that are assigned migrant status increases sharply over time, to nearly 60% by the 2001 survey.[17]

1988 Processing Changes

Following a change in the CPS processing system in 1988, the 1988 survey data were re-released, having been processed using the new system. There are therefore two files containing 1988 data, the first of which was used by the Census Bureau to produce their reports, and the second, known as the "bridge" file (or, alternatively, as the 1988 rewrite file or the 1988B file), intended to facilitate comparisons to subsequent years. Since the input data – the pool of respondents, the survey questions and answers – are identical across the two 1988 files, one should in principle be able to use either one for analysis. Unfortunately, these processing changes were not completely benign. Among the changes made in the re-processing were changes to the imputation procedures for missing data (Bureau of the Census 1991).

While the demographic characteristics of respondents (age, sex, occupation, marital status, and race) are statistically indistinguishable across the two files, as seen in Table 2.12, migration data are disconcertingly different in these two surveys, by a statistically significant margin. Note that the p-value is given here only as a reference, since the differences between the two data sets are all due to imputation. As these are not actually two separate samples, it is not clear what interpretation, if any, a t-test can be given in this case.

The change in the migration estimation is due to 445 records which were coded as movers

[17] Allocation flags were not included in the data prior to the 1988 bridge file.

Table 2.11: Allocations and Migration in CPS Data

Survey	Observations	Migration Allocations	Allocation Weight	Allocated Migration[a]
1988B	66,828	0	0	n/a
1989	62,477	8	<0.001	0
1990	68,121	5	<0.001	0
1991	68,341	6	<0.001	0
1992	67,613	0	0	n/a
1993	67,179	0	0	n/a
1994	63,822	5	<0.001	0
1996	55,000	1,125	0.023	0.213
1997	55,666	1,217	0.025	0.228
1998	56,259	1,052	0.020	0.209
1999	56,524	1,160	0.023	0.409
2000	56,718	1,149	0.022	0.561
2001	54,754	1,003	0.021	0.596

[a] Fraction of allocated observations that are assigned migrant status

Table 2.12: Summary Statistics for 1988 Surveys

Variable (I=Indicator)	Mean 1988[a]	Mean 1988B	p-Value for Equality
Interstate Migration (I)	0.025	0.028	0.002
Age	39.75	39.74	0.789
Male (I)	0.487	0.486	0.845
White (I)	0.858	0.859	0.797
High School Dropout (I)	0.178	0.178	0.922
High School Graduate (I)	0.378	0.378	0.921
Some College, No Degree (I)	0.215	0.215	0.905
College Degree or Beyond (I)	0.229	0.229	0.930
Observations	66,504	66,828	

Notes: Means include past and present DC residents. All means are weighted. Means are reported for non-student civilian adults ages 25-60. Hypothesis tests assume equal variance across surveys.

in one file but as non-movers in the other. These include 116 that were coded as movers in the original file (but not in the bridge file), and another 329 that were coded as movers in the bridge file (but not in the original file). A frightening 684 additional records are coded as migrants in both files, but their state of origin differs across files. According to the Census Bureau, these changes had the effect of "making migration recodes more consistent with residence fields" (US Census Bureau 1991). In the data analysis I rely exclusively on the bridge file for 1988 data (1988B), in the hope that the processing changes improved the data quality.

References

Bartel, Ann P. (1979). "The Migration Decision: What Role Does Job Mobility Play?" *American Economic Review* 69:5.

Blanchard, Olivier J. and Lawrence F. Katz (1992). "Regional Evolutions." *Brookings Papers on Economic Activity.*

Coulson, N. Edward, Derek Laing, and Ping Wang (2001). " Spatial Mismatch in Search Equilibrium." *Journal of Labor Economics* 19:4.

DaVanzo, Julie (1978). "Does Unemployment Effect Migration? Evidence from Micro Data." *Review of Economics and Statistics* 60:4.

Gabriel, Stuart A., Janice Shack-Marquez, and William L. Wascher (1993). "Does Migration Arbitrage Regional Labor Market Differentials?" *Regional Science and Urban Economics* 23:2.

Greenwood, Michael J. (1975). "Research on Internal Migration in the United States: A Survey." *Journal of Economic Literature* 13:2.

Greenwood, Michael J. (1993). "Migration: A Review." *Regional Studies* 27:4.

Mauro, Paolo, and Antonio Spilimbergo (1999). "How Do the Skilled and the Unskilled Respond to Regional Shocks?: The Case of Spain." *IMF Staff Papers* 46:1.

Pissarides, Christopher A., and Jonathan Wadsworth (1989). "Unemployment and the Inter-regional Mobility of Labour." *Economic Journal* 99:397.

Schlottmann, Alan M., and Henry W. Herzog, Jr. "Employment Status and the Decision to Migrate." *Review of Economics and Statistics* 63:4.

Spilimbergo, Antonio, and Luis Ubeda (2001). "A Model of Multiple Equilibria in Geographic Labor Mobility." Mimeograph.

Steiner, Clyde and Shari (1999). *Steiner's Complete How-to-Move Handbook.* San Francisco: Independent Information Publications.

U.S. Department of Commerce, Bureau of the Census (1991). *Current Population Survey: Annual Demographic File, 1988 Rewrite Supplement.* Washington, D.C.: U.S. Department of Commerce, Bureau of the Census.

Chapter 3

Do Temporary Shocks Have Permanent Effects? Evidence from the Agricultural Sector

> The crops we grew last summer weren't enough to pay the loans \
> Couldn't buy the seed to plant this spring and the Farmers Bank foreclosed \
> Called my old friend Schepman up to auction off the land \
> He said John it's just my job and I hope you understand \
> – John "Cougar" Mellencamp, *Rain On The Scarecrow*

3.1 Introduction

Do exogenous temporary shocks – such as recessions – have permanent effects on the distribution of resources? Asked differently, do recessions lead to permanent reallocation of resources from credit-constrained agents to agents who are not constrained? This paper tries to address these questions empirically using data on shocks to crop yields and crop prices and their effect on farmland distribution. I use data from the quintennial (five-year) Agricultural Census to isolate changes in farmland ownership – by size of farm and legal form of organization (family- vs. corporate-ownership) – at the county level. I then use variation in crop yields across crops and counties, and national-level variation in crop prices over time, to estimate the effect of income shocks on changes in farmland ownership. I find that negative deviations of crop revenues from long-run trends reduce the number of small and family-owned farms.

The number of farms, especially small and owner-operated farms, has been declining in the United States since the 1930s; this decline has raised concerns since its beginning (see Black and Allen 1937). Figure 3-1 shows the total number of farms (in thousands) in the US from 1910 to the present; the dotted lines bracket the period 1978-1997, which is the focus of most of this paper.[1]

As the number of farms has declined – and the size of the average farm has grown – farm revenues have fallen. This trend appears to be driven primarily by increased productivity in the agricultural sector. Yield per acre has been increasing steadily for many crops – due to better harvesting technology, disease-resistant seeds, and other improvements – while demand remains relatively inelastic. The result is a decline in revenue per acre for most traditional crops. Figure 3-2 shows the decline average revenue per acre for corn over the period 1972-1998. Other crops follow similar trends, with revenues declining by an average 2-4% annually. To remain profitable, farms have had to reduce their costs similar rates.

Declining profitability may be an important part of the explanation for the decline in the number of farms, at least in the last 30 years. This paper asks whether, and to what extent, annual fluctuations in profitability – above and beyond the trend – have contributed to the decline. Have credit constraints forced farms that were viable long-run concerns to exit following especially bad years?

My study focuses on fourteen crops that, among them, cover just under 30% of US farmland and approximately 65% of cropland. I use data from the quintennial Census of Agriculture to estimate the effect of revenue shocks on the number of family-owned, small and large farms at the county level. I identify county-level shocks to average revenue per acre by relying on two sources of variation across counties: variation in the allocation of farmland to different crops – which results in differential impact of crop-price fluctuations on average revenue – and differential innovations to crop yields, most likely due to location-specific weather shocks.

I find an effect of negative shocks to revenues on exit of farms, especially small and family-

[1]While the number of farms – especially the number of small and family-owned farms – has declined in most states over the period 1978-1997, it has increased in a few states, mainly in the Northeast (Massachusetts, Rhode Island, New Jersey, Connecticut) and the Southwest and West (California, Oregon, Utah, Montana, New Mexico, Nevada, Colorado).

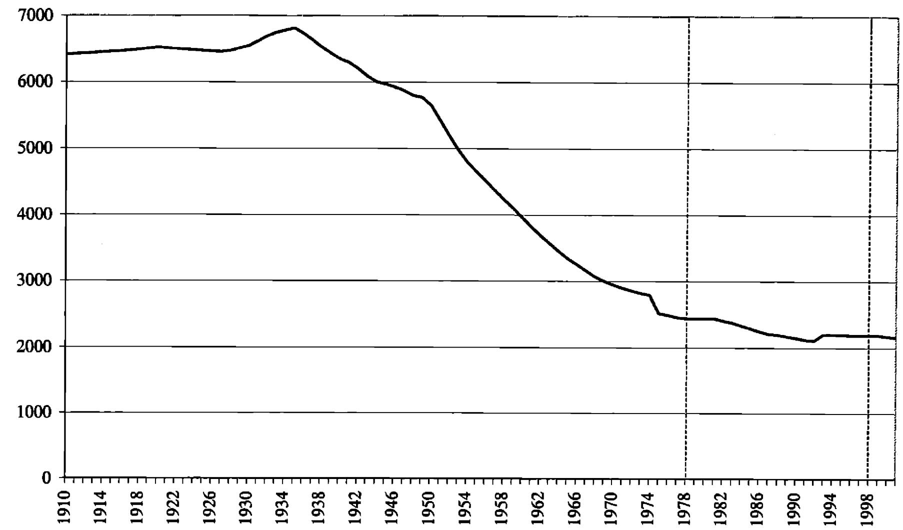

Figure 3-1: Number of Farms, 1910-2001 (Thousands)

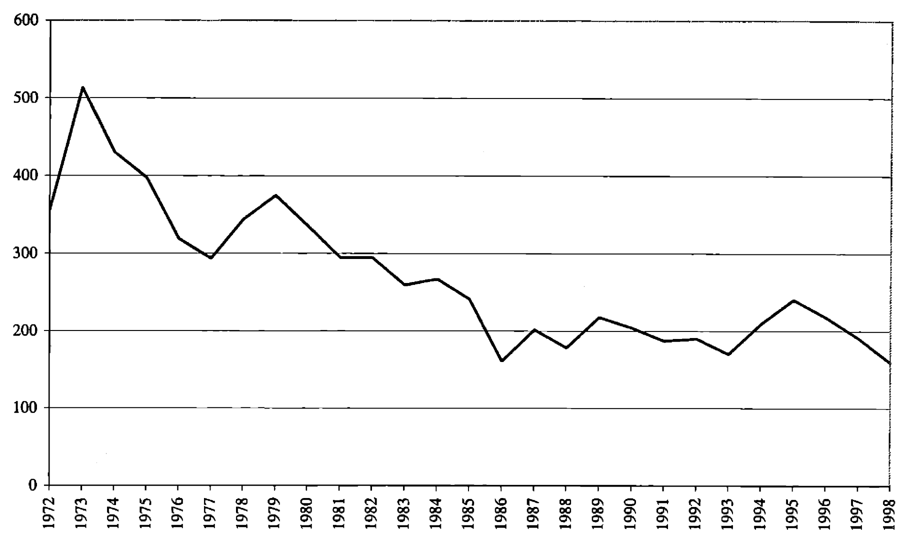

Figure 3-2: Averge Corn Revenue per Acre, 1972-1998

owned farms; this effect is extremely persistent and robust to changes in specification than the effect of shock on number of large farms. The effect on the number of large farms is ambiguous, and changes with the specification. This finding is consistent with recent analysis by Roberts and Key (2002), who find that small farms contract following negative yield shocks while large farms expand following the same negative shocks. This finding does not suggest a strong prior for the effect of shocks on the number of large farms though it does predict that the number of small farms declines following negative shocks. The effect on the number of large farms depends on whose land they buy: if already-large farms buy smaller farms in bad times, the number of large farms will remain unchanged; if they buy other large farms, the number of large farms will decrease during bad times; and if prosperous medium-sized farms expand to become large farms during bad times, the number of large farms will increase.

Several mechanisms are in place to protect farmers from some of the annual variation in revenue they inevitably experience. Variation in yields due to floods, hail damage, disease, drought, and the like, may be mitigated by crop insurance. Aggregate data on crop insurance show that between one quarter and one half of the acreage planted with the crops studied here is covered by crop insurance; for most of the crops studied, this fraction is closer to one quarter than to one half (US Department of Agriculture, various years). The federal government provides additional protections to small farmers from income fluctuations, in part through the Commodity Credit Corporation (CCC), which provides loans to farmers and acquires crops under price support programs. This paper focuses, therefore, on income fluctuations that are not protected by these programs.

The effect of idiosyncratic income shocks on land sales has been documented in the development literature. Using data on land sales in rural India, Rosenzweig and Wolpin (1985) find that households that experienced two consecutive years of drought were 1.5 times more likely to sell land than other families, though a single year of drought did not significantly increase the probability of a sale. Sarap (1995) analyzes land transactions in a small Indian village over 44 years, and finds that, among small and poor farms, land sales were mostly due to financial distress or used for debt repayment, whereas large farmers sold land to liquidate funds for more varied reasons, including consolidation of land holdings and purchase of equipment. In the US, Bierlen and Featherstone (1998) study machinery investment by large farms in the 1970s and 1980s, and find that during the farm recession of the 1980s financially distressed farmers became borrowing-constrained and were forced to reduce investments; they find no evidence of borrowing constraints in the boom years of the 1970s.

The remainder of the chapter is organized as follows. The theoretical framework is laid out in Section 3.2. Section 3.3 describes the data used in the analysis, and Section 3.4 lays out the empirical methodology. The results are presented in Section 3.5, and Section 3.6 concludes.

3.2 Theoretical Framework

> "Last year we had our crops in so early. Everything looked so nice, and we didn't get any rain – 2.4 inches of rain from May 3rd to the 8th of July. I didn't have very good oats. I didn't have any beans. Just couldn't make up that operating loan. That's why I didn't farm this year. They wouldn't give me any money [...]."
> – former Minnesota farmer, quoted in Rosenblatt, *Farming is in Our Blood*

3.2.1 Setup

In this section, I present a simple model that illustrates one way in which recessions can exacerbate borrowing constraints leading to (potentially inefficient) farmer exit. In the model, borrowing constraints prevent some farmers from borrowing to finance seeds and other up-front expenses at the beginning of the planting season, and can therefore force farmers to exit. The model studies the same basic forces as Bernanke and Gertler (1989), but differs from their model in its focus, which is on the effect of borrowing constraints on the ownership of a scarce resource, namely land.

To keep things simple, consider an economy that lasts only one period. There are four risk-neutral agents, all of whom consume at the end of the period: a poor farmer, two rich farmers and a bank. The poor farmer begins the period with wealth w^P and the rich farmers begin each with wealth $w^R \gg w^P$.

There is a single, indivisible, plot of land which will produce a crop if tended by one of the farmers. If owned by the poor farmer, the land will produce output y^P; if owned by either of the rich farmers, it will produce output y^R. To produce a crop a farmer must incur an up-front cost of seeds and other inputs I. For convenience, normalize the interest rate to zero and ignore discounting. Assume that

$$y^R, y^P > I$$

to ensure that it is profitable for any farmer to farm the land. Since the economy lasts for only one period, a sufficient condition on the rich farmers' wealth that ensures either one could buy the land from the poor farmer and produce without the need to borrow is

$$w^R > y^P. \tag{3.1}$$

Assume also that the poor farmer cannot afford to invest without borrowing, even if he owns the land:

$$w^P < I.$$

The poor farmer's ability to borrow to fund production is constrained by the fact that only a fraction $\lambda < 1$ of output can be pledged to the bank. This condition can be thought of as a reduced-form representation of a standard moral-hazard problem in which the farmer must retain a sufficiently high share of his output in order to align his incentives with his creditor's. Alternatively, it may be a reduced-form representation of a model in which the farmer receives some utility from farming his own land; by its very nature such utility is not pledgible. Because the economy ends at the end of the period, land is worthless at that point and therefore cannot serve as collateral.

Assume for convenience that the economy begins with the poor farmer owning the land (the same logic will carry over to the case in which a rich farmer owns the land). The purpose of having two rich farmers to one poor farmer is to ensure that the poor farmer has all the market power in the market for land, which also provides price determinacy.

The timing of the economy is shown in Figure 3-3. At the beginning of the period, there is a market for land in which the poor farmer may elect to sell the land to either rich farmer. The farmer who owns the land following a transaction (if any) can borrow from the bank, if needed, to finance the up-front cost I. After the farmer has invested in the crop, production/harvest takes place. The farmer then repays the loan, if any; finally, all agents consume and the period ends.

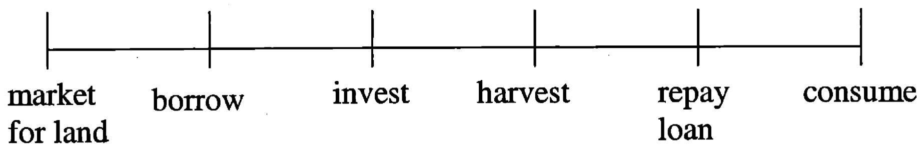

Figure 3-3: Timing in the Model

3.2.2 First Best

In the first-best allocation, the farmer with the highest productivity owns the land. In other words the poor farmer owns the land iff

$$y^P \geq y^R, \tag{3.2}$$

and either one of the rich farmers owns the land otherwise.

3.2.3 Equilibrium

First, suppose that the poor farmer does not sell the land. He will be able to borrow sufficiently to fund production iff

$$\lambda y^P \geq I - w^P.$$

Define the cutoff level of initial wealth needed for the poor farmer to produce to be

$$\overline{w}^P \equiv I - \lambda y^P. \tag{3.3}$$

The poor farmer will be able to fund production iff $w^P \geq \overline{w}^P$; otherwise he will eat only his endowment. His consumption, conditional on retaining ownership of the land, will therefore be

$$c^P\left(w^P; \text{keep}\right) = \begin{cases} w^P - I + y^P & \text{if } \; w^P \geq \overline{w}^P \\ w^P & \text{if } \; w^P < \overline{w}^P \end{cases}.$$

Next, consider the case where the poor farmer sells his land to one of the rich farmers. The poor farmer will then consume

$$c^P\left(w^P;\text{sell}\right) = w^P + p$$

where p is the price of the land. Since they will bid away any rents, the price of the land will be such that either rich farmer is indifferent between owning and not owning it. A rich farmer who owns the land will consume

$$c^R\left(w^R;\text{buy}\right) = w^R - p - I + y^R;$$

a rich farmer who does not own the land will consume

$$c^R\left(w^R;\text{pass}\right) = w^R.$$

Since the rich farmers will bid away any profit from the land, the equilibrium price will equate these values:

$$p^* = y^R - I.$$

We can now determine sale conditions. The poor farmer will sell the land iff this increases his consumption level. If the farmer is able to borrow (i.e., if $w^P \geq \overline{w}^P$), this condition holds iff

$$p > y^P - I$$

which can be rewritten as the first-best condition, Equation (3.2). In other words, if the poor farmer is able to borrow sufficiently to fund production he will sell iff it is socially optimal to do so (that is, if the rich farmers are more productive than the poor farmer).

Alternatively, suppose that the borrowing constraint binds, i.e., $w^P < \overline{w}^P$. In that case the farmer will sell at any positive price. Whether this solution replicates first best depends on the relative productivity of the two farmers. Specifically, as long as $y^R > y^P$ the sale will be efficient, but if the converse holds – if the poor farmer is more efficient – the sale will be welfare-reducing. In this case the borrowing constraint leads to an allocation of land that does not maximize each agent's productivity.

3.2.4 The Effect of Recessions

In light of the above model we can now think about the effect of recessions – referring, loosely, to negative shocks to farmers' net worth.

Revenue shocks, such as those studied in this paper, are likely to have a negative effect on farmers' resources precisely through the mechanism outlined above: since the cost of production is fixed – the fields need to be prepared and sowed, the crops need to be watered, tended, and harvested – negative shocks to either yield or crop prices will drive down farmers' profits and may cause profits to turn negative. A farmer who makes a loss will necessarily have a smaller balance to put towards the next year's investment. Whether the balance suffices or not depends, of course, both on initial wealth and on the size of the shock.

We may assume that the rich farmers are sufficiently rich that a negative shock does not impede their ability to buy and invest in the land (i.e., the condition of Equation (3.1) continues to hold). Nevertheless such a negative shock may well push the poor farmer's wealth from a level above $\overline{w}^P$ to some point below this threshold value, forcing the poor farmer to sell his land regardless of his relative productivity. Thus, by increasing the incidence of binding credit constraints, recessions can exacerbate small (poor) farmer exit.

3.3 Data

3.3.1 Crop Acreage and Prices

I use data on 14 common crops by county over the period 1977-1997 are available. For each crop and county, annual data on acreage and yield are obtained from USDA National Agricultural Statistics Service (NASS) *Crop County Data.* I limit the sample to crops that are grown continuously in a county over the entire period. Approximately half of the counties that have grown a given crop at some point over the period studied have grown it every year over this period; the other half are counties that have moved in and out of production of the crop. In all, 2,335 counties are included in the analysis. National-level price series for each crop are obtained from USDA NASS *Track Records: United States Crop Production.* I use national price series rather than state-level prices, which are available for many crops, because they are more likely to be exogenous to the individual farmers' behavior.

Table 3.1 lists the crops, along with their average annual acreage in the US over the period studied, the average number of counties in which each crop was grown over the period and the number of counties that have grown it continuously, as well as the average real price per unit of production over the period 1972-1998.[2,3]

3.3.2 Census of Agriculture

The Census of Agriculture is conducted (approximately) every five years, and includes information, at the county level, on the number of farms – including by size categories and by ownership. I use data from the 1978, 1982, 1987, 1992 and 1997 Census. Table 3.2 lists summary statistics for the counties in my sample. Although small farms and family-owned farms still dominate most counties – both in numbers and acreage – they have declined considerably over the last 20 years.

3.4 Empirical Methodology

3.4.1 Computing Shocks

Revenue shocks are made up of two components: shocks to crop prices and shocks to crop yields. The two are of course related, since high yields lead to low prices, and vice versa; and while yields have been increasing for all the crops included in this study over the past three decades, prices have fallen sharply over this period.[4] As noted in the Introduction, revenues have been declining as yields have increased, presumably due to relatively inelastic demand, possibly combined with increased production in other parts of the world. Figure 3-4 shows corn yields, measured in bushels per acre, for three counties – Kit Carson County, Colorado,

[2] For some years, crops, and states, the number of acres planted with a crop is not available; I use the number of *harvested* peanut acres instead for those observations. This correction introduces additional measurement error because the number of harvested acres is chosen *ex post*, after the yield is revealed.

[3] Corn and sorghum may be grown for grain or for silage. Since the decision harvest the crop for grain or silage is made *ex post*, after the yield and market price are revealed, I use the total number of acres planted with the crop rather than the number of acres harvested for grain alone, although the price used is the grain price.

[4] The increasing trend in crop yields may be attributable to a number of factors. These include improvements in technology – such as the adoption of disease-resistant crop varieties – that are available to a large number of farmers; the discriminate exit of the least productive land from agriculture; and possibly higher productivity of large corporate farmers *vis-a-vis* small farmers. The latter possibility is explored, and tentatively rejected, in Appendix 3.A.

Table 3.1: Crop Summary Statistics, 1972-1998

	Acres (Thousands)	Producing Counties	Unit of Production	Price per Unit[a]
Barley	9,002	1080	Bushels	2.48
Corn	76,500	2,399	Bushels	2.71
Cotton	12,600	556	100 Lbs.	63.54
Flaxseed	642	95	Bushels	6.68
Oats	11,200	1,527	Bushels	1.59
Peanuts	1,513	237	100 Lbs.	26.36
Potatoes	866	183	100 Lbs.	5.34
Rice	2,797	110	100 Lbs.	9.89
Rye	873	317	Bushels	2.52
Sorghum	13,900	1199	Bushels	2.48
Soybeans	61,400	1,764	Bushels	6.80
Sunflower	2,533	145	100 Lbs.	0.10
Tobacco	798	554	100 Lbs.	159.6
Wheat	71,400	2,336	Bushels	3.69

[a] Average price in constant 1982-1984 dollars

Table 3.2: Changes in Farmland Distribution, 1978-1997

Variable	1978	1997	growth rate
Average number of farms per county	740	628	-15%
Average number of family-owned farms	645	540	-16%
Average number of farms under 500 acres	624	517	-17%
Average number of farms above 2000 acres	20	23	+15%

Source: Author's computation from US Census Bureau, *Census of Agriculture*

Coffee County, Georgia and Adams County, Wisconsin – from 1972 to 1998. The solid line shows Adams County, the dashed line shows Kit Carson County, and the dotted line shows Coffee County. In addition to the increasing trend in yields, the figure also shows permanent differences in yields across counties and large fluctuations in yields from year to year.

I define a revenue shock to crop c in county j and year t as the deviation of log revenue per acre from the county average, after removing a linear trend; in other words, it is the residual from a regression of log revenue per acre on county fixed effects and a common linear time trend:

$$\mathrm{u}_{cjt} \equiv \ln\left(\mathrm{revenue}_{cjt}\right) - \sum_{j} \widehat{\gamma}_{cj}\mathrm{county}_j - \widehat{\theta}_c \mathrm{t} \tag{3.4}$$

where $\mathrm{revenue}_{cjt}$ is revenue per acre planted with crop c in county j in year t, county_j is a county fixed effect and t is a linear time trend,

$$\mathrm{t} = \mathrm{year} - 1977,$$

and the regression is estimated using annual data from 1978-1997. Revenue per acre is computed as

$$\mathrm{revenue}_{cjt} = \mathrm{yield}_{cjt} \cdot \mathrm{p}_{ct}$$

with p_{ct} the US-average per-unit price of crop c in year t in constant 1982-1984 dollars. Fluctuations from the linear trend can be thought of as being largely due to exogenous sources, such as demand and weather shocks.

Finally, crop-specific shocks are aggregated to the county level, weighted by the relative importance of each crop to the county:

$$\mathrm{shock}_{jt} = \sum_{c} \left(\frac{\mathrm{acres}_{cj,C(t)}}{\mathrm{farmland}_{j,C(t)}} \right) \mathrm{u}_{cjt}, \tag{3.5}$$

where $\mathrm{acres}_{cj,C(t)}$ is the number of acres devoted to crop c in county j in the previous Census year (1978, 1982, 1987 or 1992) and $\mathrm{farmland}_{j,C(t)}$ is the total number of farmland acres in county j in the previous Census year. This formulation allows crop-specific shocks that contribute a large fraction of the county's income – because they take up a larger fraction of the county's farmland – to be weighted more heavily than shocks to relatively marginal crops. Figure 3-5 shows the county-level shocks for 1979-1998, computed for the three counties – Kit Carson County, Colorado, Coffee County, Georgia and Adams County, Wisconsin – whose corn yields were shown in Figure 3-4.[5]

To analyze the effect of these shocks on five-year changes in farm numbers, I compute:

$$\begin{aligned}
\text{avgshock}_{jt} &= \frac{1}{5}\sum_{\tau=t-4}^{t} \text{shock}_{j\tau} && (3.6)\\
\text{worst}_{jt} &= \min_{t-4\leq\tau\leq t} \{\text{shock}_{j\tau}\} && (3.7)\\
\text{badyears}_{jt} &= \sum_{\tau=t-4}^{t} \mathbb{I}\left(\text{shock}_{j\tau} < 0\right), && (3.8)
\end{aligned}$$

where $\mathbb{I}(A)$ is an indicator function equal to 1 if A is true and 0 otherwise. In words, avgshock_{jt} is the average of the intercensal-period shocks to county j, badyears_{jt} is the number of these shocks that are negative, and worst_{jt} is the worst shock during the intercensal period.

While these measures are all closely related – and highly correlated – they do capture subtly different aspects of the shocks. The average shock measure captures the effect of the sequence of shocks. If a farmer has some cash reserves (or a loose credit constraint) she may be able to withstand small fluctuations in revenue if mean revenue is sufficiently high; under these circumstances, only if the positive shocks are small in absolute terms relative to the negative shocks over a given period will the farmer's credit constraint begin to bind. On the other hand, the average shock fails to capture the possibility that a single very negative shock may force a farmer into selling his land regardless of the other shocks – as might be the case, for example, if the farmer has very low reserves. This effect will be captured in regressions using the worst shock measure. Finally, the number of negative shocks is intended to capture something between these two cases, though the magnitude of these shocks is lost in the discretization.

[5] Shocks to barley, sorghum and wheat as well as corn revenues per acre are included in the Kit Carson County, CO aggregate shock; in Coffee County, GA, oats, peanuts, rye, soybeans, tobacco and wheat are included along with corn; and in Adams County, WI, the aggregate shock is a composite of revenue shocks to corn, oats, potatoes, and soybeans. The median number of crops for which a complete series exists is 3 per county.

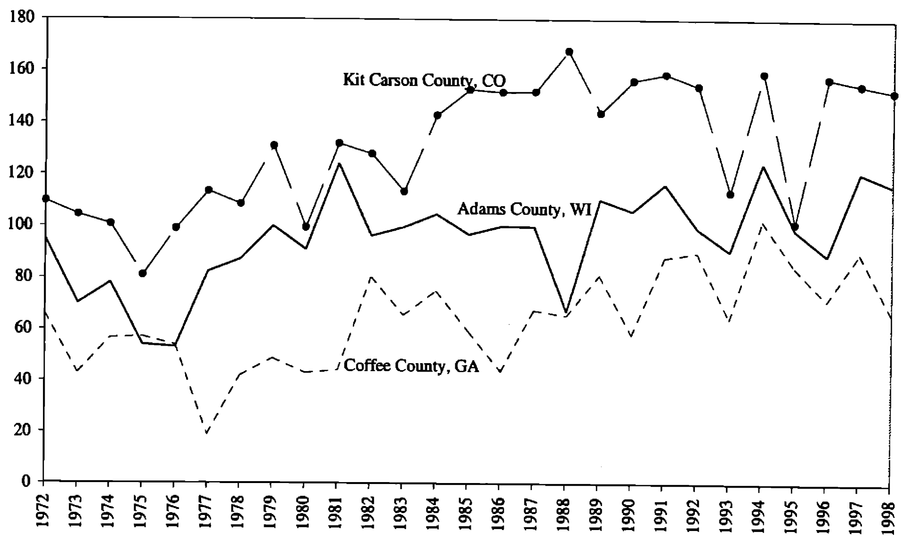

Figure 3-4: Corn Yield per Acre in 3 Counties

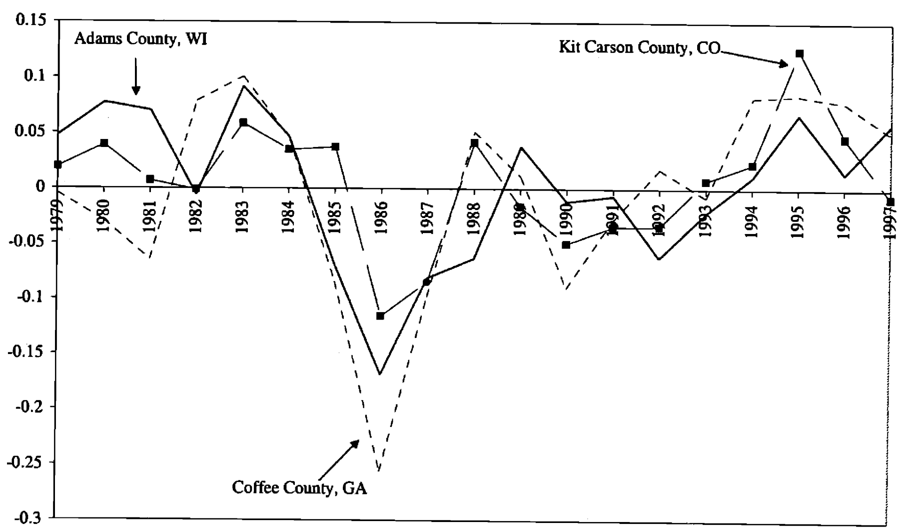

Figure 3-5: County Level Shocks in 3 Counties

3.4.2 Regressions

I regress five-year changes in the number of farms of various types – all farms, family-owned farms, and small farms – separately on each of the above five-year shock measures. For notational convenience, I renormalize the time variable so that the interval between Census dates is of length 1. Focusing on effect of shocks on the total number of farms, I estimate a fixed-effect model,

$$\Delta \ln (\text{farms}_{jt}) = \alpha + \beta \cdot f(\text{shock}_{jt}) + \sum_t \delta_t \text{year}_t + \varepsilon_{jt}, \tag{3.9}$$

where ΔX is the five-year difference in X, $f(\text{shock}_{jt})$ is one of the following: avgshock_{jt}, worst_{jt} or badyears_{jt}, and year_t is a year fixed-effect. To check robustness of the results, I also allow year fixed effects to vary by state and introduce a crop-by-crop linear trend:

$$\begin{aligned}\Delta \ln (\text{farms}_{jt}) = \alpha + \beta \cdot f(\text{shock}_{jt}) + \sum_t \delta_t \text{year}_t \\ + \sum_c \lambda_c \cdot \text{frac}_{cj,1978} \cdot \text{t} + \varepsilon_{jt},\end{aligned} \tag{3.10}$$

where $\text{frac}_{cj,1978}$ is the fraction of farmland acres planted with crop c in county j in 1978,

$$\text{frac}_{cj,1978} \equiv \left(\frac{\text{acres}_{cj,1978}}{\text{farmland}_{j,1978}} \right),$$

t is a linear time trend and u_{jt} is an error term. The crop-specific time trend is included in the regressions to allow for the possibility that farmland in areas specializing in different crops consolidate at different rates – for example because different farm sizes are better suited to different crops, and may become more or less so over time as farming technology changes.[6]

Next, to determine whether the estimated effect is temporary or permanent, I add one lags of the shock to the baseline equations:

$$\Delta \ln (\text{farms}_{jt}) = \alpha + \beta_0 \cdot f(\text{shock}_{jt}) + \beta_1 \cdot f(\text{shock}_{j,t-1}) + \sum_t \delta_t \text{year}_t + \varepsilon_{jt}. \tag{3.11}$$

[6] I estimate a lagged-dependent variable specification in addition to the fixed-effect specification presented here. Since the coefficient on the lagged dependent variable is almost never statistically distinguishable from unity, however, I conclude that the fixed-effect model is the correct specification.

The coefficient β_0 captures the immediate impact of the shock (within a five-year window), and the sum $(\beta_0 + \beta_1)$ captures the permanent effect.

To see whether the number of farms of different types – family-owned, small and large – is affected differently by the shocks, I repeat regressions (3.9)-(3.11) with LHS variables $\Delta \ln(\text{family}_{jt})$, the change in log number of family farms in county j in year t, $\Delta \ln(\text{small}_{jt})$, the change in log of the number of farms with fewer than 500 acres, and $\Delta \ln(\text{large}_{jt})$, the change in log of the number of farms with over 2,000 acres in county j in year t.

Standard errors are clustered at the county level to allow for arbitrary autocorrelation of errors in all regressions.

3.4.3 Endogeneity and Measurement Error

Before turning to the results, I flag a couple of potential pitfalls.

First, one may be concerned that deviations of revenues from trend are not entirely exogenous and may in fact vary with ownership changes, in which case it is possible that estimates from these regressions are spurious. One mechanism through which such endogeneity could emerge is crop yields. If yields increase after sale (because the average seller is less efficient than the average buyer, or because the least-productive land is diverted to non-agricultural uses), estimates of β will be downward biased in absolute value; if yields decrease after sale (for example because the seller has "specific experience" with respect to the land plots being sold (Rosenzweig and Wolpin 1995), or because of agency problems), estimates of β will be biased upwards in absolute value.[7] Alternatively, if large farms behave less as price-takers and more as price-setters, prices may be partly endogenous to ownership changes. These endogeneity concerns can probably be addressed using instrumental-variables specifications in which revenue fluctuations are identified using weather shocks.

[7]Several empirical studies of farm productivity in developing countries have found an inverse relationship between farm size and productivity, though more recent evidence suggests omitted-variable bias – specifically, the omission of data on land quality – may be the cause of this "stylized fact" (see Bhalla 1988, Bhalla and Roy 1988 and Benjamin 1995 for a review). In the U.S., anecdotal evidence suggests that corporate farms are aware of the agency problem and have attempted to reduce it with appropriate incentive programs (Haw 1987). I attempt to find a relationship between crop yields and farm size in the U.S. using the available data. While the estimated regressions are somewhat problematic since yields and ownership structure and determined simultaneously, the results, shown in Appendix 3.A, suggest that crop yields are largely uncorrelated with the size of farms and their ownership structure.

A related issue is measurement error. Due to the nature of aggregate data, the average revenue shock in a county in a given year will always be smaller than the shock to the worst-affected farmer; it is the impact of the latter shock – which may induce a farmer to exit – that we wish to study here. Shocks to average revenue per acre in the county fails to capture such variations and in that sense are measured with error. This problem is further compounded by the time-series aggregation of these shocks over 5 years. Variation in crop yields within a county may in fact induce bias if susceptibility to weather shocks varies systematically with farm structure – or if, for example, large farms tend to occupy the best land in the county. Furthermore, if different types of farms tend to specialize in different crops – large farms specializing in growing cotton, for example, while small farms specialize in tobacco – then fluctuations in crop-specific revenues, which are treated equally here, can be expected to have differential effects on ownership structure. I leave these concerns to be addressed in future work.

3.5 Results

3.5.1 Baseline Specification

I first present results from regressions without crop-specific time trends and without lagged shocks. Tables 3.3-3.6 show, respectively, OLS estimates of the effect of revenue shocks on intercensal changes in the total number of farms, the number of family-owned farms, the number of small farms (with fewer than 500 acres) and the number of large farms (above 2,000 acres) from Equation (3.9). In each table, the first three columns show results from regressions in which year fixed effects are constrained to be equal across all counties, and the next three columns show results from regressions allowing state-specific year fixed effects. Coefficients are all multiplied by 100 and can therefore be interpreted as percentages. Standard errors are clustered at the county level to allow for arbitrary autocorrelation of the error term.

For all farms, family-owned farms, and small farms, the results tell a fairly consistent story. The strongest and most consistently significant results appear when the shock measure used is the worst shock over the previous five years. Comparing the effect of a worst shock in the 25th percentile of worst shock with the effect of a worst shock in the 75th percentile of worst shocks – about -0.13 and -0.007, respectively – the estimates from Table 3.3, Columns (2) and

Table 3.3: Regression Results, Log Number of Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	11.629			-3.150		
	(2.238)			(2.633)		
worst_{jt}		23.974			17.765	
		(0.986)			(1.043)	
badyears_{jt}			-0.530			-0.430
			(0.098)			(0.103)
Year FE	Common	Common	Common	State	State	State

Notes: Coefficients are multiplied by 100. Standard errors are clustered at the county level.

Table 3.4: Regression Results, Log Number of Family Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	8.855			-4.949		
	(2.460)			(2.791)		
worst_{jt}		25.907			20.668	
		(1.078)			(1.126)	
badyears_{jt}			-0.397			-0.312
			(0.106)			(0.112)
Year FE	Common	Common	Common	State	State	State

Notes: Coefficients are multiplied by 100. Standard errors are clustered at the county level.

Table 3.5: Regression Results, Log Number of Small Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	16.349			-2.969		
	(2.944)			(3.595)		
worst_{jt}		33.332			25.893	
		(1.270)			(1.412)	
badyears_{jt}			-0.528			-0.310
			(0.138)			(0.150)
Year FE	Common	Common	Common	State	State	State

Notes: Coefficients are multiplied by 100. Standard errors are clustered at the county level.

(5), suggest that the number of farms declines by 2-3% faster in the case of the former (larger) shock than in the latter. In the median county with 570 farms, this amounts to a decline of approximately 10-15 farms. The difference in the rates of decline of small farms (Table 3.5) are larger – there is an estimated 3-4% difference in decline rates of small farms between counties experiencing shocks at the 25th- and 75th-percentile levels.

One result worth noting is that the effect of the average shock disappears in all specifications when year fixed effects are allowed to vary by state. The likeliest explanation for this result is that the shocks are sufficiently uncorrelated over time that the average shock is not a good proxy for the extent of financial distress facing a farm; the magnitude of the worst shock, and the number of negative shocks, are likely to be better measures of financial distress over a five-year period.

Unlike the estimated effects shown in Tables 3.3-3.5, the estimated effects of revenue shocks on large farms, shown in Table 3.6, are not internally consistent and raise serious concerns about the validity of the shock measures used. First, note that the sign of the effect of a negative shock is unstable across specifications. Column (1) shows a negative average shock decreases (insignificantly) the number of large farms, whereas Column (2) shows that a negative worst shock significantly increases the number of large farms. Furthermore, when year fixed effects are allowed to vary by state, both effects are large and significant (Columns (4) and (5)).

One possible cause of the strange and internally-inconsistent estimation results for large farms is the distribution of large farms. Unlike the numbers of small, family-owned, and all farms, the distribution of large farms is highly asymmetric, and has a large number of very small numbers: nearly 15% of the counties in the study have no large farms in 1978, and an additional 12.5% have only one large farm, while the mean number of large farms is 20; 9% of sample counties have no large farms in 1997. To test for the possibility that the above results are an artifact of skewed distribution of large farms, I reestimate the regressions using the raw number of large farms rather than its natural logarithm. The results are shown in Table 3.7. Since the results are in raw numbers, the coefficients are not multiplied by 100.

The results are slightly more consistent than those estimated using log number of large farms. The estimated coefficients on the average shock is only significantly different from zero in the specification with common year fixed effects (Column (1)) but not in the more flexible

Table 3.6: Regression Results, Log Number of Large Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	5.720			51.417		
	(11.980)			(14.505)		
worst_{jt}		-54.030			-30.701	
		(4.608)			(5.365)	
badyears_{jt}			0.797			-0.515
			(0.517)			(0.544)
Year FE	Common	Common	Common	State	State	State

Notes: Coefficients are multiplied by 100. Standard errors are clustered at the county level.

Table 3.7: Regression Results, Number of Large Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	3.220			3.475		
	(1.474)			(1.954)		
worst_{jt}		-3.758			-3.112	
		(0.519)			(0.656)	
badyears_{jt}			-0.038			-0.093
			(0.060)			(0.071)
Year FE	Common	Common	Common	State	State	State

Notes: Standard errors are clustered at the county level.

specification which allows year fixed effects to vary by state (Column (4)). The coefficients on the number of negative shocks (Columns (3) and (6)) are never significantly different from zero. The coefficient on the worst shock, however, suggests that the number of large farms increases in response to a negative shock. Using the same counter-factual exercise as above, comparing the effect of a 25th-percentile worst shock with the effect of a 75th-percentile worst shock, we find that the number of large farms increases by 0.4-0.5 more (in absolute numbers) in the case of the more adverse shock than in the case of the weaker shock.

I next turn to an alternative specification, allowing farmland consolidation to occur at different rates for counties specializing in different crops – to see whether controlling for differential rates of consolidation improves the consistency of the estimates.

3.5.2 Crop-Specific Time Trends

Farmland consolidation may occur at differential rates in areas that specialize in different crops – for example because different farm sizes are better suited to different crops, and may become more or less so over time as farming technology changes. If they do, and if these rates differ in a way that is correlated with the volatility of revenue of these crops (but is not directly due to the volatility of revenue) – then the estimates obtained in the previous section may be biased. I therefore estimate Equation (3.10), which includes a linear time trend in the fraction of land devoted to each of the 14 crops included in this study.

The results from these specifications, for both the restricted regressions in which year fixed effects are constrained to be constant for all counties and for unconstrained regressions in which year fixed effects are allowed to vary by state, are presented in Tables 3.8-3.11.

For all farms, family-owned farms and small farms, the results are similar to the ones estimated using the baseline equations, though somewhat smaller. Looking at Table 3.8, the difference between the rate of decline in the total number of farms in a county experiencing one negative shock in five years (the 25th percentile) relative to a county experiencing three negative shocks in the same period (the 75th percentile of the number of negative shocks) is estimated at a 1.1-1.3%. For the median county this amounts to 6-7 more farms closing in a five-year period with three negative shocks as opposed to a five-year period with only one negative shock. The estimated effect on family-owned and small farms is very similar (Tables

Table 3.8: Regression Results with Time Trends, Log Number of Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	12.103			0.720		
	(2.278)			(2.596)		
worst_{jt}		13.265			9.228	
		(1.152)			(1.233)	
badyears_{jt}			-0.649			-0.538
			(0.095)			(0.101)
Year FE	Common	Common	Common	State	State	State

Notes: Coefficients are multiplied by 100. Standard errors are clustered at the county level.

Table 3.9: Regression Results with Time Trends, Log Number of Family Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	9.489			-0.797		
	(2.509)			(2.758)		
worst_{jt}		14.571			10.767	
		(1.266)			(1.326)	
badyears_{jt}			-0.490			-0.395
			(0.104)			(0.110)
Year FE	Common	Common	Common	State	State	State

Notes: Coefficients are multiplied by 100. Standard errors are clustered at the county level.

Table 3.10: Regression Results with Time Trends, Log Number of Small Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	17.034			3.483		
	(2.983)			(3.515)		
worst_{jt}		18.891			14.787	
		(1.505)			(1.662)	
badyears_{jt}			-0.673			-0.471
			(0.136)			(0.148)
Year FE	Common	Common	Common	State	State	State

Notes: Coefficients are multiplied by 100. Standard errors are clustered at the county level.

3.9 and 3.10).

There is no consistent effect of negative shocks on the number of large farms (Table 3.11). Changing the LHS variable from log number of large farms to the (raw) number of large farms does not substantively alter these results.

3.5.3 Short-Run vs Long-Run Effects

Shocks may have temporary effects that disappear in the long run, or, more generally, may have different short- and long-run effects. To disentangle the immediate effect of a shock from its long-run effect I estimate Equation (3.11), which allows for separate short-run and long-run effects of a shock. In each regression, the coefficient on the shock captures its immediate effect, while the coefficient on the lagged shock captures additional or counter-balancing long-run effects. The results presented in this section are from regression that include crop-specific time trends, as above.

The results are presented in Tables 3.12-3.15. As before, the first three columns in each table show results from regressions in which year fixed effects are held constant across counties, whereas Columns (4)-(6) allow year fixed effects to differ by state.

Across specifications, several clear results emerge. First, the effect of shocks on the total number of farms, as well as the effect on family-owned farms and small farms, remains strong and, in all but one or two regressions in which the estimate is only marginally significant, the effect is estimated to be very significant in both the short- and longer-runs. The coefficient on the lagged shock is sometimes positive and sometimes negative, but almost always substantially smaller than the coefficient on the main effect, suggesting that most of the effect is instantaneous (within a 5-year window) and permanent. The results for large farms continue to be suspicious; the signs on several of the coefficients change when the change in the number of large farms is used on the LHS, rather than the rate of change.

3.6 Conclusion

Uncertain agricultural revenues – driven by idiosyncratic shocks to crop yields as well as shocks to aggregate yields, which affect crop prices – are a source of risk to credit-constrained farmers.

Table 3.11: Regression Results with Time Trends, Log Number of Large Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	4.501			32.576		
	(12.010)			(14.333)		
worst_{jt}		4.659			8.415	
		(5.281)			(6.482)	
badyears_{jt}			0.792			-0.237
			(0.513)			(0.539)
Year FE	Common	Common	Common	State	State	State

Notes: Coefficients are multiplied by 100. Standard errors are clustered at the county level.

Table 3.12: Regression Results with Lagged Shocks and Time Trends, Log Number of Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	32.715			28.083		
	(2.638)			(2.965)		
$\text{avgshock}_{j,t-1}$	2.134			7.168		
	(2.522)			(3.040)		
worst_{jt}		13.422			8.522	
		(1.233)			(1.335)	
$\text{worst}_{j,t-1}$		-1.645			0.476	
		(1.387)			(1.473)	
badyears_{jt}			-0.842			-0.751
			(0.104)			(0.106)
$\text{badyears}_{j,t-1}$			-0.179			-0.371
			(0.110)			(0.111)
Year FE	Common	Common	Common	State	State	State

Notes: Coefficients are multiplied by 100. Standard errors are clustered at the county level.

Table 3.13: Regression Results with Lagged Shocks and Time Trends, Log Number of Family Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	32.450			29.983		
	(2.868)			(3.151)		
$\text{avgshock}_{j,t-1}$	3.215			7.738		
	(2.848)			(3.528)		
worst_{jt}		15.507			11.127	
		(1.131)			(1.452)	
$\text{worst}_{j,t-1}$		4.943			2.372	
		(1.660)			(1.799)	
badyears_{jt}			-0.724			-0.659
			(0.114)			(0.117)
$\text{badyears}_{j,t-1}$			-0.201			-0.449
			(0.117)			(0.123)
Year FE	Common	Common	Common	State	State	State

Notes: Coefficients are multiplied by 100. Standard errors are clustered at the county level.

Table 3.14: Regression Results with Lagged Shocks and Time Trends, Log Number of Small Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	44.395			42.933		
	(3.429)			(4.020)		
$\text{avgshock}_{j,t-1}$	-11.290			-1.008		
	(3.446)			(4.225)		
worst_{jt}		20.121			15.505	
		(1.611)			(1.814)	
$\text{worst}_{j,t-1}$		-6.713			-1.736	
		(1.884)			(2.075)	
badyears_{jt}			-1.006			-0.841
			(0.146)			(0.154)
$\text{badyears}_{j,t-1}$			-0.091			-0.436
			(0.146)			(0.158)
Year FE	Common	Common	Common	State	State	State

Notes: Coefficients are multiplied by 100. Standard errors are clustered at the county level.

In extreme cases, fluctuations in revenue may force farmers into bankruptcy or into voluntary sale of their farms to agricultural firms or other buyers. In this paper I study the extent to which temporary fluctuations in agricultural revenue can cause – or expedite – exit of small and family farmers and concurrent consolidation in the agricultural sector.

I find that deviations from long-run trends in crop revenues do induce farm exit, and that to a large extent these effects are permanent; this effect is strongest for small and family-owned farms. The effect on large farms is indeterminate. It may be that in some areas, the number of large farms increases because a prosperous medium-sized farm buys one or more small farms in bad times, putting it over the threshold of 2,000 acres, while in other areas large farms merge with other large farms during bad times, reducing the number of large farms in the county (though not necessarily their acreage). I leave the reasons for these varies outcomes for future study. Another possibility is that the results presented here are an artifact of the shock specification; closer attention needs to be paid to the definition of revenue shocks in future work. Ideally, an exogenous variable, such as weather, can be used to instrument for revenue shocks.

The tentative finding presented here may be interpreted to imply that credit constraints prevent farmers from holding on to their farms even though the farms may be viable long-run concerns. An alternative interpretation of the same result is that farmers' farming skill is revealed in lean years, in which case such negative shocks may increase long-run production efficiency. Which interpretation is adopted has critical implication for policy recommendations. I leave this question for future research.

Appendix to Chapter 3: Productivity and Size

If large farms are more efficient than small farms – or if corporate-owned farms are more efficient than family-owned farms – can we find evidence of this in the data? Efficiency may be measured along multiple dimensions. More efficient managers may have lower costs, or they may better allocate land to the crops for which their comparative advantage is greatest, or they may produce higher yields holding constant the distribution of crops. In this section, I test the latter hypothesis, asking whether crop yields change as land ownership changes in a county. I do not find any effect of the distribution of farmland across different types of owners on crop yields.

I first compute the fraction of acres in county j that, in year t, belong to family farms (fracfamily), to small farms (fracsmall), and to large farms (fraclarge), where small and large farms are as defined in the text (under 500 acres and over 2000 acres, respectively):

$$\begin{aligned}
\text{fracfamily}_{jt} &= \frac{\text{famacres}_{jt}}{\text{farmland}_{jt}} \\
\text{fracsmall}_{jt} &= \frac{\text{smallacres}_{jt}}{\text{farmland}_{jt}} \\
\text{fraclarge}_{jt} &= \frac{\text{largeacres}_{jt}}{\text{farmland}_{jt}}.
\end{aligned}$$

I include these variables in a simple regression of the yield of crop c in county j in year t on year and county fixed effects:

$$\text{yield}_{cjt} = \alpha_c + \beta_c \text{fracfamily}_{jt} + \sum_j \gamma_{cj} \text{county}_j + \sum_t \delta_{ct} \text{year}_t + \text{u}_{cjt} \tag{3.12}$$

$$\text{yield}_{cjt} = \alpha_c + \beta_c \text{fracsmall}_{jt} + \sum_j \gamma_{cj} \text{county}_j + \sum_t \delta_{ct} \text{year}_t + \text{u}_{cjt} \tag{3.13}$$

$$\text{yield}_{cjt} = \alpha_c + \beta_c \text{fraclarge}_{jt} + \sum_j \gamma_{cj} \text{county}_j + \sum_t \delta_{ct} \text{year}_t + \text{u}_{cjt}. \tag{3.14}$$

Care needs to be used when interpreting results from these regressions. If, as the main part of this paper argues, the size distribution of farms is partly due to innovations in yields – specifically, there is a decline in the number of small and family-owned farms following negative yield innovations, and an increase in the number of large farms – a spurious negative relationship between size and yields will be estimated in these regressions. I nevertheless present the results as a point of reference.

Table 3.16 shows the results. Each cell in the table represents a separate regression, and shows the coefficient β_c on the fraction of farmland acres owned by family farms (Column 1), the fraction owned by small farms (Column 2), and the fraction owned by large farms (Column 3). Year fixed effects are included to capture both global weather conditions and trends in improved yields. County fixed effects are included to capture permanent differences in soil quality and climate. The family-farm regressions use data from 1978, 1982, 1987, 1992 and 1997; because county-level data on the acreage of small and large farms are not available for the 1978 and 1982 Census years, only the last three years (1987, 1992, 1997) are used in the regressions using small farms and large farms. All regressions include year and county fixed effects; standard errors are clustered at the county level.

The results are rarely significant, and certainly not consistent in any one direction. Family farms appear to have lower per-acre yields of tobacco, but not of any other crops, than corporate-owned farms. Small farms have higher yields of corn, rye, and soybeans than the average farm, and lower oat yields. Large farms have marginally-significant (p-value 0.058) lower yields of sorghum compared with the average farm, but are otherwise indistinguishable. As noted above, the coefficients on the fraction of large farms may be downward biased while the coefficients on the fractions of small and family-owned firms are biased upwards. This bias may well explain the significance of the coefficients on the fraction of small farms in several of the regressions.

References

Benjamin, Dwayne (1995). "Can Unobserved Land Quality Explain the Inverse Productivity Relationship?" *Journal of Development Economics* 46.

Bentley, Susan E., et al. (1989). "Involuntary Exits from Farming: Evidence from Four Studies." US Department of Agriculture Economic Research Service Agricultural Economic Report 635.

Bernanke, Ben and Mark Gertler (1989). "Agency Costs, Net Worth, and Business Fluctuations." *American Economic Review* 79.

Bierlen, Ralph and Allen M. Featherstone (1998). "Fundamental q, Cash Flow, and Investment: Evidence from Farm Panel Data." *Review of Economics and Statistics* 80.

Bhalla, Surjit S. (1988). "Does Land Quality Matter? Theory and Measurement." *Journal of Development Economics* 29.

Bhalla, Surjit S. and Prannoy Roy (1988). "Mis-specification in Farm Productivity Analysis: The Role of Land Quality." *Oxford Economic Papers* 40.

Black, John D. and R. H. Allen (1937). "The Growth of Farm Tenancy in the United States." *Quarterly Journal of Economics* 51.

Goodwin, Barry K. and Ashok K. Mishra (2000). "An Analysis of Risk Premia in U.S. Farm-Level Interest Rates." *Agricultural Finance Review* 60.

Haw, C. L. William (1987). "We Who Are Corporate Farmers." In Comstock, Gary, ed., *Is There a Moral Obligation to Save the Family Farm?* Ames, Iowa: Iowa State University Press.

Roberts, Michael J. and Nigel Key (2002). "Risk and Structural Change in Agriculture: How Income Shocks Influence Farm Size." Mimeograph, US Department of Agriculture.

Rosenblatt, Paul C. (1990). *Farming is in Our Blood: Farm Families in Economic Crisis.* Ames, Iowa: Iowa State University Press.

Rosenzweig, Mark R. and Kenneth I. Wolpin (1985). "Specific Experience, Household Structure, and Intergenerational Transfers: Farm Family Land and Labor Arrangements in Developing Countries." *Quarterly Journal of Economics* 100.

Sarap, Kailash (1995). "Land Sale Transactions in an Indian Village: Theories and Evidence." *Indian Economic Review* 30.

US Department of Agriculture (various years). *Agricultural Statistics.* Washington, DC: United States Government Printing Office.

US Department of Agriculture (2002a). *Crop County Data.* `http://usda.mannlib.cornell.edu/data-sets/crops/9X100/`

US Department of Agriculture (2002b). *Farms and Land in Farms.* `http://usda.mannlib.cornell.edu/reports/nassr/other/zfl-bb/fmno0202.pdf`

US Department of Agriculture (2002c). *Track Records: United States Crop Production.* `http://usda.mannlib.cornell.edu/data-sets/crops/96120/`

Table 3.15: Regression Results with Lagged Shocks and Time Trends, Log Number of Large Farms

	(1)	(2)	(3)	(4)	(5)	(6)
avgshock_{jt}	-47.375			-36.638		
	(13.400)			(16.866)		
$\text{avgshock}_{j,t-1}$	-46.026			-52.692		
	(15.883)			(19.646)		
worst_{jt}		-1.671			-2.134	
		(5.655)			(7.312)	
$\text{worst}_{j,t-1}$		-0.522			-6.461	
		(7.240)			(8.698)	
badyears_{jt}			2.029			1.034
			(0.594)			(0.640)
$\text{badyears}_{j,t-1}$			0.127			-0.715
			(0.557)			(0.620)
Year FE	Common	Common	Common	State	State	State

Notes: Coefficients are multiplied by 100. Standard errors are clustered at the county level.

Table 3.16: Yield Regression Results

	Fraction of Acres in		
	Family Farms	Small Farms	Large Farms
Barley	-4.563	-4.442	-1.672
	(4.600)	(10.040)	(4.915)
Corn	-5.899	16.867	-6.725
	(6.525)	(7.312)	(7.350)
Cotton	-85.192	-71.617	-8.447
	(59.714)	(166.41)	(67.825)
Flaxseed	4.008	32.813	27.785
	(13.578)	(123.33)	(22.825)
Oats	-2.510	-12.301	-0.273
	(4.683)	(6.244)	(5.324)
Peanuts	-295.81	693.24	-182.41
	(341.86)	(576.93)	(254.76)
Potatoes, Irish	14.609	45.784	26.862
	(45.690)	(62.038)	(38.424)
Rice	-303.24	2,634.2	520.63
	(528.89)	(1,952.4)	(853.71)
Rye	1.986	17.238	-2.878
	(3.380)	(7.754)	(5.825)
Sorghum	4.550	10.629	-12.472
	(5.011)	(11.655)	(6.568)
Soybeans	-2.342	4.322	0.726
	(1.690)	(2.146)	(2.121)
Sunflower	186.00	989.58	-558.32
	(456.76)	(2,933.6)	(1,007.6)
Tobacco	-524.86	22.986	246.34
	(207.40)	(202.21)	(214.32)
Wheat	-1.335	1.528	-3.572
	(2.477)	(4.097)	(3.296)