MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

JUN 0 8 2009

LIBRARIES

ARCHIVES

Essays on the Development of the American Economy

by

Richard A. Hornbeck
B.A. Economics
University of Chicago, 2004

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

June 2009

© 2009 Richard A. Hornbeck. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly
paper and electronic copies of this thesis document in whole or in part
in any medium now known or hereafter created.

Signature of Author...
Department of Economics
May 15, 2009

Certified by..
Michael Greenstone
3M Professor of Environmental Economics
Thesis Supervisor

Certified by.... ....
Esther Duflo
Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics
Thesis Supervisor

Certified by. ....
Daron Acemoglu
Charles P. Kindleberger Professor of Applied Economics
Thesis Supervisor

Accepted by ....
Esther Duflo
Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics
Chairman, Departmental Committee on Graduate Studies

Essays on the Development of the American Economy

by

Richard A. Hornbeck

Submitted to the Department of Economics on May 15, 2009
in partial fulfillment of the requirements for the Degree of
Doctor of Philosophy in Economics

ABSTRACT

The first essay analyzes the impact of the 1930's American Dust Bowl and investigates how much the short-term costs from erosion were mitigated by long-term adjustments. Exploiting new data collected to identify low, medium, and high erosion counties, estimates indicate that the Dust Bowl led to substantial immediate decreases in agricultural land values and revenues. Until at least the 1950's, however, there was limited reallocation of farmland away from activities that became relatively less productive due to erosion. Relative changes in agricultural land values and revenues indicate that the annualized long-term cost was 86% of the short-term cost to agriculture. Substantial out-migration reflects the large cost of the Dust Bowl, and was an important channel through which short-term costs were partly mitigated.

The second essay examines the impact on agricultural development from the introduction of barbed wire fencing to the American Plains in the late 19th century. Farmers were required to construct fences to be entitled to compensation for damage by others' livestock. From 1880 to 1900, the introduction and universal adoption of barbed wire greatly reduced the cost of fences, relative to predominant wooden fences, most in counties with the least woodland. Over that period, counties with the least woodland experienced substantial relative increases in settlement, land improvement, land values, and the productivity and production share of crops most in need of protection. This increase in agricultural development appears partly to reflect farmers' increased ability to protect their land from encroachment. States' inability to protect this full bundle of property rights on the frontier, beyond providing formal land titles, might have otherwise restricted agricultural development.

The third essay quantifies agglomeration spillovers by comparing the growth of total factor productivity (TFP) among incumbent plants in "winning" counties that attracted a large manufacturing plant to "losing" counties that were the plant's second choice. Five years after the opening, incumbent plants' TFP is 12% higher in winning counties. This effect is larger for plants with similar labor and technology pools as the new plant. We find evidence of increased wages in winning counties, indicating that profits increase by less than productivity.

Thesis Supervisor: Michael Greenstone

Title: 3M Professor of Environmental Economics

Quantifying Long-term Adjustment to Environmental Change: Evidence from the American Dust Bowl

Richard A. Hornbeck*

Abstract

The first essay analyzes the impact of the 1930's American Dust Bowl and investigates how much the short-term costs from erosion were mitigated by long-term adjustments. Exploiting new data collected to identify low, medium, and high erosion counties, estimates indicate that the Dust Bowl led to substantial immediate decreases in agricultural land values and revenues. Until at least the 1950's, however, there was limited reallocation of farmland away from activities that became relatively less productive due to erosion. Relative changes in agricultural land values and revenues indicate that the annualized long-term cost was 86% of the short-term cost to agriculture. Substantial out-migration reflects the large cost of the Dust Bowl, and was an important channel through which short-term costs were partly mitigated.

---

*I thank my advisors Michael Greenstone, Esther Duflo, Daron Acemoglu, and Peter Temin. I also thank David Autor, Abhijit Banerjee, Geoff Cunfer, Joe Doyle, Greg Fischer, Price Fishback, Claudia Goldin, Raymond Guiteras, Gary Libecap, Bob Margo, Ben Olken, Paul Rhode, Claudia Rei, Cate Rosenthal, Tavneet Suri, Rob Townsend, John Wallis, and Noam Yuchtman for their comments, suggestions, and/or generosity with data. I am grateful to Lisa Sweeney, Daniel Sheehan, and the GIS Lab at MIT for their assistance with GIS software; as well as to Christopher Compean, Lillian Fine, Phoebe Holtzman, Paul Nikandrou, and Praveen Rathinavelu for their valuable research assistance. For supporting research expenses, I thank the MIT Schultz Fund, the World Economy Lab, and UROP program. hornbeck@mit.edu

One recurrent theme within economics is that production adjusts more to a shock in the "long-run" than in the "short-run" (e.g., Samuelson's Le Chatelier principle). The short-run costs from a shock may be mitigated in the long-run by economic agents' adjustments. This is a central issue in the context of global climate change and local environmental collapse (Diamond 2005). While short-run costs from weather fluctuations may be large, the implied effect for a permanent change in climate may be an overestimate if agents are able to adjust (Deschenes and Greenstone 2007, Dell et al. 2008, Guiteras 2008).

Similar issues motivate work by Blanchard and Katz (1992) and others on regional economic adjustment to labor demand shocks and the resulting geographic allocation of labor. This relates to a general question of how quickly resources can be reallocated in response to changes in prices or relative productivity. While such responses may be slow and difficult to observe in response to modern shocks, historical shocks provide settings in which to identify economic adjustments that occur over a long period of time (Carrington 1996, Margo 1997, Duflo 2004).

Empirical challenges arise in isolating the effect of an historical shock from other economic changes. The shock must be large and observable to people at the time, but to allow identification of its effects it must also be differential across areas. Large geographically-differentiated shocks are rare, and there is typically not detailed data on what happened in the immediate aftermath of the shock. This makes it difficult to study the specific adjustment strategies that might have been adopted in reaction to the shock. This data scarcity explains why the historical literature has generally focused on long-run changes in population (Davis and Weinstein 2002, Miguel and Roland 2006, and Redding and Sturm forthcoming).

This paper analyzes the aftermath of tremendous permanent soil erosion during the 1930's American Dust Bowl. The Dust Bowl erosion was quite damaging, and its effects varied across counties within a state. Detailed data allow for an examination of many dimensions of adjustment over more than 60 years. This paper focuses on the speed and magnitude of adjustment in the agricultural sector, the resulting difference in short-run and long-run costs, and the reallocation of labor.

The Dust Bowl was a period of extreme soil erosion on the American Plains, unexpectedly brought about by the combination of severe drought and intensive land-use. Strong winds swept topsoil from the land in massive dust storms, and occasional heavy rains carved deep gullies in the land. By the 1940's, many Plains areas had cumulatively lost more than 75% of their original topsoil.

The paper begins by setting out a model of agricultural production in the short-run and long-run. The model highlights that the cost of the Dust Bowl will be immediately capitalized into land values. In the long-run, land allocations adjust toward less-affected activities. In general equilibrium, adjustment may also occur through out-migration and an expansion of local industry. Thus, there are multiple potential margins of adjustment, and this paper examines many of them.

The key to this analysis is the existence of substantial local variation in erosion following the Dust Bowl. In particular, the empirical analysis is based on a comparison of areas that became severely, moderately, or lightly eroded. Erosion levels are combined with Census data to form a balanced panel of counties from 1910 to the 1990's. The regressions compare changes in more-eroded and less-eroded counties within the same state, after adjusting for differences in pre-Dust Bowl characteristics.

The Dust Bowl imposed substantial costs on Plains agriculture. From 1930 to 1940, immediately following the Dust Bowl, more-eroded counties experienced substantial and permanent relative declines in agricultural land values. The per-acre value of farmland declined by 28% in high erosion counties and 16% in medium erosion counties, relative to changes in low erosion counties. Assuming that the low erosion counties were not affected, the decline in land values indicates a total capitalized economic loss from the Dust Bowl of $153 million in 1930 dollars ($1.9 billion in 2007 dollars). Given average land values in 1930, this is the total value of farmland equal to an area roughly the size of Oklahoma.

More-eroded counties also experienced substantial immediate declines in agricultural revenues per-acre of farmland. Under strong theoretical assumptions, the relative percent changes in agricultural land values and revenues imply that the annualized long-term cost, in present discounted value, was 86% of the short-term cost to agriculture. That is, for each 1% that revenue declined from 1930 to 1940, land values declined by 0.86%. Similar levels of cost persistence are implied by the estimated path of recovery in revenues.

One potential margin of adjustment to the Dust Bowl was in land-use decisions, but there was limited and slow adjustment. Total farming experienced little decline, consistent with an inelastic demand for land in non-agricultural sectors. Within the agricultural sector, adjustment appears to have been possible: there was a decrease in the relative productivity of land for crops (compared to animals) and for wheat (compared to hay). Until at least the 1950's, however, there was little adjustment of land from crops to pasture or from wheat to hay.

The paper then examines why land-use adjustment was slow. Credit constraints may have been important, and adjustment was greater in the more-eroded areas that had more banks prior to the 1930's, compared to those with few banks. Evidence is mixed on the role of tenant farming in slowing land-use adjustment. Government programs do not appear to have discouraged adjustment, at least not in more-eroded counties relative to less-eroded counties, as payments were not targeted toward more-eroded counties.

A second potential margin of adjustment was labor, through migration or movement between sectors. Consistent with contemporary and historical accounts, there was a substantial population decrease from 1930 to 1940 in more-eroded counties relative to less-eroded counties, at the same time as substantial out-migration from the entire region. By 1940, unemployment was slightly higher and proxies for wages were lower. Equilibrium was reestablished through further declines in population, rather than an increase in local industry.

One caveat of examining within-region variation from a shock is the inability to observe aggregate adjustment through technological innovation. Instead, this paper analyzes changes in technology choice, treating the technological possibility frontier as common to all areas.

Overall, this paper finds that the Dust Bowl imposed substantial costs on Plains agriculture. Immediate adjustment took place mainly through decreased land values and population, while agricultural land-use only adjusted slowly. There was limited mitigation of short-run agricultural losses. Migration reflected these large losses, and was an important channel through which economy-wide consequences were mitigated, both in the short-run and long-run.

The paper is organized as follows. Section I provides historical background on the Dust Bowl and aggregate trends over the 20th century. Section II models the short-run and long-run response of agricultural production to the Dust Bowl, and discusses how population, industry, and other locations might also change. Section III describes the measurement of erosion levels and presents baseline summary statistics by erosion level. Section IV outlines the empirical methodology, Section V presents the results, and Section VI concludes.

I Historical Background

I.A The Dust Bowl and Massive Soil Erosion: 1931 to 1938

Since the late 19th century, agricultural activity had increased substantially on the Plains. World War I then increased demand for American agricultural goods, leading to high prices and a boom period in Plains agriculture. Farming continued to expand and European production recovered after WW1, contributing to low prices for agricultural goods in the 1920's. Thus, despite high yields, farms were struggling even prior to the Great Depression.

By the time moderate drought began on the Plains in 1931, farms had increasingly plowed up native grasslands. Dry conditions persisted, contributing to a loss in ground cover that made the land susceptible to self-perpetuating dust storms (wind erosion). The hardened bare ground also experienced substantial runoff during the occasional heavy rain storm (water erosion).

In 1934 and 1936, there were devastating droughts and widespread crop failure. From 1934 to 1936, especially large dust storms blew topsoil across the Plains. The occasional fierce rain, especially in 1935, carved deep gullies in farmland. This period of massive erosion (the Dust Bowl period) is considered to have ended by 1941 at the latest, along with the US entry into WW2, the end of the Depression, and the return of generally wetter weather.

There is limited precise information, however, on when erosion occurred in particular areas, as the stock of topsoil was not recorded in comparable units at different points in time. The Soil Erosion Service (SES), established in 1933 and expanded to become the Soil Conservation Service (SCS) in 1935, published erosion maps from 1933 to December 1937, at which point there is a gap in the National Archives collection of soil erosion maps until the late 1940's. The SCS marked a very general boundary for which areas experienced severe wind erosion from 1935 to 1936, in 1938, and from 1935 to 1938. There are few historical accounts or references to severe erosion occurring after 1938, despite some incorrect information that has circulated.[1] Thus, I consider the Dust Bowl erosion to have occurred between 1931 and 1938, i.e., between the collection of census data in 1930 and 1940.

---

[1] Worster (1979, p.30) presents a map of the SCS wind erosion regions (1935-1936, 1938, 1935-1938) that also adds a substantial wind erosion area in 1940. Geoff Cunfer discovered the source of this error: a 1940 USDA document, which cites a December 8, 1939 Washington Evening Star newspaper article, which in turn cites information provided by SCS technicians. The displayed 1940 wind erosion region was projected in 1939, and the 1940 USDA document states that these projections proved incorrect and that there turned out to be little blowing. The 1939 newspaper article map also displays how the Dust Bowl region evolved from 1935-1939. In addition to the common 1935-1936 and 1938 regions, this map shows a 1937 region that is similar to and contained within the 1935-1936 region and a 1939 region that is at the center of the 1938 region and is much smaller. Neither Cunfer nor I has found these 1937 and 1939 regions in other archived publications. My interpretation is that these regions were included by the newspaper to fill in the time series from 1935-1939, rather than as a reflection of particular new blowing in those years.

I.B Causes of the Dust Bowl and Government Policy

The Dust Bowl was caused by a combination of prolonged severe drought and intensive land-use, though it is debated how much weight should be attributed to each factor.[2] New Deal government officials emphasized that an unsustainable amount of Plains land had been plowed for crops and that the remaining grassland was over-grazed by livestock. There was a fear that the region would become the once-imagined "American Desert" if private land-use practices were allowed to continue (Science 1934, News-week 1936). Attributing the Dust Bowl to unsustainable land-use became an important rationalization for New Deal agricultural policies.

One explanation for the Depression was farmers' lack of purchasing power, due to low prices for agricultural goods. The 1933 Agricultural Adjustment Act was designed to raise prices by paying farmers to reduce planted acreages. The government also purchased livestock at above-market prices, canning some meat and destroying the rest (Leuchtenburg 1963, Saloutos 1982).

As these initial policies became less popular and the Supreme Court declared aspects unconstitutional, the policies' stated purpose shifted toward conservation. The 1935 and 1936 Soil Conservation and Domestic Allotment Acts were partly an effort to combat erosion and preserve the country's farmland resources, but they also represented the new banner under which to reduce agricultural production (Rasmussen and Baker 1979, Phillips 2007). The 1938 Agricultural Adjustment Act combined efforts to maintain higher prices with this new focus on erosion prevention, nominally making a portion of acreage reduction payments conditional on conservationist land-use adjustments. The SCS established soil erosion control projects, aimed at demonstrating the effectiveness of their recommended soil conservation techniques (SCS 1937). These adjustments were mainly plowing along contour lines, retaining crop residues to provide ground cover, fallowing land with protective cover, planting

---

[2] The typical historical view emphasizes the over-exploitation of land (SCS 1955, Worster 1979, Hurt 1981), as did most contemporaries (SCS 1935 and 1939, Hoyt 1936, Wallace 1938). Cunfer (2005) emphasizes the impact of severe weather.

alternate strips of cash crops and drought-resistant crops, or converting cropland back to grassland.[3]

I.C Responses to the Dust Bowl

The Dust Bowl is associated with substantial out-migration, most famously described in Steinbeck's The Grapes of Wrath. There is little empirical evidence, however, that separates Dust Bowl migration from other sources of change in Plains states' population. All migrants from the region may have been inappropriately (and derogatorily) lumped together as "Okies."[4]

Most historical discussions of agricultural land-use have focused on its role in causing the Dust Bowl, but many give the impression that there was substantial adjustment following the Dust Bowl. Historical accounts of government programs generally suggest that these encouraged substantial land-use adjustment.[5]

Hansen and Libecap (2004) present evidence that externalities played an important role in causing the Dust Bowl. Farmers could discourage wind erosion by fallowing land or converting land to grasslands and pasture, but much of the benefit would be captured by neighboring farms. They argue that increased farm sizes and soil conservation districts lowered these coordination problems, and that resulting production adjustments prevented the Dust Bowl's reoccurrence.[6]

Agricultural adjustment to the Dust Bowl would be expected to reflect the substantial decrease in productivity due to soil erosion. In particular, erosion may have changed the

---

[3] The SCS also took greater responsibility over the administration of Civilian Conservation Corps work camps, which employed workers on public projects to reduce unemployment. Some of this work focused on preventing erosion: terracing land, rehabilitating gullies, and planting trees as wind breaks.

[4] Worster (1979, p51) describes interviews with migrants registering at Federal Emergency Relief offices around the country: only 12% of families from Oklahoma attributed their migration to farm failure, 17% of families from Kansas, and 16% of families from Colorado.

[5] For example, see Hurt (1985). In contrast, Cunfer (2005) notes that overall land-use patterns appear surprisingly stable from 1920 to 2000. Yet Cunfer's aggregate analysis of land-use may overlook important adjustments within agricultural production (Hansen 2005).

[6] This is in line with much of the historical literature However, Worster (1979) argues that the Dust Bowl did reoccur to some degree, and that conservation never replaced capitalism as the dominant farming mentality.

relative returns to different productive activities. If the production of crops is more sensitive to soil quality, then farmers may shift more-eroded land into pasture for animals. Similarly, land may be shifted from soil-sensitive crops (wheat) to more erosion-resistant crops (hay/grass). Agricultural adjustment will also interact with migration and adjustment in non-agricultural sectors, as discussed later.

I.D Aggregate Trends in Agriculture and Population on the Plains

Figure 1 presents aggregate changes in agriculture and population on the Plains over the 20th century.[7] Panels A and B show changes in the per-acre value of farmland and farm revenue in logs.[8] Average farmland became less valuable and less productive from 1910 to the 1940's. This partly reflects an increase in total farmland (Panel C), which expanded into less valuable areas. Average land values and revenues rose after total farmland stabilized in the 1940's.

Panel D shows that the total amount of cropland, as a fraction of total farmland, was constant from 1925 to 1930, decreased from 1930 to 1945, and was then mostly constant. This decrease is consistent with accounts that Plains farming had become too intensive prior to the 1930's and then found it necessary to scale back cultivation. However, this decrease is also consistent with an expansion of farming from 1930 to 1945 into areas less-suited for growing crops.

Panel E shows that there was a substantial increase from 1930 to 1940 in the proportion of cropland kept fallow. Much of this increase was short-lived, but fallowing remained above its level in 1930. Panel F shows that there was a gradual increase in average farm sizes after

---

[7] Aggregate numbers are calculated by summing over the same 769 Plains counties in every reported period. The only exception is in Panel E. Data on land fallowing are available from 1925 to 1945 for all 769 counties, when fallow land is defined as cropland left idle or in cultivated summer fallow. In 1950 and after, data are only available in 587 counties for land in cultivated summer fallow. Following Hansen and Libecap (2004), this series is treated as comparable and the sample is restricted to the 587 counties over the entire period.

[8] The value of farmland is divided by the CPI, as land is an asset whose value will change with the price of consumption. The value of revenue is divided by the PPI for all farm products, to focus on changes in productivity rather than output prices.

1930. The aggregate changes in panels E and F are consistent with results presented by Hansen and Libecap (2004): increased farm sizes may have allowed farmers to internalize the positive externalities from preventing wind erosion by fallowing land and decreasing cultivation shares.

Panel G shows that total population grew at a roughly constant rate from 1910 to 2000, with somewhat slower growth from 1930 to 1950. Note that, from 1930 to 1940, total population declined in the five central Plains states (Oklahoma, Kansas, Nebraska, South Dakota, North Dakota) between 3% and 7%. Panel F shows that the fraction of population in rural areas declined at a constant rate from 1910 to 1940, and a somewhat quicker rate from 1940 to 1970. Some of this decline is mechanical: as population grows, more people will be in non-rural areas (places with more than 2,500 inhabitants).

Figure 1 provides an overall empirical context in which to interpret the estimated relative changes. However, it is difficult to isolate the effect of the Dust Bowl in these aggregate changes. Even relative to other areas of the country, other shocks may substantially affect these outcomes (e.g., WW1, the Great Depression, WW2, technological change, price changes). The empirical analysis that follows will attempt to isolate the effects of the Dust Bowl, controlling for state-specific changes over each time period.

II Theoretical Framework

This section outlines a model in which permanent erosion during the Dust Bowl may affect agricultural production differently in the "short-run" and "long-run." In the short-run, farmers can only adjust variable inputs. In the long-run, farmers can adjust product choices or fixed inputs at some cost. This model helps to interpret the speed and magnitude of adjustment in agricultural production decisions, and to compare short-run and long-run costs using observed changes in land values, revenue, and inputs.

The effects on agriculture are then interpreted within a basic model of long-run supply and demand across locations, sectors, and factor inputs. This model guides the empirical analysis of migration and non-agricultural outcomes, and informs how cross-county spillovers influence the estimated relative effect of the Dust Bowl within states.

II.A A Model of Agricultural Production

Assume perfectly competitive markets, in which a farmer chooses production decisions in every period to maximize the present discounted value of profits. The farmer produces one good and allocates a share θ of land between two production technologies, $\mathrm{F}_1(\theta, V_1)$ and $\mathrm{F}_2(1-\theta, V_2)$, that are each a function of variable input choices V. The variable inputs reflect factors that can be adjusted quickly, such as labor, and are purchased in each period at a constant price.[9] The two technologies reflect different production methods or choices of fixed inputs that can only be adjusted slowly, such as land characteristics (cropland vs. pasture) or product choices (wheat vs. hay).

For any given land allocation, the farmer chooses variable inputs to maximize profits. Define $\Pi_1(\theta)$ and $\Pi_2(1-\theta)$ as the maximum obtainable profit after allocating land to each technology. Given these maximum obtainable profits, the farmer chooses θ to maximize total profits: $\Pi_1(\theta)+\Pi_2(1-\theta)$. Assuming that these two functions are differentiable and concave, an optimal interior $\overline{\theta}$ is chosen such that $\Pi_1'(\overline{\theta}) = \Pi_2'(1-\overline{\theta})$.[10]

The farmer obtains an initial profit $\pi^I = \Pi_1(\overline{\theta})+\Pi_2(1-\overline{\theta})$. The value of land is assumed to equal the present discounted value of profits, $\frac{\pi^I}{1-\beta}$, where β is a constant discount factor.

---

[9] This baseline model assumes that variable inputs are supplied in competitive markets, in which agriculture is only a small part. This rules out general equilibrium effects where input prices depend on agricultural production.

[10] This analysis will focus on interior solutions, where technological change is simply an adjustment along an existing production possibility frontier. If the initial equilibrium were a corner solution and only one technology were used, then subsequent adoption of the other technology could be interpreted as technological growth.

This assumption is consistent with competition and free entry in land markets, with the marginal land buyer having access to credit at a fixed interest rate.[11]

II.B Agricultural Production After a Permanent Shock

An unexpected and permanent shock at $t = 0$ is assumed to decrease the relative profitability of the first technology. Define periods $t < T$ as the "short-run" when variable inputs can adjust but the land allocation is constrained at its previous level ($\theta = \overline{\theta}$). The land allocation is then completely adjusted when $t = T$, after which is the "long-run." Assume initially that adjusting the land allocation is costless, but the model will then highlight the impact of adjustment costs.

This artificial distinction between the "short-run" and "long-run" can be interpreted as a simplified case of an unconstrained optimal adjustment model in which several factors could delay adjustment. First, adjustment costs could be convex, e.g., if needed capital or other inputs have an upward sloping supply curve in each period. Second, adjustment costs could decline over time due to learning-by-doing or other positive spillovers in technological adoption (Griliches 1957, Foster and Rosenzweig 1995). Third, the potential for learning to resolve uncertainty about optimal adjustments could delay investment responses (Dixit and Pindyck 1994, Bloom et al. 2007). Fourth, the original land allocation could have required fixed investments that depreciate over time, and would be lost during adjustment (Chari and Hopenhayn 1991).

Change in Profits. By assumption, the relative profitability of the first technology decreases in the short-run, given the new optimal choice of variable inputs. For analytical convenience, assume that it decreases by a constant percentage: $\Pi_1(\theta)$ decreases to $\delta\Pi_1(\theta)$, where $\delta \in (0,1)$. The land allocation constraint binds because $\delta\Pi_1'(\overline{\theta}) < \Pi_2'(1-\overline{\theta})$. The farmer earns a short-run profit $\pi^{SR} = \delta\Pi_1(\overline{\theta}) + \Pi_2(1-\overline{\theta})$. In the long-run, a new optimal

---

[11] If this assumption were relaxed and the Dust Bowl temporarily lowered the discount factor, then land values would decline more in the short-run and increase more in the long-run.

$\widehat{\theta}$ is chosen such that $\delta\Pi_1'(\widehat{\theta}) = \Pi_2'(1-\widehat{\theta})$. The farmer earns a long-run profit $\pi^{LR} = \delta\Pi_1(\widehat{\theta}) + \Pi_2(1-\widehat{\theta})$.

Given an initially binding land allocation constraint, it must be that $\widehat{\theta} < \overline{\theta}$ and $\pi^{SR} < \pi^{LR} < \pi^{I}$. Profit increases from the short-run to the long-run by $\int_{\widehat{\theta}}^{\overline{\theta}} \Pi_2'(1-x) - \delta\Pi_1'(x)dx$; intuitively, the difference in marginal profit is regained for each unit adjusted. Taking a first-order Taylor expansion of each marginal profit function around $\overline{\theta}$, this term simplifies to $\frac{1}{2}(\overline{\theta}-\widehat{\theta})(1-\delta)\Pi_1'(\overline{\theta})$. The value of this "adjustment triangle" corresponds to one-half the change in land allocation multiplied by the initial decrease in marginal return. The approximation is exact if $\Pi_2'(1-\overline{\theta})$ and $\Pi_1'(\overline{\theta})$ are linear. As one would expect, profits fall in the short-run and partially recover in the long-run as the land allocation adjusts.

Change in Land Values. Land values are derived as the net present value of profits, given a profit stream of π^{SR} until period T and π^{LR} thereafter. In each period $0 \leq t \leq T-1$, the value of land is $\sum_{i=t}^{T-1} \pi^{SR}\beta^{i} + \sum_{i=T}^{\infty} \pi^{LR}\beta^{i}$; for $t \geq T$, the value of land is $\frac{\pi^{LR}}{1-\beta}$.

This has three implications that will be estimated from the data. First, land values initially fall by a smaller percentage than profits. This is due to the expected long-run recovery in profits. Rearranging terms, the value of land at $t=0$ is $\frac{\pi^{SR}}{1-\beta} + (\pi^{LR}-\pi^{SR})\sum_{i=T}^{\infty}\beta^{i}$. This expression is greater than the PDV of short-run profits (the first term) when the long-run recovery in profits is larger and the long-run is achieved sooner.

Second, land values recover somewhat over time as periods of short-run profits are past and periods of long-run profits become more immediate. Rearranging terms, the value of land in each period $0 \leq t \leq T-1$ is $\frac{\pi^{LR}}{1-\beta} - (\pi^{LR}-\pi^{SR})\sum_{i=t}^{T}\beta^{i}$. Land values increase in each period, though by a smaller percent than profits increase when the long-run is reached.

Third, the value of land at $t=0$ capitalizes the full PDV economic loss associated with the shock. Rearranging terms, the value of land at $t=0$ is $\frac{\beta^{T}\pi^{LR}+(1-\beta^{T})\pi^{SR}}{1-\beta}$. Intuitively, this is a weighted average of long-run profits and short-run profits, where the weights are the relative value of each.

Change in Land Allocation. In the long-run, land is allocated away from the production technology that experienced a relative decrease in profitability. Rearranging terms from the profit function Taylor expansion, the change in land allocation $(\widehat{\theta} - \overline{\theta})$ equals the initial decrease in marginal return $((1 - \delta)\Pi_1^{'}(\overline{\theta}))$, divided by the summed slopes of the two technologies' marginal returns at the initial equilibrium $(\Pi_2^{''}(1 - \overline{\theta}) + \Pi_1^{''}(\overline{\theta}))$. The change in land allocation will be greater when there is a larger decrease in marginal return, and smaller when marginal returns are more sensitive to changes in land allocation.

Change in Revenue and Inputs. There are no general theoretical predictions about changes in revenue or inputs, because they depend on the functional form of the production function. However, a Cobb-Douglas functional form provides a useful benchmark: because factor shares are constant, there would be proportional decreases in profits, revenue, and inputs.

All major inputs are not reported in the data, so it will not be possible to analyze profits directly. Rosenzweig and Wolpin (2000) illustrate how unobserved inputs, particularly family labor, can bias estimates of the effect of environmental shocks on profits (revenues minus observed input costs). The analysis here will instead focus on percent changes in revenue, which would proxy for percent changes in profits assuming a Cobb-Douglas functional form. This assumption is quite strong, so an empirical analysis of observed inputs will explore the potential direction of bias.

The Effect of Adjustment Costs. Consider now the differences in these effects when the farmer must pay a non-recoverable cost $C_{12}(L)$ to shift land L from technology 1 to technology 2. For simplicity, assume that the farmer will pay this one-time cost in period T.

Assuming that some land adjustment is still optimal, the new long-run land allocation is chosen in period T to satisfy the condition: $(1 - \delta)\Pi_1^{'}(\widehat{\theta}) = \Pi_2^{'}(1 - \widehat{\theta}) - (1 - \beta)C_{12}^{'}(\overline{\theta} - \widehat{\theta})$. This condition assumes that the farmer has perfect access to credit so that only $(1-\beta)$ of the adjustment cost is paid in each period; if instead the farmer faces some capital constraints, then full initial adjustment will be more costly.

Compared to the model with no adjustment costs, there will be less adjustment in land allocation. Thus, there will be less long-run recovery in profit, in addition to the direct reduction in profits from the paid adjustment cost. The value of land will initially decrease by an additional term, equal to the PDV of the future adjustment cost. This additional term increases over time as the adjustment cost becomes more immediate. At $t = T$, the value of land then increases by the full value of the adjustment cost.[12] The sum total of these effects is clear in the special case when the adjustment cost is only ϵ less than the PDV long-run recovery in profits: the value of land decreases at $t = 0$, remains constant, and then partially recovers at $t = T$ when it increases by the full amount of the adjustment cost.

If adjustment costs can be fully recovered, such as in purchasing general machinery or livestock, then the value of land follows the same pattern initially and does not increase at $t = T$ when these costs are paid. Thus, observed changes in land values can indicate whether adjustment takes place on margins that require fixed or mobile investments. Note, however, that this model assumes perfect foresight. Any systematic errors about the future costs of erosion, potential government subsidies, or other factors will also influence the evolution of land values.

II.C Empirical Predictions for Agricultural Production

(1) The immediate decrease in land values reflects the present discounted value of the economic cost to agriculture.

(2) Profits recover partially in the long-run, as fixed production decisions adjust.

(3) If profits are expected to recover in the long-run, then land values fall by a smaller percentage than profits in the short-run.

(4) Under strong functional form assumptions, percent changes in revenue may proxy for percent changes in profits. If revenue recovers due to an increase in inputs rather than

---

[12] This refers to the gross value of land, which is reported in the data. The value of land, net of any mortgage debt incurred, would increase gradually as that debt is repaid.

technological adjustment, then the percent recovery in revenue will overstate the percent recovery in profits.

(5) Adjustment costs reduce or slow recovery in profits. If adjustment costs are fixed to the land, land values increase when the costs are paid.

II.D A Model of Adjustment Across the Broader Economy

If the Dust Bowl substantially affected agricultural production, this could have general equilibrium effects in non-agricultural sectors and non-Dust Bowl areas. This section outlines a basic model of supply and demand across locations, sectors, and factor inputs. This is a model of long-run supply and demand, in which many margins are free to adjust.

In this model, there are two locations (Dust Bowl and non-Dust Bowl). For example, this could be two counties within Oklahoma that are more- and less-eroded, or one eroded county in Oklahoma and one non-eroded county in California. There are two sectors (agriculture and industry) that produce freely tradeable goods using two homogeneous factors (land and labor). The supply of land is fixed in each location. Labor is supplied by workers who pay a cost to change location or sector. Workers must supply labor in the same location that they live. Workers consume land (housing), agricultural goods, and industrial goods. Assuming perfectly competitive markets, all prices (land values, wages, and prices for agricultural and industrial goods) are set such that each market clears.[13]

For the comparative statics, assume that agricultural productivity declines in the Dust Bowl area. This decreases the demand for agricultural land. Consider the case in which the supply of land for agriculture is inelastic, i.e., workers and industrial firms (current or entering) have little use for additional marginal lands. In this case, adjustment in the land market mainly occurs through decreased land prices with little change in total farmland.

Assume further that the Dust Bowl decreases the productivity of agricultural labor. Workers then have an incentive to move to the non-Dust Bowl area and/or switch to the

---

[13] This model is similar to that derived by Roback (1982), which focuses on urban amenities.

industrial sector. Equilibrium wages remain relatively lower in the Dust Bowl area (particularly in the agricultural sector) to the extent that there are costs to moving (and switching sectors). Workers consume some land that is now cheaper, so paid wages fall by more than local price-adjusted wages.

The industrial sector has an incentive to expand in the Dust Bowl area for two reasons. First, cheaper land is attractive if the sector has any use for additional lands. Second, the industrial wage would decrease to the extent that workers switched sectors or consumed cheaper land.[14] Lower wages encourage higher labor-capital ratios in the agricultural sector, and discourage adjustment toward production methods that use less labor (such as animals or hay, relative to crops or wheat).

In the non-Dust Bowl area, wages decrease in response to an in-migration of workers. Because the Dust Bowl area is less productive, agricultural output prices increase and the agricultural sector expands in the non-Dust Bowl area. As a result, land prices increase. The industrial sector contracts due to higher land prices and lower output prices.

II.E Empirical Predictions for the Broader Economy

(1) Workers migrate from the Dust Bowl area, though this incentive is tempered by lower housing prices and moving costs. If workers migrate to the compared non-Dust Bowl area, the relative comparison in population will overstate the number of migrants.

(2) A greater share of workers become employed in the industrial sector, which also expands due to lower land values and wages.

(3) If industry and workers have inelastic demand for agricultural land, then adjustment in land markets occurs mostly through lower prices rather than a contraction in total farmland.

(4) Lower wages encourage labor-intensive production in the agricultural sector.

---

[14] The industrial sector could contract if it was supplied inputs by the local agricultural sector, or if it sold outputs locally to the agricultural sector. This is not captured in the model, but could be empirically relevant.

(5) Decreased agricultural production in the Dust Bowl area raises land values in the non-Dust Bowl area, causing the relative comparison of land values to overstate the total cost of the Dust Bowl.

(6) Discouraged production activities in the Dust Bowl area are encouraged in the non-Dust Bowl area, due to increased output prices. Relative comparisons in production across sectors or agricultural products would overstate adjustment in the Dust Bowl area.

III Data and Summary Statistics

III.A Measuring Erosion

This paper uses new detailed data on which areas were most eroded after the 1930's, which I gathered by digitizing soil erosion survey maps currently in the National Archives. In response to the Dust Bowl, these maps were produced by the Soil Erosion Service (SES) and its successor, the Soil Conservation Service (SCS). The SES published the 1934 Reconnaissance Erosion Survey, which detailed the severity and type of erosion across the United States based on measurements taken during visits to each county (SCS 1935). In August 1936, the SCS published a map of which general areas had been affected to different degrees and a second more-detailed map in December 1937.

This first type of erosion map indicates whether areas were slightly, moderately, or severely eroded; whether it was wind erosion (topsoil being swept from the land) or sheet erosion (water generally shifting the topsoil off the land); and whether there were many gullies (formed by small streams of runoff carving paths in the land). There are two main limitations to this map type. First, it only reflects the stock of erosion at the time of the survey. Since there is no detailed baseline data and the erosion category definitions change over time, it is not possible to measure directly the erosion that occurred during the Dust

Bowl. Second, it is difficult to compare erosion across categories and generate a single index for how much an area was affected.[15]

The SCS prepared a second type of map that indicates dated regions of wind erosion on the Plains. These show broad areas that experienced blowing from 1935 to 1936 and in 1938, which correspond closely to areas with the highest wind speeds (Chepil et al. 1962). These regions have less local variation, however, and it is not clear how the exact areas were determined. There was substantial wind erosion at other times during the Dust Bowl, as well as water erosion, and these maps give little sense of the cumulative effect on the soil.

My analysis focuses on a third type of SCS erosion map, which shows cumulative erosion damage after the Dust Bowl. The map classifies areas as slightly eroded (less than 25% of topsoil lost), moderately eroded (25% to 75% of topsoil lost), and severely eroded (more than 75% of topsoil lost). Of all the SCS erosions maps, this map has the most detailed local variation in erosion. Documentation is sparse, but this map (along with most others) is recorded as "based on the 1934 Reconnaissance Erosion Survey and other surveys." The National Archives have three copies of this same map with publishing dates in 1948, 1951, and 1954 – prior to the next substantial period of erosion in the mid-1950's.

The main limitation of this cumulative erosion map is that it does not just reflect erosion that occurred during the 1930's Dust Bowl. Total erosion would have begun to increase along with the settlement of the Plains. Large-scale detailed erosion surveys only began in the 1930's, so there is no systematic baseline data on erosion and topsoil levels prior to the Dust Bowl. To focus on the change in erosion over the 1930's, the empirical analysis controls for a range of county characteristics in 1930, 1925, 1920, and 1910 that might predict baseline erosion levels.

To focus on the extreme erosion of the 1930's, the sample is restricted to counties within the 12 Plains states: Colorado, Iowa, Kansas, Minnesota, Montana, Nebraska, New Mexico,

---

[15] Hansen and Libecap (2004) assigned erosion categories based on the 1936 map, as they are focused on specifically wind erosion.

North Dakota, Oklahoma, South Dakota, Texas, and Wyoming.[16] In the cumulative erosion map, many areas in non-Plains states were also classified as moderately or severely eroded. Non-Plains states had only severe water erosion indicated in other maps, however, which is more likely to reflect gradual erosion over many previous decades. As a robustness check, the empirical analysis shows that more-eroded non-Plains counties did not experience the substantial relative drop in land values during the 1930's that occurred in more-eroded Plains counties. This suggests that the cumulative erosion measure is an effective proxy for the extreme Dust Bowl erosion on the Plains in the 1930's.

III.B Combining County Outcomes with Erosion Data

The core dataset is compiled from US Census data on agriculture, population, and industry, aggregated to the county level. Among the key variables are the value and quantity of agricultural land, agricultural revenue, total cropland and pasture, revenue from crops or animals, production and acreage for a variety of crops, the number of farms, population, rural population, farm population, retail sales, manufacturing workers and value added, and unemployment.[17] Other data sources include banking data from the FDIC; New Deal expenditures from the Office of Government Reports; and drought data from the National Climatic Data Center.[18]

The data are combined into a panel of US counties, where the unit of observation is held fixed at 1910 county boundaries. Census data are used from 1910, 1920, and then

---

[16] The results are similar when the sample is restricted further to "Great Plains counties," as defined by the Great Plains Population and Environment Database (which provided additional historical data on these 12 states). This reduces the sample from 769 counties to 356 counties concentrated in western Kansas and Oklahoma, northwest Texas, central and western Nebraska, eastern Colorado, and the Dakotas.

[17] An unusual aspect of this data concerns the calculation of agricultural revenue in 1920, 1925, and 1930. In these periods, revenues from animals sold and animal products sold are not reported. Only the total value of animals is reported. In 1910 and 1940, both of these variables are reported. The ratio of revenue to stock is calculated for each county and used to impute the county's value of animal revenue in the intervening years. The results are not sensitive to using the revenue/stock ratio in 1910, 1940, or a weighted average of the two – the results presented here use a weighted average (with the 1910 ratio weighted by two-thirds in 1920, one-half in 1925, and one-third in 1930).

[18] I thank Price Fishback for providing the New Deal data (Fishback et al. 2005) and the drought data (Boustan et al. 2007).

approximately every 5 years through 1997, though data availability for some variables is more restricted. Data from 1935 is dropped, due to ongoing erosion from the Dust Bowl. When counties split or borders were otherwise adjusted in subsequent periods, the reported Census data are aggregated or adjusted to reflect the original 1910 county border (Hornbeck 2008). The base year was chosen as 1910 in a tradeoff between more sample counties and a longer pre-period for analysis: changing from 1900 to 1910 substantially increases the number of sample counties in Oklahoma and other areas, while moving beyond 1910 has less effect on sample size and the pre-period becomes less informative about differential trends.

The core results focus on the same sample of counties, for which data on key variables are available in every period of analysis.[19] This allows for a clearer interpretation of estimated changes over time and across specifications, because the sample composition is not also changing.

Figure 2 displays these 769 sample counties, along with the mapped cumulative erosion levels. The white area represents low erosion, the light gray is medium erosion, and the dark gray is high erosion (the crossed out areas are not in the sample).[20] The figure was constructed by overlaying the cumulative erosion map on a map of United States historical county borders using Arcview Geographical Information System (GIS) software. To create the data, the erosion areas were traced and merged with 1910 county borders to assign each county the percent of its area that falls within each erosion category.

---

[19] In a tradeoff between timespan and sample size, the sample is restricted to counties that have data on: land values, revenue, and farmland through 1992; cropland through 1974; population and rural population through 1969. The sample also excludes a small number of counties with less than 1000 acres of farmland and a few counties with extreme values due clearly to measurement error and/or the adjustment to maintain constant boundaries.

[20] The main reasons why areas are dropped from the sample are: the area was not sufficiently settled by 1910 for data to be reported; data are unavailable in 1925; agricultural revenue could not be imputed for 1920, 1925, or 1930; and border adjustments left data unavailable for any piece that made up more than 1% of the original county's total area.

III.C Summary Statistics, by Erosion Level

Table 1 presents baseline differences between sample counties in 1930, based on their assigned post-Dust Bowl erosion level.[21] For comparison, column 1 presents mean characteristics for all counties in 1930. Columns 2 and 3 report coefficients from a single regression of each outcome variable on the percent of a county in medium erosion and high erosion, controlling for state fixed effects.[22]

Counties with more erosion after the Dust Bowl tended to have previously: higher land values, denser population, more but smaller farms, a larger fraction of cropland in corn or cotton (as opposed to wheat and hay), and more animals. Column 4 reports the average difference between a county with high erosion and a county with medium erosion. These counties were more similar, though high erosion counties have more cropland allocated to corn.

Measured erosion levels are not randomly assigned for two reasons. First, part of this erosion occurred prior to the Dust Bowl and could be caused by, reflected in, or otherwise jointly determined with pre-Dust Bowl county characteristics. Second, the Dust Bowl's intensity was partly determined by the intensity of cultivation and other county land-use practices. The main empirical challenge is that counties with these different characteristics may have changed differently after the 1930's, even if the Dust Bowl had not occurred.

To address this issue, the empirical analysis will focus on specifications that control for differential changes over time that are correlated with state, county characteristics in 1930, and lagged values of those characteristics. The empirical analysis will not control for county characteristics after 1930, as these are potential outcomes of the Dust Bowl. A robustness check will instrument for county erosion levels using intensity of drought during the 1930's.

---

[21] Given the construction of the sample, there are no missing values for panels A, B, and C. For panels D and E, missing values are assumed to be zero (generally in areas where these products are not produced).

[22] The means and regressions are weighted by county farmland levels, as the empirical analysis will be focused on the effect for an average acre of farmland. These variables will all be included as controls in later weighted regressions, so their weighted difference is more relevant to interpreting their importance as control variables.

IV Empirical Framework

The empirical analysis is based upon comparing changes in outcomes for counties that experience different levels of erosion. For a graphical analysis of the data over the entire sample period, outcome Y_{ct} in county c and year t is regressed on the percent of the county in medium erosion (M_c) and high erosion (H_c) areas, a state-by-year fixed effect (α_{st}), linear functions of county characteristics (X_c), and an error term (ϵ_{ct}):

$$(1) \qquad Y_{ct} = \beta_{1t} M_c + \beta_{2t} H_c + \alpha_{st} + \theta_t X_c + \epsilon_{ct}.$$

The effects of erosion and each county characteristic are allowed to vary in each year. The included county characteristics are each of the variables in panels B – E of table 1, and their lagged value from all available previous years. In each regression, every county included has data in every analyzed period.

This regression is a special case, in which each regressor is fully interacted with each time period. Thus, controlling for a county fixed effect would not change the difference between estimated β's; rather, it simply normalizes the β's relative to zero in some base period. County fixed effects are omitted in this regression in order to observe the pre-Dust Bowl difference in outcome levels across counties with different erosion levels, after controlling for the other pre-Dust Bowl characteristics.

Graphing the estimated β's shows how counties with each erosion level changed over the sample period, relative to a county with low erosion. To interpret these changes as the relative effect of the Dust Bowl, the primary identification assumption is that counties with different erosion levels would have changed the same if not for the Dust Bowl. In practice, this assumption must hold after controlling for differential changes over each time period that are correlated with each state and pre-Dust Bowl county characteristic.

From the estimation of equation (1), it is possible to observe whether counties with different erosion levels were on similar trends from 1910 to 1930 (after controlling for other covariates). Because counties may not be on the same trend prior to the Dust Bowl, the empirical analysis then controls for these differences in pre-trends.

For estimating the numerical results, two modifications are made to equation (1). First, the regression controls for the outcome level in all available pre-Dust Bowl periods, L_c. This is a flexible way to control for differential pre-trends, without assuming that trends prior to the 1930's would have continued predictably – which is theoretically inappropriate for asset prices such as land values, and is empirically questionable for other outcomes.

Second, the outcome Y_{ct} is differenced from its value in 1930, so that each coefficient can be interpreted as the relative change since 1930. Differencing the data also improves the precision of the estimates by absorbing any fixed county characteristics; yet, as before, the estimated coefficients are not changed by differencing the data or by including a county fixed effect.[23] The regression is otherwise the same:

$$(2) \qquad Y_{ct} - Y_{c1930} = \beta_{1t} M_c + \beta_{2t} H_c + \alpha_{st} + \theta_t X_c + \gamma_t L_c + \epsilon_{ct}.$$

As before, note that the effects of erosion and each county characteristic are allowed to vary in each year. In order to condense the results, a specification is sometimes estimated that pools the erosion variables for some combination of later time periods. The control variables are still allowed to vary in each year, so this amounts to averaging the estimated β's across years.

Regressions for agricultural outcomes are weighted by county farmland (or an analogous land measure) in 1930 in order to estimate the average effect for an acre of farmland. Regressions for labor outcomes are weighted by county population in 1930 in order to estimate the average effect for a person. Standard errors are clustered at the county level, to adjust for heteroskedasticity and within-county correlation over time. To check for spatial correla-

---

[23] Because the sample is balanced and the regressors are fully interacted with each time period, estimation of the regression is completely separable across year pairs. The methods are equivalent in the case of two time periods. While the coefficients are identical, the methods have slightly different standard errors. This differencing was done throughout for ease of interpretation and computational speed.

tion among counties, Conley standard errors are estimated for changes in land values from 1930 to 1940 (Conley 1999). These standard errors are similar to standard errors clustered at the county level, indicating that the impact of the Dust Bowl was not highly spatially correlated.[24]

V Results

V.A Agricultural Land Value and Revenue

Agricultural Land Value. Figure 3 graphs the β's from estimating equation (1) for the per-acre value of farmland. High and medium erosion counties had slightly higher land values than low erosion counties in 1910, after adjusting for the control variables. These counties were more similar by 1920 and then trended similarly through 1930.

After the Dust Bowl, high erosion counties experienced an immediate, substantial, and persistent relative decrease in land values. Medium erosion counties experienced a smaller but substantial relative decrease. Comparing the two lines, high erosion counties also experienced a persistent decrease relative to medium erosion counties.

To account for differences in the level and trend of county land values prior to the 1930's, the numerical results are obtained by estimating equation (2). This regression controls for county land values in 1930, 1925, 1920, and 1910, interacted with each year.[25] These controls absorb the graphed pre-1930's differences by erosion level, and are a flexible way of adjusting for differential pre-trends in land values. Table 2, column 1, presents the results.

---

[24] The Conley method allows for outcomes to be correlated among nearby counties, with the degree of correlation declining linearly until some cutoff distance. Relative to when there is restricted to be no cross-county correlation, the percent increase in standard errors for changes in high erosion vs. low erosion counties is: 0% (50 mile cutoff), 5% (100 mile), 16% (300 mile), 30% (500 mile), 29% (700 mile), 23% (900 mile). For changes in medium erosion vs. low erosion counties, the Conley standard errors decline by 0% to 17%. The Conley specifications are not also weighted by county farmland levels, so the comparison in standard errors is relative to when not weighting the regression.

[25] In this case, the regression also controls for agricultural revenues per-acre in each pre-period, as the later analysis will combine the specifications in a seemingly unrelated regression framework.

Panel A reports that, from 1930 to 1940, land values fell by 27.8% in high erosion counties relative to low erosion counties. Land values fell by 16.7% in medium erosion counties relative to low erosion counties. If we assume that low erosion counties were not affected by the Dust Bowl, the total capitalized cost can be approximated by multiplying the percent decline in land values by the original value of those acres. These estimates indicate that the Dust Bowl erosion imposed an economic cost on agriculture of $153 million in 1930 dollars ($1.9 billion in 2007 dollars).[26] If low erosion counties were also damaged by the Dust Bowl, this number would understate the total cost. If the Dust Bowl made less-affected land in the same state more valuable, e.g., through higher output prices, this number would overstate the total cost.[27] Given average 1930 land values, this reflects an area of farmland the size of Oklahoma becoming worthless.

Panel B reports the estimated change in land values from 1930 to 1945. Column 2 expresses this change as a fraction of the change in land values from 1930 to 1940. These numbers reflect the amount of decrease from 1930 to 1940 that persisted into 1945: 84% for high erosion and 61% for medium erosion. For changes in high erosion counties relative to medium erosion counties, the persistence was 119% (land values declined further). There is no strong *a priori* reason for one of these three relative comparisons to reflect a more or less persistent shock. Taking the efficient weighted average of these three parameters, the average persistence is 76.5%.[28]

---

[26] By multiplying 1930 county farmland levels by the share of its area in each erosion category, there were approximately 51 million farm acres in high erosion areas and 170 million farm acres in medium erosion areas. The per-acre value of farmland was $3.60 in high erosion counties and $3.62 in medium erosion counties, weighting by farmland. The total cost is found by multiplying the acres affected, the value of the acres, and the percent decline in land values: $51 million from high erosion and $102 million from medium erosion.

[27] There is no obvious indication of higher prices from national trends in the ppi for farm products and the cpi for urban consumers: farm output prices were lower from 1938 to 1940 than in the 1920's – both absolutely and relative to the urban cpi. All prices declined in the 1930's and increased during WW2, while farm output prices increased during the severe droughts from 1934 to 1936.

[28] The variance-covariance matrix for these three estimates is known, so the efficient average is estimated through GLS (regressing the three values on a constant). Intuitively, this procedure gives more weight to the more-precisely estimated coefficients, and less weight to those coefficients are are more correlated with each other.

Panels C, D, and E report later changes in land values, pooling the estimates over 2, 3, and 4 census periods. There is little indication of a systematic recovery in land values. This suggests that the overall percent recovery in profits is not large, and that adjustment did not take place through fixed improvements in land that would be capitalized in its value.

Agricultural Revenue. Figure 4, panel A, graphs the β's from estimating equation (1) for the per-acre value of farm revenue. From 1930 to 1940, more-eroded counties experienced substantial declines in farm revenue. Revenue recovered partly from 1940 to 1945 and then the paths diverge: high erosion counties have lower revenue compared to low erosion counties, while revenue gradually recovers in medium erosion counties. Revenue had been increasing in more-eroded counties from 1910 to 1925, and then decreased from 1925 to 1930. Table 2, column 3, presents the numerical results for farm revenue, controlling for these pre-1930 differences.

Panel A reports the relative decrease in farm revenue from 1930 to 1940: 31.6% for high erosion counties and 20.2% for medium erosion counties. Panels B to E report the changes for later periods, and column 4 expresses these changes as a fraction of the decrease from 1930 to 1940. From 1940 to 1945, more-eroded counties experienced a substantial recovery in revenues, though averaged levels were still lower than in 1930. Over later periods, however, only medium erosion counties experienced a substantial recovery in revenues. On average, approximately 70% of the initial decline in revenues persisted between 1950 and the 1990's.

One way of interpreting these results is to take the present discounted value of all lost revenue after 1940. This calculation assumes an interest rate of 5%, linearly interpolates annual values between each census period, and assumes that revenue is constant after 1992. To obtain a measure of cost persistence, this PDV is expressed as a fraction of the PDV of lost revenue assuming no recovery after 1940.[29] The estimated persistence in losses (and standard error) is: 0.762 (0.132) for high erosion relative to low erosion, 0.687 (0.152) for medium erosion relative to low erosion, and 0.896 (0.351) for high erosion relative to medium

---

[29] To do this calculation, annual coefficients for column 4 are linearly interpolated. Given an interest rate r, the estimated persistence in year t is multiplied by $\frac{1}{(1+r)^t}$. These values are summed and divided by $\frac{(1+r)}{r}$.

erosion. If percent changes in revenue were a proxy for percent changes in profits in every period, then these numbers represent the full economic loss as a fraction of the short-run loss. However, if the recovery in revenue is due to an increase in production inputs, then these numbers will understate the persistence in lost profit.

The value of capital machinery and equipment is the primary production input for which comparable data are available across periods in all sample counties. Figure 4, panel B, graphs the results from estimating equation (1) for the log value of capital per acre of farmland. Capital inputs fell substantially from 1930 to 1940, though not by as much as revenues.[30] Capital inputs partially recovered from 1940 to 1945 and then diverged: inputs increased over time for medium erosion counties and declined again for high erosion counties. Thus, part of the observed recovery in revenue may be due to greater inputs, overstating the recovery in profits – particularly for the medium erosion counties in which revenue recovered substantially.

Agricultural Land Value vs. Revenue. Comparing the relative changes in land values and revenues provides a second estimate of the persistence in short-run costs. The percent decline in land values should reflect the true economic loss from the Dust Bowl, anticipating and appropriately discounting the future recovery in profits. Under strong functional form assumptions, the immediate percent decline in revenue would proxy for the immediate percent decline in profits. The true economic loss (land value) can then be expressed as a fraction of the short-run economic loss (revenue or profit).

Column 5 of table 2 reports this ratio, estimated for each relative comparison between erosion levels.[31] If the theoretical assumptions hold, these ratios have a clear structural interpretation: for every 1% lost in the short-run (1940), the full economic loss is between 0.82% and 0.98%, with a weighted average of 0.86%. Theoretically, this ratio should be weakly less than one, though the estimated parameters need not be. Because none are statistically

---

[30] From estimating equation (2) for capital inputs, the decrease was 0.225 (0.055) for high erosion and 0.121 (0.037) for medium erosion, relative to low erosion.

[31] To obtain the standard error of this ratio, both coefficients are estimated through seemingly unrelated regression, where the same control variables are included in each regression.

less than one, these estimates fail to reject the hypothesis that long-run adjustments did not mitigate the short-run costs. These coefficients are all statistically greater than zero, which rejects the hypothesis that short-run costs were associated with no PDV cost. If the theoretical assumptions fail to hold, these numbers still have an appealing reduced-form interpretation: the full economic loss (land value) is scaled by the magnitude of the shock (revenue).

The main limitations in interpreting this ratio are the strong assumptions required for immediate percent changes in revenue to proxy for immediate percent changes in profits. Capital inputs fell substantially from 1930 to 1940 but somewhat less than revenues, so profits may have fallen by a larger percent than revenues. In this case, the estimated ratios in column 5 would overstate the persistence in costs (the denominator should be larger). By contrast, the persistence in costs may have been understated by the estimated recovery in revenues because capital inputs increased after 1940. The difference between revenues and profits may then bias the two cost persistence estimates in opposite directions. The range of cost persistence estimates from the two approaches is: between 0.762 and 0.881 for high erosion vs. low erosion, between 0.687 and 0.824 for medium erosion vs. low erosion, and between 0.896 and 0.981 for high erosion vs. medium erosion.

V.B Measured Erosion as a Proxy for the Dust Bowl

The empirical analysis relies on mapped erosion levels indicating areas that were more or less eroded during the Dust Bowl, after adjusting for state fixed effects and pre-Dust Bowl county characteristics. This section presents further results on changes in land values, in order to explore the validity of these erosion measures.

Plains vs. Non-Plains. One concern is that some of the measured variation in erosion levels occurred prior to the 1930's. Indeed, the nationwide map of cumulative erosion indicates that there were many severely eroded areas in the Eastern United States, where the Dust Bowl did not cause additional erosion. If counties with different baseline erosion levels would have changed differently over the 1930's, this would be confounded with the impact of the Dust Bowl.

To explore this possibility, figure 5 compares estimated relative changes in land values for more-eroded counties in Plains and non-Plains states. For each year, panel A (panel B) graphs the estimated difference in land values between high and low erosion counties (medium and low erosion counties), controlling only for state-by-time fixed effects.[32]

Before 1930, more-eroded Plains counties have higher land values than less-eroded Plains counties; and this difference is greater than the difference between more-eroded and less-eroded counties in non-Plains states. One explanation is that "more-eroded" Plains counties had yet to experience much of that erosion prior to the Dust Bowl. Following the Dust Bowl, more-eroded counties on the Plains experience a large relative decrease in land values. By contrast, land values in more-eroded and less-eroded counties in non-Plains states were relatively unchanged. These counties are otherwise on similar relative trends before 1930 and after 1960, at which point the level differences are also similar.

These estimates have two implications. First, the cumulative erosion measure appears to proxy for the particularly destructive erosion on the Plains during the Dust Bowl. Second, differentially eroded counties experienced similar changes in land values over the 1930's when the Dust Bowl did not contribute to further erosion.

Instrumenting for Erosion with Drought. A related concern is that observed erosion levels partly reflect counties' land-use, both prior to 1930 and during the Dust Bowl. When analyzing the effect on land value, OLS estimates could be negatively biased if farmers in counties otherwise becoming less valuable did not protect their land from Dust Bowl erosion. OLS estimates could be positively biased if counties otherwise becoming more valuable were farmed more intensely and this caused greater Dust Bowl erosion. Measurement error in erosion levels could also attenuate the estimated effects.

---

[32] The other control variables are not available for all non-Plains counties. The sample is expanded to include all Plains and non-Plains counties with available land value data in each period: 867 Plains counties and 1840 non-Plains counties.

Some portion of the Dust Bowl erosion was due to severe drought on the Plains during the 1930's. As a robustness check, the analysis instruments for erosion levels using the number of months during the 1930's that a county was in extreme drought, severe drought, moderate drought, and the average level of drought over the 1930's.[33] The analysis now also controls for these four drought variables from 1895 to 1929, in order to capture general differences in susceptibility to drought.

The identifying assumption is that 1930's weather was not systematically correlated with future changes in land values, aside from its impact through changing erosion. These drought variables are sometimes correlated with pre-1930's county characteristics, so the empirical analysis continues to control for these variables. Temporary weather shocks should not have a direct effect on land values, though fixed investments in land could decrease in response to drought if income falls and credit is difficult to obtain.[34]

Table 3, columns 1 and 2, report the first-stage results. Column 1 reports estimates from regressing the percent of a county in high erosion areas on the drought instruments, controlling for drought from 1895 to 1929, state fixed effects, and the same pre-1930's county characteristics as before. There is more high erosion in counties with more months of extreme drought, and a weaker effect in counties with more severe drought. High erosion is not predicted by months of moderate drought or average drought over the 1930's. Column 2 reports analogous estimates for the percent of a county in medium erosion areas. Medium erosion is higher in counties with higher average drought levels and more moderate drought, with slightly weaker effects from severe and extreme drought. The instruments are jointly significant, but the overall explanatory power is low and raises weak instrument problems. Ongoing work is gathering additional potential instruments.[35]

---

[33] Drought is measured by the Palmer Drought Index, which is defined on a running basis depending on the arrival of rainfall and the loss of moisture. Extreme, severe, and moderate reflect conventional cutoffs in this Index. The drought data were collected at weather stations, and the PDI is reported by state and district (10 districts to a state). Details are available at ftp://ftp.ncdc.noaa.gov/pub/data/cirs/

[34] This exclusion restriction appears more likely to be violated for other county outcomes, so the instrumental variables specifications are only estimated for changes in land values.

[35] Average wind speeds may predict wind erosion, and slopes may predict water erosion. There are direct measures of soil characteristics and its susceptibility to erosion, but these data are only available in the

Table 3, column 3, reports the estimated relative changes in land values from 1930 to 1940, instrumenting for erosion levels with the 1930's drought variables. As a basis for comparison, column 4 presents the OLS results for the same sample and control variables. When instrumenting, higher erosion counties experienced a larger relative decline in land values (63% instead of 28%) and medium erosion counties experienced a somewhat smaller relative decline (12% vs. 17%). The implied cost of the Dust Bowl increases from $153 million to $187 million (in 1930 dollars). These IV estimates are much less precise, but they are consistent with the previous OLS findings: measured erosion levels appear to proxy for the impact of the Dust Bowl and decreased land values.[36]

V.C Adjustment in Agricultural Production

Agricultural costs from the Dust Bowl were estimated to be persistent. There are three main interpretations for this finding. The first interpretation, which later estimates support, is that adjustment in agricultural production was possible but costly and slow. The second interpretation is that adjustment to the Dust Bowl was not possible. The third interpretation is that adjustment was fast and had already taken place by 1940.

This section attempts to identify examples of possible margins of adjustment, and then what adjustments occurred. Results from estimating equation (1) are graphed in Figures 6 – 9, in order to observe relative pre-trends in each agricultural outcome. Table 4 reports the numerical results from estimating equation (2), which controls for differences in relative pre-trends.

Total Farmland. Figure 6 graphs the estimated changes in total farmland. This extensive margin of farming is fairly stable: immediately after the Dust Bowl, there are no sub-

---

modern era and may be an outcome of the Dust Bowl erosion. Hansen and Libecap (2004) use these average wind speeds and soil characteristics as instruments. The instruments could potentially be interacted with each other or pre-1930 county land-use. Because the instruments must predict the transition from low erosion to medium and then high erosion, non-linear functions of the instruments may be appropriate. In the case of the drought instruments, including their squared values yields similar 2SLS results.

[36] As a falsification test, the same specification can be used to examine pre-differences in land values: more-eroded counties have similar or lower land values than in 1930, instrumenting for erosion with 1930's drought. In periods after 1940, the estimates are quite variable but fluctuate around their value in 1940.

stantial or statistically significant relative changes in farmland. These small relative changes in total farmland suggest that the previous estimates of changing land values, revenue, and capital inputs do not reflect large changes in the underlying composition of farmland.

High erosion counties later experience a gradual small decline in farming, which by the 1950's is lower by a statistically significant 3% (table 4, column 1). These declines in farming are then largely reversed, however, and medium erosion counties are relatively unchanged.[37]

One explanation for the persistence in overall land allocation is that non-agricultural sectors of the local economy might have an inelastic demand for land, at least in the short-run. In this case, the supply of land to the agricultural sector would be relatively inelastic. When agricultural demand for land decreased following the Dust Bowl, we would then expect large declines in land values and little change in the extensive margin of farming.

After the original settlement of the Plains slowed, the aggregate allocation of land to agriculture was fairly constant (figure 1). Assume that aggregate demand for agricultural land is constant, and that county-level fluctuations are driven by county-specific demand shocks for non-agricultural land. If these demand shocks have little effect on the county allocation of land to agriculture, then the agricultural sector may face a relatively inelastic supply of land. Indeed, the standard deviation of county-level fluctuations is relatively small (0.057).[38] There may be something about these counties in the 1930's and 1940's that limited adjustment, and later analysis will explore the possibility of credit constraints and other factors.

Crops vs. Animals. Higher quality land is generally thought to have a comparative advantage for growing crops, as opposed to raising animals. There is a reasonable *a priori* expectation that erosion from the Dust Bowl would reduce the productivity of land for crops by more than for animals. Thus, we might expect long-run conversion of land from growing

---

[37] These estimates do not reflect an increase in abandoned farmland, proxied by the difference between total farmland and cropland or land in pasture. Relative changes in this measure are never greater than 1% of the total land in farms, in either direction.

[38] This is calculated by first-differencing county farmland shares from 1945 to 1992, and taking the standard deviation of these differences.

crops to raising animals. Indeed, these arguments were made strongly by contemporaries, and this land conversion was a major goal of government policy, advocated as both good for the farmer and good for society.[39]

To explore this potential margin of adjustment, it is possible to estimate relative changes in the productivity of land for crops and animals. This has three main limitations. First, changes in relative productivity need not imply changes in relative profitability if input quantities or prices also change. Second, it is only possible to estimate changes in the productivity of the average unit of land, whereas adjustment incentives depend on the marginal unit of land. Third, if land allocations adjust, then changes in land composition would be confounded with productivity changes.

Crop productivity is defined as the total value of crops sold, divided by acres of cropland. Animal productivity is defined as the total value of animals sold and animal products sold, divided by acres of pasture. These measures of productivity are problematic because some unknown portion of cropland would be used to grow crops that are fed to that farm's animals. The analysis would overstate the relative decline in crop productivity if farmers in higher erosion counties became more prone to use cropland to feed their own animals.

Figure 7, panels A and B, present the relative changes in crop and animal productivity for high erosion and medium erosion counties, respectively.[40] Both productivities are normalized to zero in 1930, in order to allow a clearer comparison between the relative changes. For both high erosion and medium erosion counties, crop productivity declined more than animal productivity from 1930 to 1940. This relative decrease only persists for medium erosion counties.[41] Table 4, columns 2 and 3, report the corresponding numerical results. Crop

---

[39] The SCS in 1955 was still strongly advocating conversion of land from cropland to grassland (Allred and Nixon 1955).

[40] The crop and animal productivity regressions are weighted instead by 1930 levels of county cropland and pasture, respectively. This allows the estimates to be interpreted as the percent change for an average unit of that land.

[41] Note that there was no relative change in productivity for high erosion counties relative to medium erosion counties. This may reflect a particular functional form through which erosion affects crops. This is consistent with earlier results on land values and revenue, where the high vs. medium shock had more persistent costs – there may be fewer opportunities to adjust production.

productivity appears to have been more sensitive to erosion than animal productivity. The exception is for high erosion counties after 1945, in which animal productivity also declined substantially.

However, farmers did not begin systematically shifting land from cropland to pasture until the 1950's (figure 7 panel C and D, table 4 column 4). High and medium erosion counties both then shifted land. These estimates are consistent with adjustment costs declining in the long-run. Counterintuitively, the largest shift occurred in high erosion counties, for which there were not estimated relative productivity differences in later periods. This may reflect poor quality land being shifted to pasture, which would account for the accompanying decrease in pasture productivity. Alternatively, the quality of the empirical comparison may be deteriorating over time.

Wheat vs. Hay. A similar empirical exercise compares relative changes in the productivity of wheat and hay, and the relative allocation of land to each. These are the two most widely grown crops for which comparable data are available over a long period of time. There is also an *a priori* expectation that erosion would discourage the production of wheat, relative to hay. Wheat production is more sensitive to soil quality and more likely to cause erosion, whereas hay is cultivated grass. Hay is less sensitive to soil quality, less prone to cause erosion, and an input in the production of animals. In contemporaneous and later writings, it was generally recommended that farmers shift land from wheat to hay (or to native grasslands and pasture).

Productivity for each crop is defined as the total quantity produced divided by the total acreage harvested. Because acreage is only available for harvested land, complete crop failure would cause the analysis to understate declines in productivity. Many counties do not report wheat data for 1940, but the reason is unclear.[42] The same caveats as before apply: productivity may not reflect changes in profitability; changes in the average acre may not proxy for changes in the marginal acre; and compositional changes may affect productivity.

---

[42] Data are often unavailable in one of the analyzed periods, so restricting the analysis to a balanced panel reduces the sample size by roughly half.

Figure 8, panels A and B, present the relative changes in wheat and hay productivity for high erosion and medium erosion counties, respectively.[43] From 1930 to 1940, wheat productivity decreased substantially. In 1950, when both production and acreage data are again surveyed for hay, wheat productivity remains persistently less than hay productivity.[44] Columns 5 and 6 of table 4 report the numerical results. Wheat productivity recovers substantially after 1964, but it remains relatively lower than hay productivity.

Figure 8, panels Panels C and D, show that farmers did not begin systematically shifting land from wheat to hay until the 1950's or 1960's. Table 4, column 7, reports the numerical results. After 1964, there were substantial and statistically significant declines in the amount of land devoted to wheat, as a fraction of the total land devoted to wheat and hay. Much of this adjustment occurs after data is no longer available for cropland and pasture, so it is difficult to compare the amount of adjustment along each land-use margin.

Average Farm Size and Land Fallowing. Hansen and Libecap (2004) present evidence that the presence of small farms limited efforts to prevent wind erosion during the Dust Bowl, due to substantial externalities. If more-eroded counties remained more prone to wind erosion, one potential margin of adjustment would be the consolidation of land holdings to internalize these externalities. Indeed, estimates indicate that more-eroded counties experienced greater increases in average farm size during the 1940's (figure 9 panel A, table 4 column 8). However, there may be other reasons for farm sizes to adjust.

If increased farm sizes primarily reflected a desire to incorporate these externalities, there should be a corresponding increase in efforts to prevent erosion. Previous estimates show that cropland did not decrease relatively in more-eroded counties during the 1940's. As another method of preventing erosion, Hansen and Libecap (2004) analyze the fraction of cropland kept fallow. More-eroded counties, however, experienced relative decreases in

---

[43] The wheat and hay productivity regressions are weighted by 1930 acres of land devoted to that crop. The other crop and animal variables are omitted as controls, as they directly change along with wheat and hay.

[44] As in the case for crop and animal productivity, there is not much relative change between high and medium erosion counties.

fallowing during the 1940's (figure 9 panel A, table 4 column 9). It does not seem that more-eroded counties relatively adjusted production to prevent future erosion.

These estimates imply that more-eroded counties became more intensely cultivated from 1940 to 1945, similar to the estimated increases in capital inputs. This would cause the observed percent recovery in agricultural revenues to overstate the percent recovery in profits.

V.D Why Was Adjustment in Agricultural Production Limited?

This section explores three possible explanations for the observed slow and limited adjustment in agricultural production: credit constraints, low incentives for tenants to adjust, and government payment programs that might discourage adjustment.

Credit Constraints. During the Great Depression, access to credit may have been generally restricted. In particular, counties with higher erosion experienced substantial decreases in land values, so land-owning farmers would have lost potential collateral. Poorly performing local mortgages may have restricted banks' ability to lend. The most-eroded counties also experienced more bank failures during the 1930's (figure 10). Bank weakness and bank failures can lead to persistent decreases in the supply of credit, especially during the 1930's (Bernanke 1983, Calomiris and Mason 2003, Ashcraft 2005).

Without easy access to credit, it may have been difficult for farmers to adjust agricultural production in response to the Dust Bowl erosion. Raising animals or shifting crops would require the purchase of livestock and perhaps different machinery. Other local sectors would require capital to shift land from agriculture to industrial uses or residential housing.

To explore whether restricted access to credit might explain the limited adjustment in agricultural land-use, I estimate whether counties affected by the Dust Bowl adjusted more if they had more banks prior to the 1930's.[45] The estimated equation is a modified version of equation (2). It adds interaction terms between the log number of banks (B) at the end of 1928 and the percent of a county in high and medium erosion areas. For ease of

---

[45] Subsequent banking is potentially an outcome of the Dust Bowl. The results are similar when using the amount of deposits prior to the Dust Bowl.

interpretation, the log number of banks is normalized to have a mean of zero and a standard deviation of one. The specification also controls for the main effect from banks, and an interaction between banks and each pre-Dust Bowl outcome level to allow for differential pre-trends by banking levels. The coefficients on these terms are allowed to vary over each time period.

$$\begin{aligned} \text{(3)} \qquad Y_{ct} - Y_{c1930} &= \beta_{1t} M_c + \beta_{2t} H_c + \beta_{3t} B_c \\ &+ \beta_{4t} B_c \times M_c + \beta_{5t} B_c \times H_c \\ &+ \alpha_{st} + \theta_t X_c + \gamma_{1t} L_c + \gamma_{2t} B_c \times L_c + \epsilon_{ct}. \end{aligned}$$

The coefficients of interest are β_{4t} and β_{5t}. These indicate whether counties with more banks adjusted differently over each time period than similarly eroded counties with fewer banks.

The main identification assumption is that counties with more original banks had better access to credit during and after the Dust Bowl, but would otherwise have adjusted agricultural land-use similarly. Pre-Dust Bowl banking levels may be correlated with other county characteristics (credit demand, education levels, overall financial development), so it must be assumed that these other characteristics do not predict differential responses to erosion.

Table 5, columns 1 to 3, presents the results. More-eroded counties with more banks immediately shifted toward greater pasture (column 1).[46] By contrast, there was no relative shift from wheat to hay (column 2). This could reflect increases in animal production requiring more credit than increases in hay production, due to large capital purchases of livestock. Estimated changes in total farmland are mixed (column 3). High erosion counties with more banks had greater declines in farmland over time, but medium erosion counties experienced some relative increases in farmland.

It is surprising that pre-1930's banking levels would continue to predict differential adjustment into the 1960's, presumably after the recovery of local credit markets. There may

---

[46] When estimating equation (4) without the interaction between banking and pre-Dust Bowl land allocation, there is no pre-trend toward greater pasture in more-eroded counties with more banks.

be lock-in effects once early adjustments differ, but the interaction effect is strengthening slightly. This highlights the possibility that banking levels may be correlated with other characteristics that predict adjustment to erosion.

Low Incentives for Tenants to Adjust. When land is rented with imperfect contracting, tenants may have inefficiently low incentives to make permanent investments in the land. Contemporaries emphasized that tenants' focus was on their crop rather than the land (McDonald 1938). Thus, the presence of tenant farmers may have reduced adjustments to restore land productivity or prevent future erosion.[47]

This hypothesis is tested by estimating equation (4), but replacing the number of banks with the share of farmland managed by tenants in 1930 (normalized to have a mean of zero and a standard deviation of one). Table 5, columns 4 to 6, present the results on the interaction term. There are few systematic patterns in the data, though there is some evidence of counties with less tenant farmland switching more from wheat to hay, which would also be less likely to worsen erosion.

For a narrow window around 1930 to 1950, there is data on tenants' land value, cropland, and equipment. Comparing counties with different erosion levels, tenants and non-tenants are estimated to have had similar changes in the per-acre value of farmland, the share of cropland in total farmland, and per-acre equipment. These estimates suggest that tenants may have had similar adjustment incentives.

If tenant contracts give inefficient adjustment incentives, then overall farming should adjust away from land being managed by tenants. Figure 11 presents the results from estimating equation (1) for the share of land farmed by tenants. Tenant shares did not decline immediately, but there is some indication of declines in the 1950's and after. This shift away from tenancy, along with the move toward larger farm sizes, could indicate that long-run adjustment in land-use required the reorganization of land ownership.

---

[47] Tenant farmers could also be more credit constrained, but this is not obvious: land owners would have lost substantial capital from lower land values, while tenants' assets might have been more liquid and allowed greater adjustment.

Government Payment Programs. There was a substantial increase in government payment programs during the 1930's, and many programs were targeted at the agricultural sector. To the extent that these payments were also targeted to higher erosion counties, they could have distorted farmers' incentives to adjust agricultural production. The New Deal agricultural payment programs tended to encourage the analyzed production adjustments, as well as the general continuation of farming.

Table 6, panel A, presents information on average payments for different New Deal programs, and the correlation between payments and county erosion levels, after controlling for state fixed effects and the usual county characteristics. There is little systematic relationship between program payments and erosion levels. It is possible, however, that farmers expected future payments.

Parts of the New Deal programs were temporary, but some parts were extended or changed into other forms of agricultural payments. Panel B of table 6 presents information on all government payments, beginning again in 1969. Higher erosion counties did not receive more payments in 1969, but began to in 1974. Total agricultural payments were much higher in 1987, but higher erosion counties received substantially less. This is likely due to the 1985 introduction of the Conservation Reserve Program (CRP), which paid farmers to take low-quality and erosion-prone land out of production. This precluded other farm payments. In 1992 and 1997, the higher erosion counties receive more CRP payments.

Overall, government payments were not targeted to more-eroded counties, so these programs seem unlikely to have discouraged agricultural adjustment to the Dust Bowl erosion. However, these estimates are consistent with a persistent drop in land quality. In the 1990's, a higher fraction of government payments were through the CRP, which targets low-quality erodible lands. These lands were still disproportionally in those counties more eroded during the Dust Bowl.

V.E Adjustment in Population, Industry, and Wages

In response to the substantial costs of the Dust Bowl, there should be strong incentives to reallocate labor across locations and/or sectors. Given that the agricultural sector was slow to adjust, this may create extra pressure for adjustment along the available margin of labor. However, if the agricultural sector was slow to adjust away from labor-intensive activities, this persistence could slow the reallocation of labor.

Population. Panel A of figure 12 graphs changes in log population from estimating equation (1). Prior to the 1930's, population was increasing in high erosion counties relative to low erosion counties. In medium erosion counties, population increased relatively from 1910 to 1920 and then decreased from 1920 to 1930. Following the Dust Bowl, population decreased substantially in both high and medium erosion counties.

In order to control for pre-1930 differences in levels and trends, column 1 of table 7 presents the numerical results from estimating equation (2). During the 1930's, high erosion counties experienced an 8.1% decrease in population, relative to low erosion counties in the same state. Medium erosion counties had a 6.5% relative decrease in population. In both county types, population continued to decline through the 1950's. By comparison, state-wide populations decreased 3% - 4% in Oklahoma, Kansas, and Nebraska from 1930 to 1940.

To explore selection in which populations decreased, panel B of figure 12 graphs changes in the fraction of population in rural areas of the county (areas with fewer than 2,500 inhabitants) and column 2 of table 7 presents the numerical results. The fraction of population in rural areas was mostly unchanged, indicating that decreases in population did not occur predominately in rural areas. Similarly, estimates in column 3 of table 7 show that the decrease in population was not especially among those living on farms.[48] There was sometimes a smaller decrease among farm populations, which may otherwise be less likely to migrate. The widespread population decrease in these counties may also indicate important spillover effects between agriculture and industry, perhaps through locally supplied inputs.

---

[48] Data on farm populations are not available before 1930.

Historical survey data provide some evidence on the selection of Dust Bowl migrants. Janow and McEntire (1940) describe results from a special survey of 1930's migrants to California, many of whom came from Oklahoma.[49] Comparing the distribution of migrants' original occupations in Oklahoma to the distribution of all Oklahoma residents occupations in 1930, there is only a small tendency for migration from the agricultural sector.[50] By contrast, migrants from other regions were less likely to be from the agricultural sector. The Dust Bowl may have shifted the typical selection of migrants toward the agricultural sector, but not much beyond proportional levels.

Unemployment and Manufacturing Employment. From 1930 to 1940, the overall unemployment rate increased by 0.71 percentage points in high erosion counties (column 1, table 8), and this increase was gone by 1950. Medium erosion counties had no relative increase in unemployment by 1940. These estimates suggest that surplus labor was reallocated relatively quickly.

For a restricted sample of counties, it is possible to analyze within-county shifts in labor toward the manufacturing sector. I focus on the manufacturing sector, instead of wholesale and retail, because it would tend to be less directly affected by local changes in demand. However, there may be important spillovers between agriculture and manufacturing, as suggested by the widespread decrease in population.

Column 2 of table 8 reports changes in the fraction of the labor force employed in manufacturing.[51] Column 3 reports changes in the fraction of the population employed in manufacturing. In high erosion counties, there were small increases from 1930 to 1940 in the

---

[49] The survey was conducted by the Bureau of Agricultural Economics, in cooperation with the California Department of Education. They surveyed families with school children that had migrated to California in the 1930's. This survey covered 116,000 families with 187,000 school children, which the surveyors estimated to be 84% of all such families. The results are described by Janow and McEntire (1940). I have not found the original data, but have converted the published figures back to aggregated numerical data when possible.

[50] Comparing the fractions in each occupation for migrant male heads of families vs. all employed males: 31% vs. 28% farmer; 19% vs. 15% farm laborer; 13% vs. 6% semi-skilled laborer; 12% vs. 12% skilled laborers; 12% vs. 15% other laborer; 5% vs. 10% clerks; 5% vs. 8% owner/manager; 1% vs. 2% servant; 0.5% vs. 3% professional.

[51] The labor force is defined as all employed workers, laidoff workers, and unemployed workers searching for a job.

proportion of workers or population employed in manufacturing, though these are not statistically significant. These estimates represent large percent increases in manufacturing employment of 11% and 15%, but they do not account for much overall movement of labor because manufacturing was a small sector of the economy. In medium erosion counties, there was no immediate shift in labor, but there were some later increases.

Columns 4 and 5 of table 8 report that total manufacturing establishments and value added did not increase following the Dust Bowl, though the coefficients are imprecisely estimated. Manufacturing may have been too slow to expand, perhaps due to the Depression, to attract workers before they left the county. Even after the Depression, there was no increase in manufacturing and the reallocation of labor continued through population declines. This overall pattern of results is similar to that found by Blanchard and Katz (1992) for state-level responses to a labor demand shock in the second half of the 20th century.

Proxies for Wages. Completing the empirical picture of labor market adjustment requires knowledge of changes in wage rates. Micro data on changes in wages are not available, but per-capita retail sales may serve as a proxy for local income (Fishback et al. 2005). This measure would differ from wages if there is net savings or borrowing, and would more closely reflect income if individuals' labor supply changed.

Column 1 of table 9 reports that high erosion and medium erosion counties experienced a 9.8% and 6.3% decrease in per-capita retail sales from 1930 to 1940. If this were the decrease in total income, it would be partly offset by lower land prices. A rough calculation indicates that individuals spent about 1% of their budget on land for housing.[52] Decreased land values would then compensate workers for a decrease of 2.78% and 1.67% in high and medium erosion counties, respectively.

Per-capita retail sales recovered from 1940 to 1958, partially in high erosion counties and completely in medium erosion counties. If we assume that income or wages recovered by a

---

[52] Assume that annual rental rates on residential land are 7% of the average value of one acre of agricultural land. The budget share is found by dividing this number by average per-capita retail sales. The number is similar whether unweighted or weighting by farmland or population.

similar percent, it is useful to check whether the magnitude of recovery would require relative changes in labor demand or if it could be explained by the observed decreases in population over the same time period. Reinterpreting estimates by Borjas (2003) on immigration would imply that a 10% decrease in population would increase wages by 3% to 4%, or increase income by 6.4% if individuals' labor supply is allowed to increase. In high erosion counties, the population decrease from 1940 to 1960 implies a 2% to 2.7% increase in wages and a 4.3% increase in income – compared to an observed 4% recovery in per-capita retail sales from 1940 to 1958. In medium erosion counties, the population decrease implies a 3.2% to 4.2% increase in wages and a 6.7% increase in income – compared to a 6.8% recovery in per-capita retail sales.

For a restricted sample of counties, additional data is available on total payroll and workers in the manufacturing, wholesale, and retail sectors. Columns 2 to 4 of table 9 report estimated changes in the log of payroll divided by workers, as a proxy for salary income. In the retail sector, for which most counties are available, there is an insignificant immediate decline in salary and the long-run declines are similar or more negative than for per-capita retail sales. In the wholesale and manufacturing sectors, there are greater long-run declines in salaries. Given the large decreases in population, this may reflect changes in the underlying skill distribution of the workforce. Otherwise, relative wages are at their lowest only when relative changes in population cease.

Interactions with Agriculture. Changes in the labor market potentially interact with adjustment incentives in the agricultural sector. There is some evidence that wages decreased in more-eroded counties, and this might be especially true in the agricultural sector if switching sectors is costly.

Consistent with temporarily lower wages, estimates indicate that the ratio of labor to capital machinery increased temporarily in the agricultural sector. This ratio is approximated by dividing the number of people living on farms by the value of equipment and machinery on farms. Changes in the log of this ratio are estimated by equation (2). For high erosion counties, the changes (and standard error) after 1930 are: 14% (4.8%) by 1940, 2% (4.8%) by 1945, and 9% (5.5%) by 1969. For medium erosion counties, the changes after 1930 are: 13% (4.1%) by 1940, 4.5% (4.1%) by 1945, and 2.4% (4.3%) by 1969.

Despite the immediate decreases in population, there may still have been surplus labor in 1940. This is reflected in somewhat higher unemployment and the subsequent decreases in population. Some of this surplus labor appears to have been absorbed in the agricultural sector. This may have discouraged switching land from crops to animals or from wheat to hay, which use less labor. It was in the 1950's, as population declines ceased, that agricultural adjustment began to appear along these margins.

VI Conclusion

The 1930's American Dust Bowl provides a unique context in which to examine economic adjustment to a permanent change in the environment. Dust Bowl erosion imposed substantial short-run costs on Plains agriculture. Adjustment in agricultural land-use was slow and limited, despite evidence of potentially productive adjustments. Estimates imply that the long-run cost to agriculture was 86% of the short-run cost.

In other settings where short-run production costs from a shock are substantial, slow and limited adjustment would imply that the long-run production costs may also be large. This paper provides some evidence on the role of credit and land ownership in predicting adjustment, and an important area for further research is understanding what conditions facilitate long-run adjustment. Further research on historical shocks may help anticipate the long-run effects of modern shocks for which only short-run effects can be observed.

If long-run production costs are large, the overall welfare consequences may be partly lessened if people can leave the area. The Unites States is known to have high labor mobility, so migration may be more limited in response to shocks in other countries. In particular, cross-country migration may be severely restricted by political constraints. If short-run production costs persist and migration is limited, shocks that have large short-run costs may impose similarly large long-run costs.

References

Allred, B.W. and W.M. Nixon. 1955. "Grass for Conservation in the Southern Great Plains." Farmer's Bulletin No. 2093.

Ashcraft, Adam. 2005. "Are Banks Really Special? New Evidence from FDIC-induced Failures of Healthy Banks." *American Economic Review,* Vol. 95, No. 5, pp. 1712-1730.

Bernanke, Ben S. 1983. "Nonmonetary Effects of the Financial Crisis in the Propagation of the Great Depression." *American Economic Review,* Vol. 73, No. 3, pp. 257-276.

Blanchard, Olivier Jean, and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity,* Vol. 1992, No. 1, pp. 1-61.

Bloom, Nick, Stephen Bond, and John Van Reenen. 2007. "Uncertainty and Investment Dynamics." *Review of Economic Studies,* Vol. 74, pp. 391-415.

Borjas, George J. 2003. "The Labor Demand Curve *is* Downward Sloping: Reexamining the Impact of Immigration on the Labor Market." *Quarterly Journal of Economics,* Vol 118, No. 4, pp. 1335-1374.

Boustan, Leah Platt, Price V. Fishback, and Shawn E. Kantor. 2007. "The Effect of Internal Migration on Local Labor Markets: American Cities During the Great Depression." NBER WP#13276.

Calomiris, Charles W., and Joseph R. Mason. 2003. "Consequences of Bank Distress during the Great Depression." *American Economic Review,* Vol. 93, No. 3, pp. 937-947.

Carrington, William J. 1996. "The Alaskan Labor Market during the Pipeline Era." Journal of Political Economy, Vol. 104, No. 1, pp. 186-218.

Chari, V. V. and Hugo Hopenhayn. 1991. "Vintage Human Capital, Growth, and the Diffusion of New Technology." *Journal of Political Economy,* Vol. 99, No. 6, pp. 1142-1165.

Chepil, W.S., F. H. Siddoway, and Dean V. Armbrust. 1962. "Climatic Factor for Estimating Wind Erodibility of Farm Fields." *Journal of Soil and Water Conservation,* Vol. 17, pp. 162 – 165.

Conley, T.G. 1999. "GMM estimation with cross sectional dependence." *Journal of Econometrics,* Vol. 92, pp. 1-45.

Cunfer, Geoff. 2005. <u>On the Great Plains: Agriculture and Environment</u>. Texas A&M University Press, College Station.

Davis, Donald R., and David E. Weinstein. 2002. "Bones, Bombs, and Break Points: The Geography of Economic Activity." *American Economic Review,* Vol. 92, No. 5, pp. 1269-1289.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. 2008. "Climate Change and Economic Growth: Evidence from the Last Half Century." NBER WP#14132.

Deschenes, Olivier, and Michael Greenstone. 2007. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather." *American Economic Review,* Vol. 97, No. 1, pp. 354-385.

Diamond, Jared. 2005. <u>Collapse: How Societies Choose to Fail or Succeed</u>.

Dixit, A. and R. Pindyck. 1994. <u>Investment Under Uncertainty</u>. Princeton University Press. Princeton, New Jersey.

Duflo, Esther. 2004. "The medium run effects of educational expansion: evidence from a large school construction program in Indonesia." *Journal of Development Economics,* Vol. 74, pp. 163-197.

Fishback, Price V., William Horrace, and Shawn Kantor. 2005. "Did New Deal Grant Programs Stimulate Local Economies? A Study of Federal Grants and Retail Sales During the Great Depression." *Journal of Economic History,* Vol. 65, No. 1, pp. 36-71.

Foster, Andrew D., and Mark R. Rosenzweig. 1995. "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture." *Journal of Political Economy,* Vol. 103, No. 6, pp. 1176-1209.

Griliches, Zvi. 1957. "Hybrid Corn: An Exploration in the Economics of Technical Change." *Econometrica,* Vol. 25, No. 4, pp. 501-522.

Guiteras, Raymond. 2008. "The Impact of Climate Change on Indian Agriculture." Mimeo.

Hansen, Zeynep K. 2005. "Book Review of On the Great Plains: Agriculture and Environment by Geoff Cunfer." *Journal of Economic History,* Vol. 65, No. 4, pp. 1162-1163.

Hansen, Zeynep K. and Gary D. Libecap. 2004. "Small Farms, Externalities, and the Dust Bowl of the 1930s." *Journal of Political Economy,* Vol. 112, No. 3, pp. 665 – 694.

Hornbeck, Richard. 2008. "Good Fences Make Good Neighbors: Evidence on the Effects of Property Rights." Mimeo.

Hoyt, John C. 1936. "Droughts of 1930 – 1934." Water-Supply Paper 680, U.S. GPO.

Hurt, Douglas R. 1981. The Dust Bowl: An Agricultural and Social History. Nelson Hall, Chicago.

Hurt, Douglas R. 1985. "The National Grasslands: Origin and Development in the Dust Bowl," in The History of Soil and Water Conservation, eds. Douglas Helms and Susan L. Flader. The Agricultural History Society, Washington D.C., pp. 144 – 157.

Janow, Seymour J., and Davis McEntire. 1940. "The Migrants: Migration to California." *Land Policy Review*, July-August, pp. 24-36.

Leuchtenburg, William E. 1963. Franklin D. Roosevelt and the New Deal, 1932 – 1940. Harper & Row, Publishers, New York.

Margo, Robert A. 1997. "Wages in California during the Gold Rush." NBER WP on Historical Factors in Long Run Growth, Historical Paper 101.

McDonald, Angus. 1938. "Erosion and its Control in Oklahoma Territory." Soil Conservation Service, Miscellaneous Publication No. 301, GPO.

Miguel, Edward, and Gerard Roland. 2006. "The Long Run Impact of Bombing Vietnam." Mimeo.

National Archives (College Park, MD), Record Group 114, Cartographic Records of the Soil Conservation Service, Entry #149.

News-week. 1936. "Drought: A Merciless Sun and a Scourge of Insects Destroy Crops, Cattle and Men – Two-thirds of the Country Afflicted." July 18.

Phillips, Sarah T. 2007. This Land, This Nation: Conservation, Rural America, and the New Deal. Cambridge University Press.

Pressly, Thomas J. and William H. Scofield. 1965. Farm Real Estate Values in the United States by Counties, 1850 – 1959. University of Washington Press, Seattle.

Rasmussen, Wayne D. and Gladys L. Baker. "Price-Support and Adjustment Programs From 1933 Though 1978: A Short History." Agricultural Information Bulletin No. 424, U.S. Department of Agriculture, U.S. Government Printing Office.

Redding, Stephen J., and Daniel M. Sturm. Forthcoming. "The Costs of Remoteness: Evidence from German Division and Reunification." *American Economic Review.*

Roback, Jennifer. 1982. "Wages, Rents, and the Quality of Life." *Journal of Political Economy*, Vol. 90, No. 6, pp. 1257 – 1278.

Rosenzweig, Mark R., and Kenneth I. Wolpin. 2000. "Natural 'Natural Experiments' in Economics." *Journal of Economic Literature*, Vol. 38, No. 4, pp. 827-874.

Saloutos, Theodore. 1982. The American Farmer and the New Deal. Iowa State University Press.

Science. 1934. "The Recent Destructive Dust Cloud." New Series, Vol. 79, May 25, p. 473.

Soil Conservation Service. 1935. "Soil Erosion: A Critical Problem in American Agriculture." Part V of the Supplementary Report of the Land Planning Committee to the National Resources Board. US GPO.

Soil Conservation Service. 1937. "Soil Conservation Districts for Erosion Control." Miscellaneous Publication No. 293, GPO.

Soil Conservation Service. 1939. "Land Facts on the Southern Plains." Miscellaneous Publication No. 334, GPO.

Soil Conservation Service. 1955. "Facts About Wind Erosion and Dust Storms on the Great Plains." Leaflet No. 394. US Department of Agriculture.

US Department of Agriculture. 1940. "The Dust Bowl: Agricultural Problems and Solutions," Editorial Reference Series, No. 7, July 15.

Wallace, Henry A. 1938. "Forward" in Soils and Men: Yearbook of Agriculture. US Department of Agriculture, Washington D.C., Government Printing Office.

Worster, Donald. 1979. Dust Bowl: The Southern Plains in the 1930s. Oxford University Press, Oxford.

Figure 1. Aggregate Changes on the Plains in Agriculture and Population

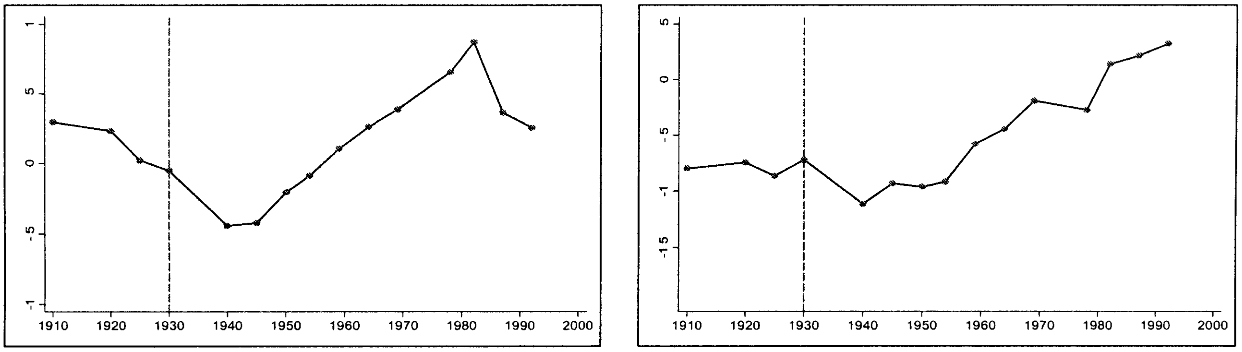
A. Farmland Value, per farm acre (Log/CPI)

B. Agricultural Revenue, per farm acre (Log/PPI)

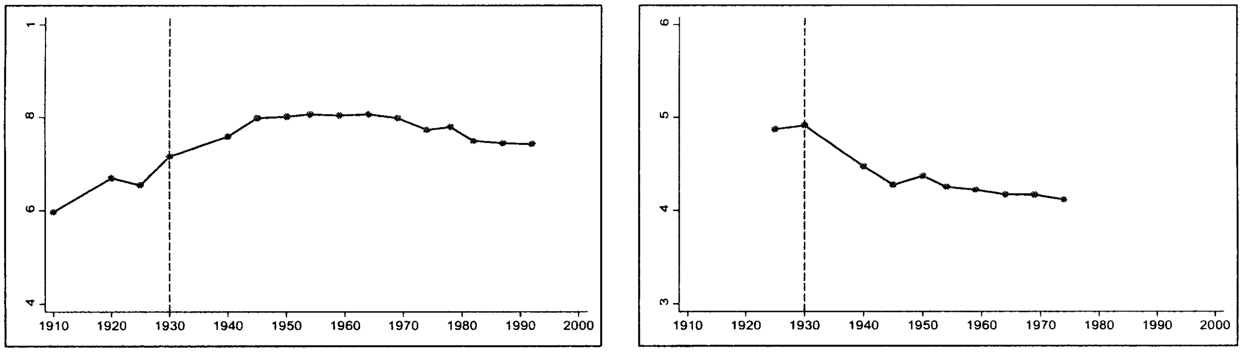
C. Farmland Acres, per county acre

D. Cropland Acres, per farm acre

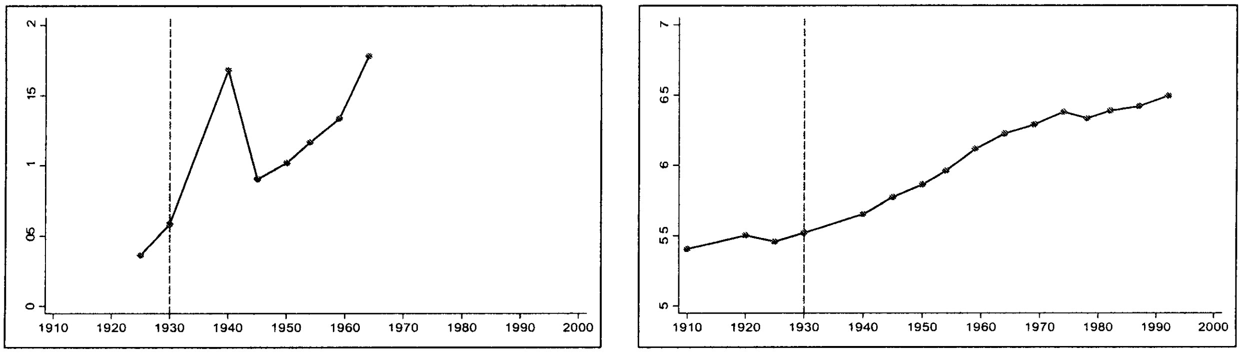
E. Fallow Acres, per cropland acre

F. Average Farm Size (Log)

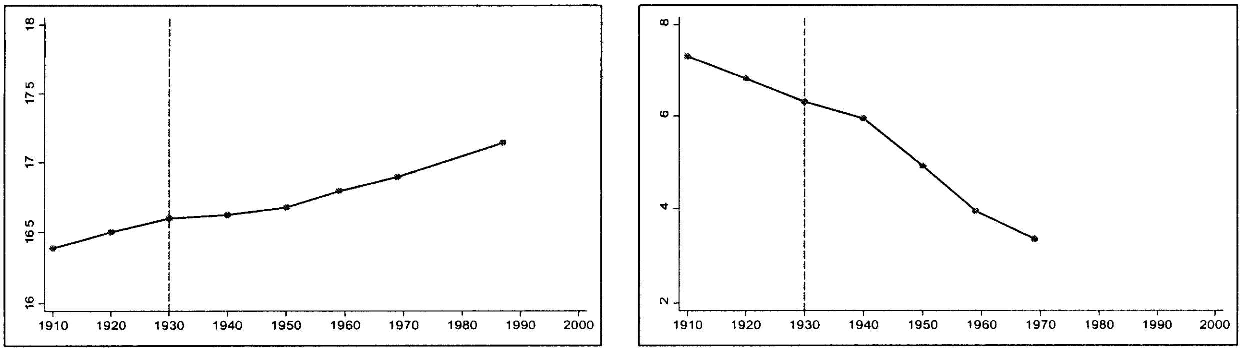
G. Population (Log)

H. % Population in Rural Areas

Notes: All panels (except E) report values aggregated over the same 769 Plains counties in every period for the following outcomes: Panel A, the log per-acre value of farmland divided by the consumer price index; Panel B, the log per-acre agricultural revenues divided by the producer price index for farm products; Panel C, the fraction of county land in farms; Panel D, the fraction of farmland that is cropland; Panel E, the fraction of cropland that is fallow (available for only 587 counties); Panel F, the log acres of farmland divided by the number of farms; Panel G, the log population; Panel H, the fraction of population living in areas with fewer than 2,500 inhabitants.

Figure 2. Sample Counties, Shaded by Erosion Level

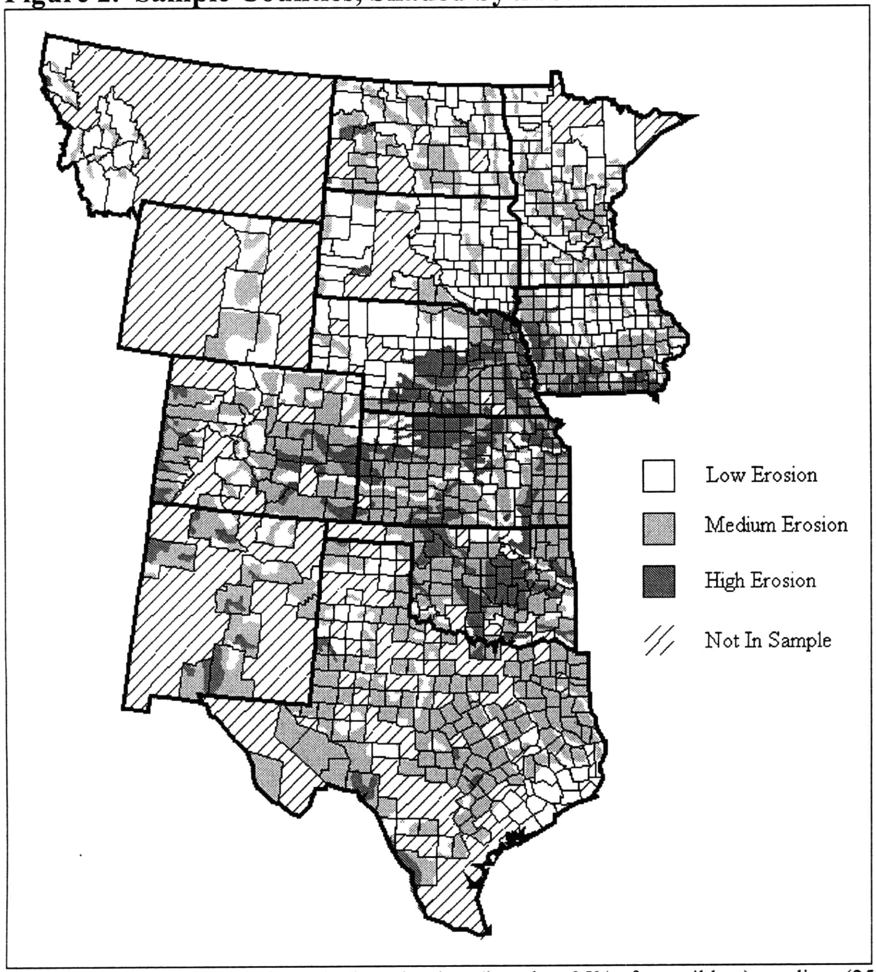

Notes: Mapped erosion levels are indicated as low (less than 25% of topsoil lost), medium (25%-75% of topsoil lost and may have some gullies), or high (more than 75% of topsoil lost and may have numerous or deep gullies). Thin lines denote 1910 county borders, corresponding to the sample of 769 counties described in Table 1. Thick lines denote state boundaries. Crossed out areas are not in the sample.

Figure 3. Log Changes in the per-Acre Value of Agricultural Land, by Erosion Level

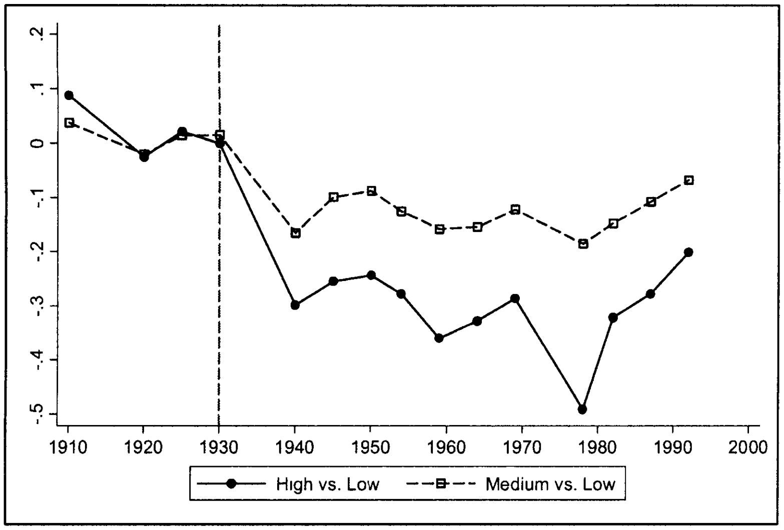

Notes: This figure graphs the estimated coefficients (β) from equation (1) in the text. For each year, the solid circles report differences in the log per-acre value of farmland for high erosion counties, relative to low erosion counties. The hollow squares report differences for medium erosion counties, relative to low erosion counties. These coefficients are estimated by regressing the per-acre value of farmland on the percent of a county in a high erosion area (solid circle) and the percent of a county in a medium erosion area (hollow square), controlling also for state-by-year fixed effects, the interaction between each year and each county characteristic in panels B – E of table 1, and the interaction between each year and the available lagged values of each county characteristic in panels B – E of table 1.

Figure 4. Log Changes in Agricultural Revenue and Capital, by Erosion Level

A. Agricultural Revenue, per Acre of Farmland

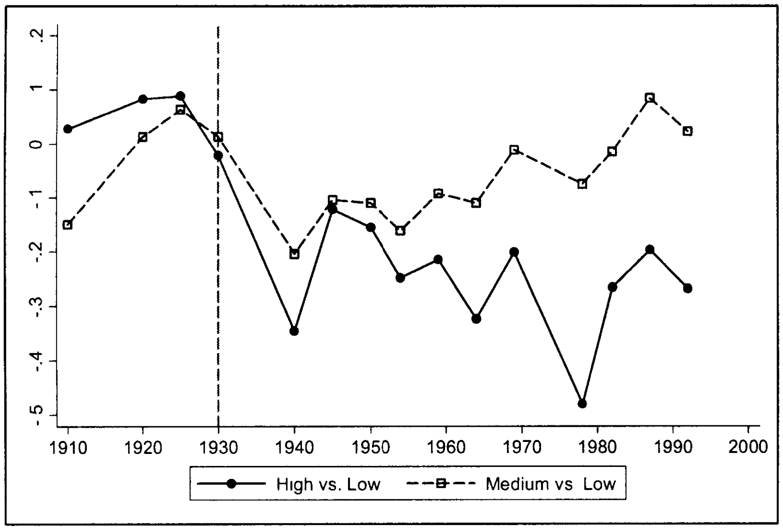

B. Value of Capital (Machinery & Equipment), per Acre of Farmland

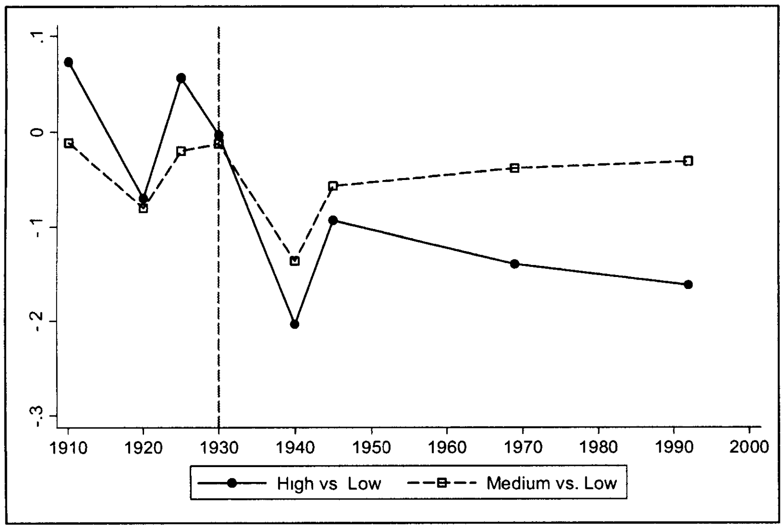

Notes: For the indicated outcome variable, each panel graphs the estimated coefficients (β) from equation (1) in the text, as described in the notes to Figure 3.

Figure 5. Log Changes in Land Value, Plains Counties vs. Non-Plains Counties

A. High – Low Erosion

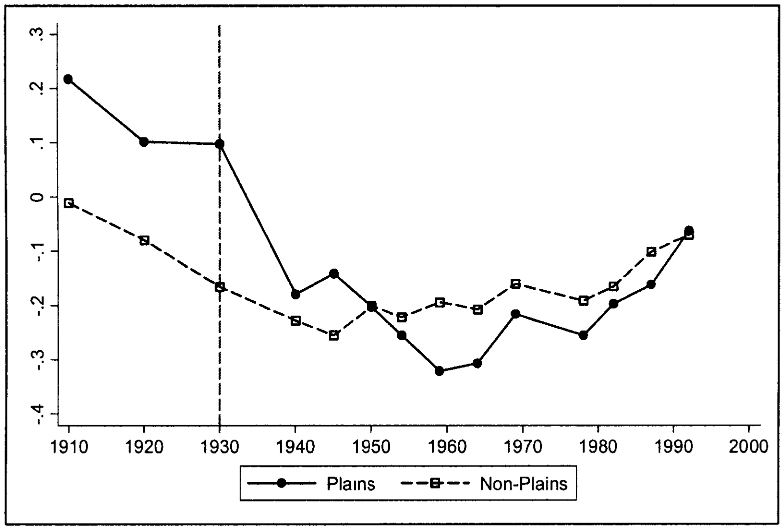

B. Medium – Low Erosion

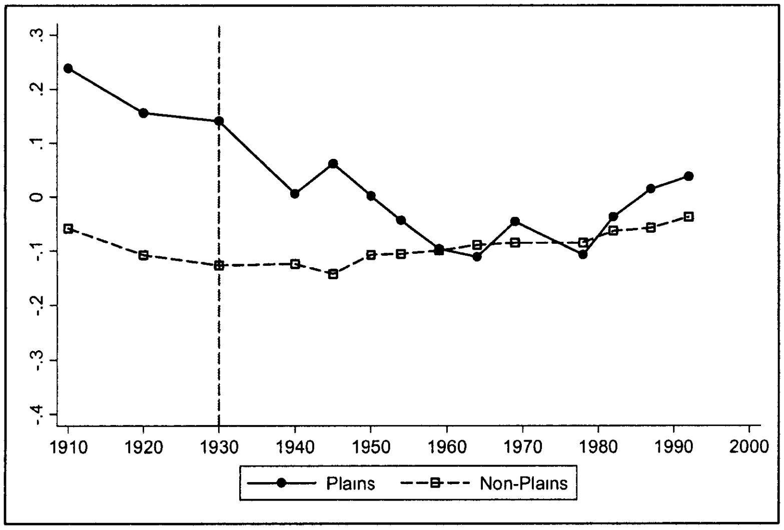

Notes: The estimates in both panels are from a single regression of equation (1) in the text, modified to include only state-by-time fixed effects as controls. Panel A reports the difference between high and low erosion counties, for both Plains and non-Plains states Panel B reports the difference between medium and low erosion counties, for both Plains and non-Plains states. Plains states are those depicted in Figure 2. The sample includes 867 Plains counties and 1840 non-Plains counties.

Figure 6. Changes in Farmland Acres per county acre, by Erosion Level

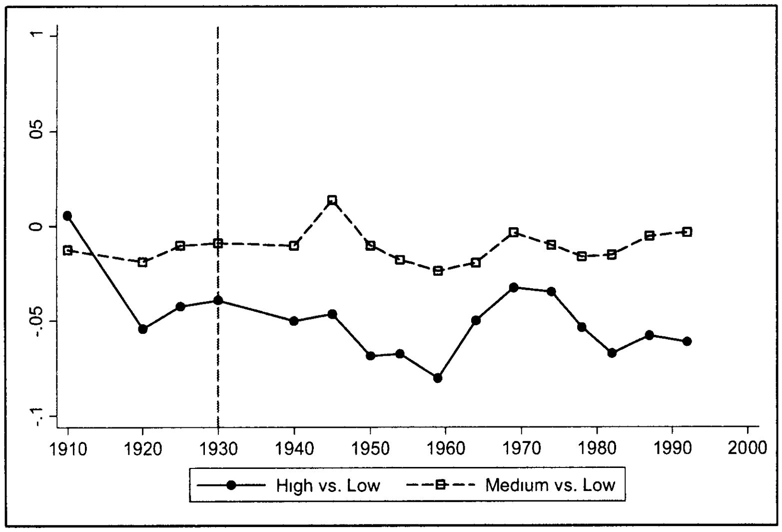

Notes. This figure graphs the estimated coefficients (β) from equation (1) in the text, as described in the notes to Figure 3. Omitted as controls are the 1930 and lagged values of farmland per county acre.

Figure 7. Log Changes in Crop & Animal Productivity, and Changes in Cropland as a Fraction of Cropland and Pasture

A. Crop and Animal Productivity: High – Low Erosion

B. Crop and Animal Productivity: Medium – Low Erosion

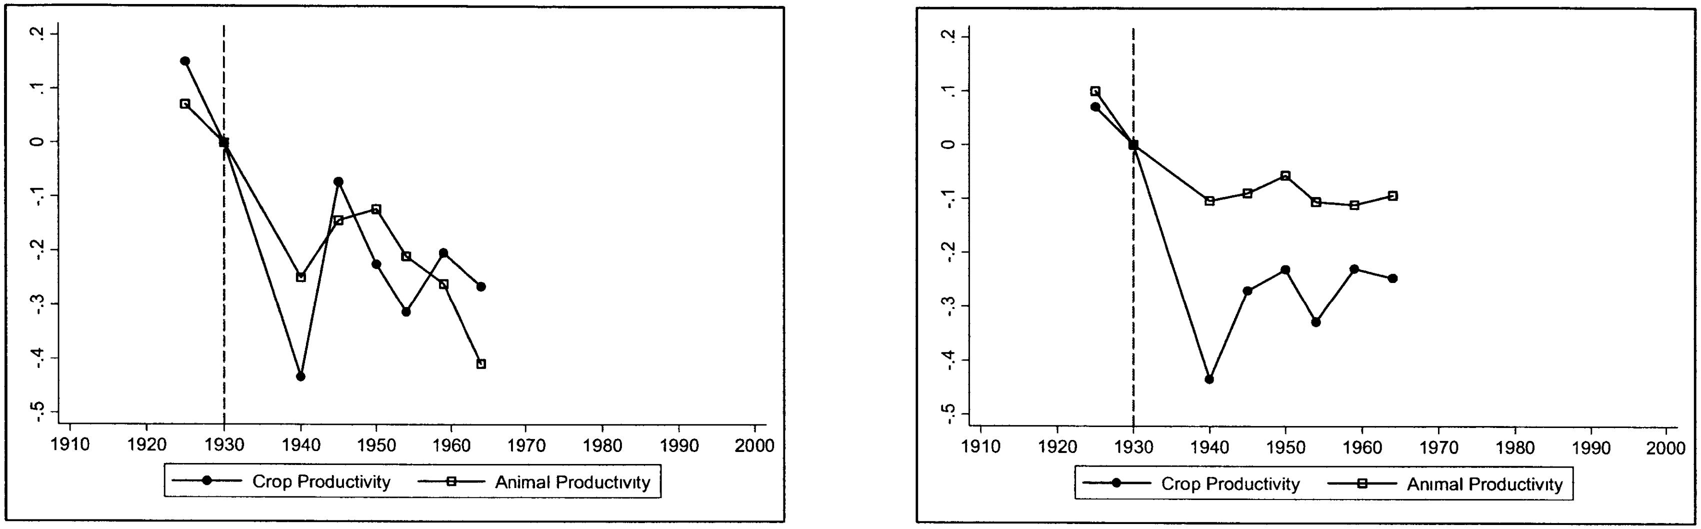

C. Allocation of Land to Crops: High – Low Erosion

D. Allocation of Land to Crops: Medium – Low Erosion

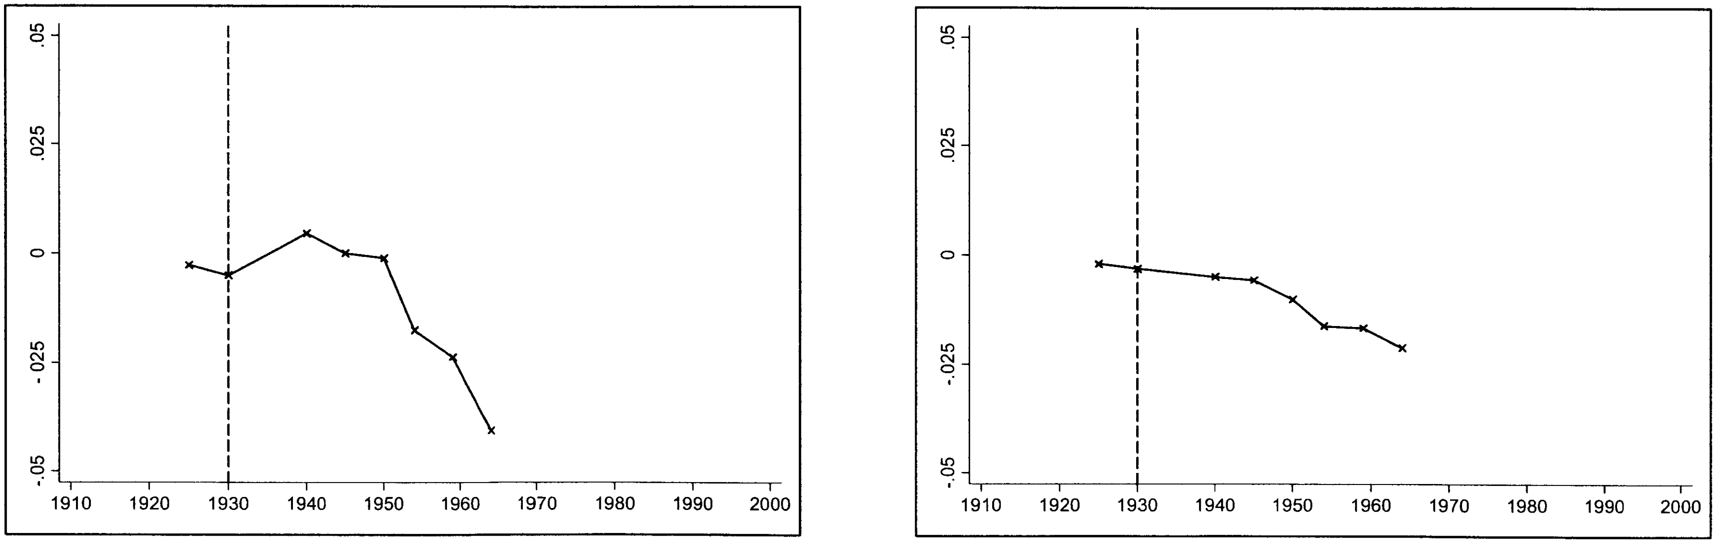

Notes: Panels A and C graph coefficients from equation (1), estimated separately for crop and animal productivity. Panel A reports the change in high erosion relative to low erosion, Panel B reports changes in medium relative to low erosion (normalized to zero in 1930). Panels C and D graph estimated changes in the relative fraction of land allocated to crops.

Figure 8. Log Changes in Wheat & Hay Productivity, and Changes in Wheat Land as a Fraction of Land in Wheat and Hay

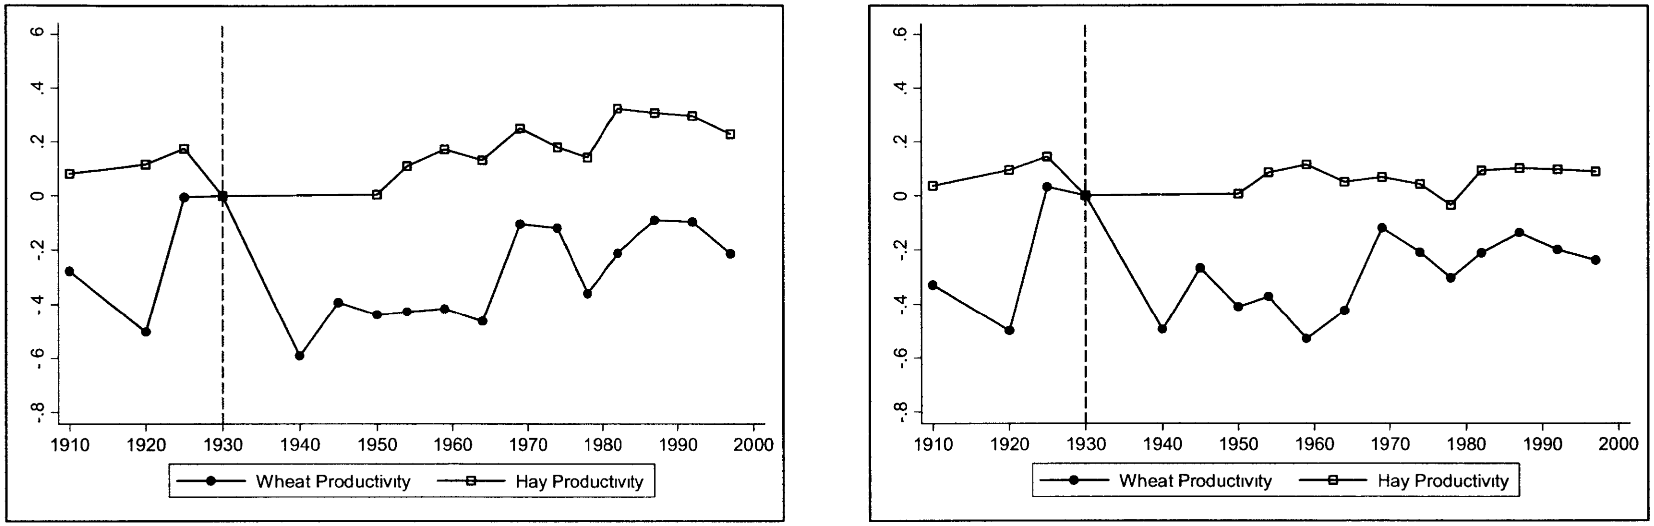

A. Wheat and Hay Productivity: High – Low Erosion

B. Wheat and Hay Productivity: Medium – Low Erosion

C. Allocation of Land to Wheat: High – Low Erosion

D. Allocation of Land to Wheat: Medium – Low Erosion

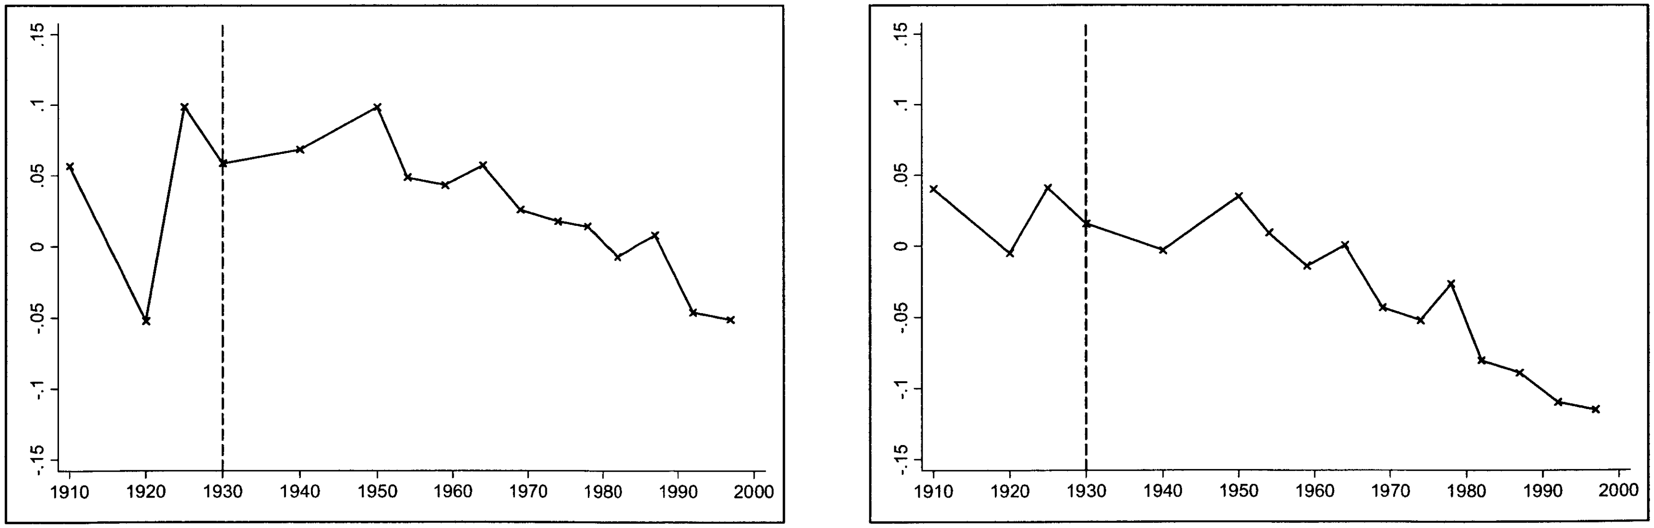

Notes: Panels A and C graph coefficients from equation (1), estimated separately for wheat and hay productivity. Panel A reports the change in high erosion relative to low erosion, Panel B reports changes in medium relative to low erosion (normalized to zero in 1930). Panels C and D graph estimated changes in the relative fraction of land allocated to wheat.

Figure 9. Changes in Average Farm Size and Land Fallowing, by Erosion Level

A. Log Average Farm Size

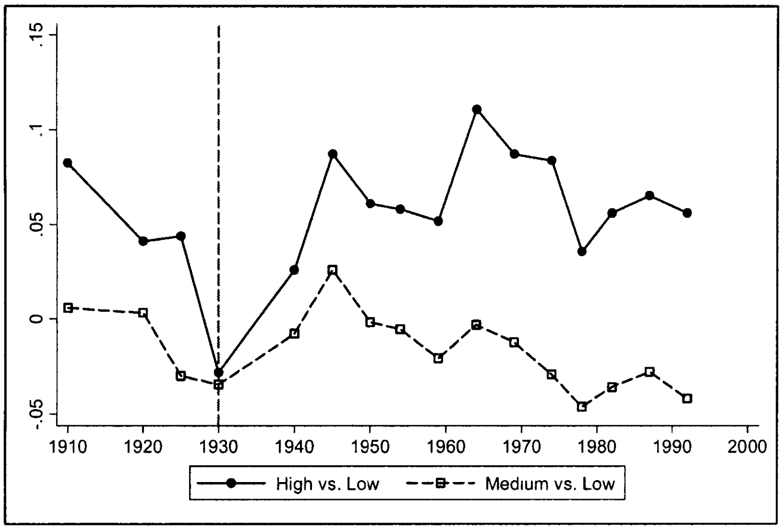

B. Acres of Fallow Land, per acre of cropland

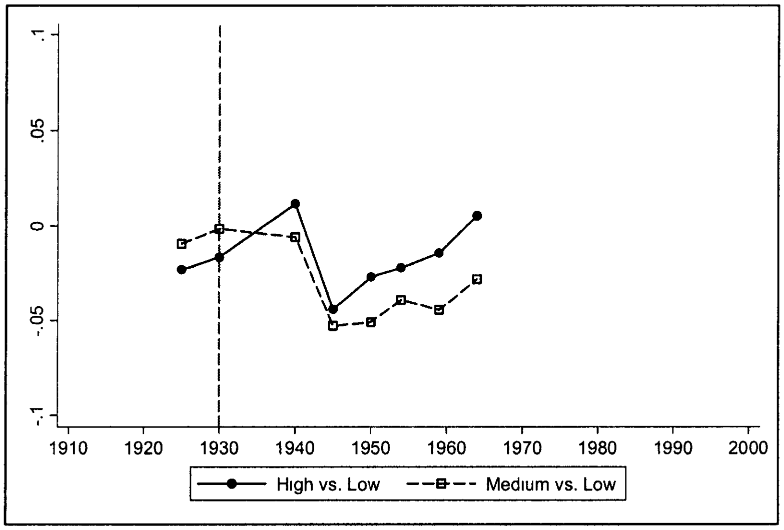

Notes: For the indicated outcome variable, each panel graphs the estimated coefficients (β) from equation (1) in the text, as described in the notes to Figure 3. Panel A omits 1930 and lagged values of average farm size as a control variable.

Figure 10. Log Changes in the Number of Active Banks, by Erosion Level

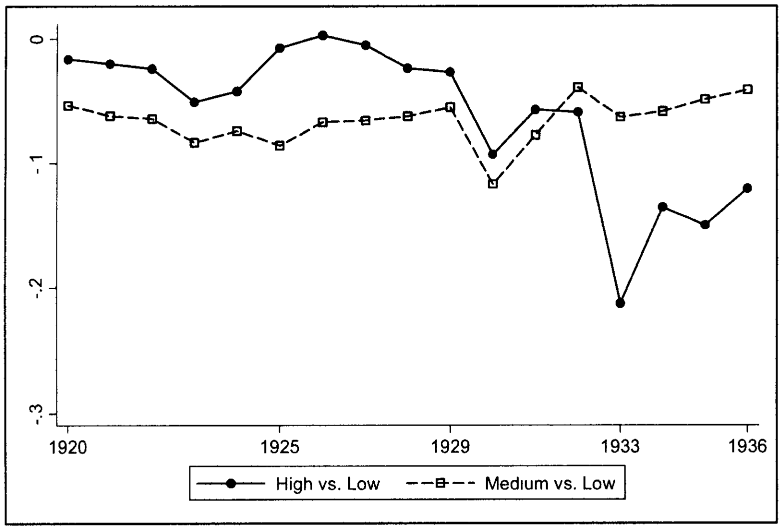

Notes: This figure graphs the estimated coefficients (β) from a modified version of equation (1), described in the notes to Figure 3. Banking data is available annually, so the regression controls for state-by-year fixed effects and county characteristics interacted with every year.

Figure 11. Changes in the Tenant Share of Farmland, by Erosion Level

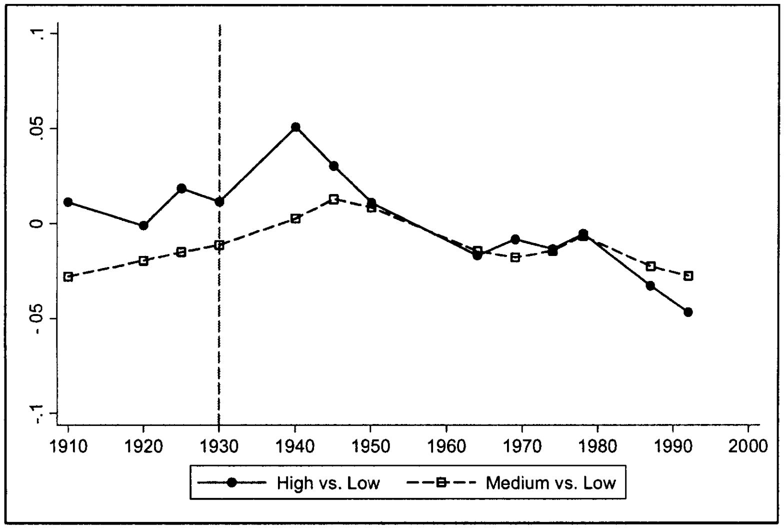

Notes: This figure graphs the estimated coefficients (β) from equation (1) in the text, as described in the notes to Figure 3.

Figure 12. Changes in Population and Rural Population, by Erosion Level
A. Log Population

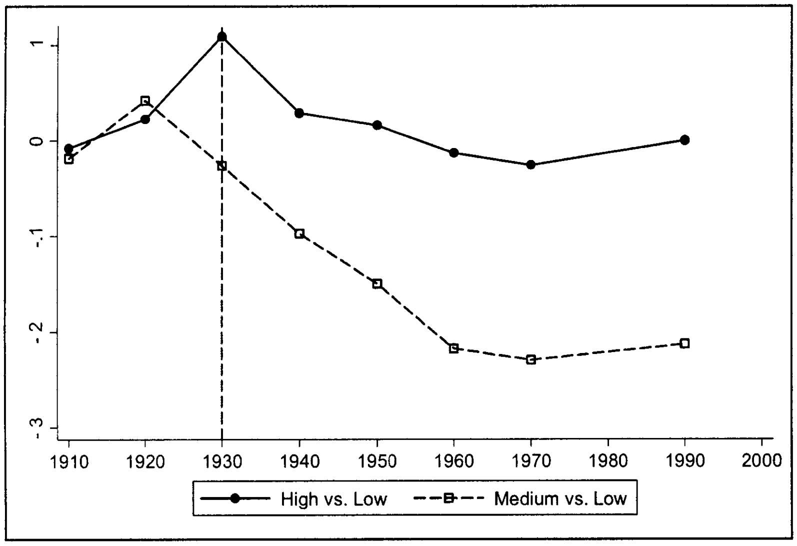

B. Fraction of Population in Rural Areas of County

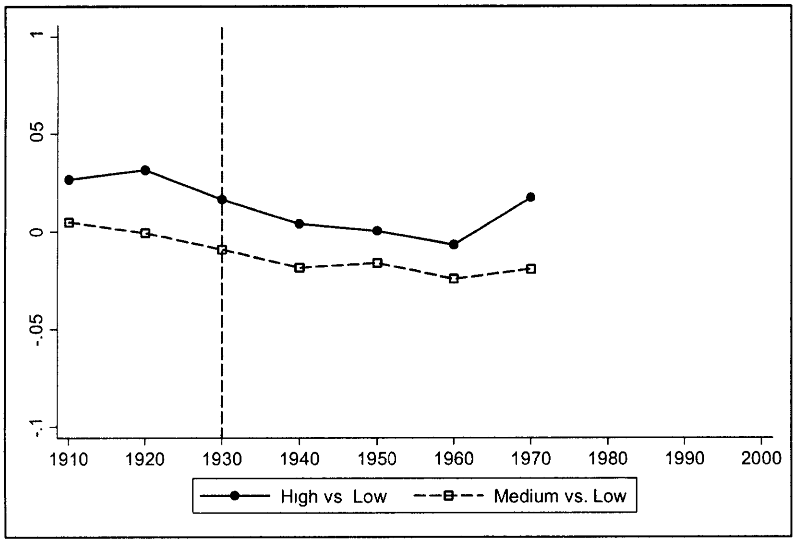

Notes: For the indicated outcome variable, each panel graphs the estimated coefficients (β) from equation (1) in the text, as described in the notes to Figure 3. In panel A, the 1930 and lagged values of population per acre are omitted as controls. In panel B, omitted controls are the 1930 and lagged values of fraction of population in rural areas.

Table 1. County Characteristics in 1930, by Post-Dust Bowl Erosion Level

		Relative to Low Erosion:		
	All Counties (1)	Medium Erosion (2)	High Erosion (3)	(3) – (2) (4)
Panel A: Value				
Value of farm land & buildings,	3.527	0.244*	0.225	-0.019
per acre in farms	[0.853]	(0.114)	(0.154)	(0.106)
Value of farm land,	3.309	0.217	0.202	-0.014
per acre in farms	[0.819]	(0.111)	(0.150)	(0.103)
Value of all farm products,	2.107	0.258**	0.202	-0.055
per acre in farms	[0.789]	(0.100)	(0.132)	(0.100)
Panel B: Land-use				
Acres of land in farms,	0.819	0.035	-0.004	-0.039
per county acre	[0.187]	(0.019)	(0.023)	(0.020)
Acres of cropland,	0.491	0.035	-0.008	-0.043
per acre in farms	[0.229]	(0.031)	(0.040)	(0.028)
Panel C: Population and Farms				
Population, per 100 county acres	3.524	1.289	1.262	-0.027
	[7.018]	(0.824)	(0.826)	(0.773)
Percent of population, rural areas	0.818	-0.022	0.034	0.057
	[0.221]	(0.030)	(0.043)	(0.042)
Percent of population, on farms	0.542	0.020	0.040	0.020
	[0.175]	(0.024)	(0.032)	(0.032)
Farms, per 100 county acres	0.328	0.163**	0.148**	-0.015
	[0.246]	(0.026)	(0.031)	(0.031)
Average farm size (acres)	613	-233	-307*	-74
	[1145]	(184)	(152)	(142)
Panel D: Cropland Allocation				
% Corn	0.188	0.070**	0.192**	0.122**
	[0.166]	(0.017)	(0.027)	(0.023)
% Wheat	0.181	-0.043	-0.102**	-0.059
	[0.214]	(0.027)	(0.035)	(0.034)
% Hay	0.157	-0.044	-0.099*	-0.056*
	[0.177]	(0.030)	(0.044)	(0.022)
% Cotton	0.104	0.053**	0.020	-0.033
	[0.209]	(0.017)	(0.019)	(0.019)
% Oats, Barley, and Rye	0.134	-0.005	-0.033*	-0.028
	[0.116]	(0.009)	(0.013)	(0.010)
Panel E: Animal Production				
Cows, per county acre	0.056	0.011**	0.013**	0.003
	[0.033]	(0.003)	(0.004)	(0.003)
Pigs, per county acre	0.066	0.040**	0.055**	0.015
	[0.096]	(0.008)	(0.012)	(0.011)
Chickens, per county acre	0.296	0.115**	0.116**	0.001
	[0.279]	(0.025)	(0.034)	(0.031)

Notes: Column 1 presents average values for the 769 sample counties, where counties are weighted by the acres of farmland in 1930 (counties based on 1910 borders; with at least 1000 acres of farmland in every period; and with data on each variable in Figure 1 for each period shown). The standard deviation is reported in brackets.

Columns 2 and 3 report coefficients from a single equation that regresses the indicated county characteristic in 1930 on the percentage of the county in a medium erosion area and in a high erosion area (low erosion is the omitted category), conditional on State fixed effects and weighted by acres of farmland in 1930. Column 4 reports the difference between the coefficients in columns 2 and 3. Robust standard errors are reported in parentheses. ** denotes statistical significance at the 1% level, * at the 5% level.

Table 2. Log Changes in Land Value and Revenue per farm acre, by Erosion Level

	Land Value		Revenue		
	Change After 1930 (1)	% Persisting After 1940 (2)	Change After 1930 (3)	% Persisting After 1940 (4)	Ratio of Change (1) / (3) (5)
Panel A. 1940					
High - Low	-0.278		-0.316		0.881
	(0.041)		(0.055)		(0.120)
Medium - Low	-0.167		-0.202		0.824
	(0.029)		(0.039)		(0.136)
High - Medium (calculated)	-0.111		-0.114		0.981
	(0.032)		(0.048)		(0.313)
Averaged Value (GLS)					0.861
					(0.112)
Panel B. 1945					
High - Low	-0.234	0.841	-0.122	0.386	
	(0.037)	(0.087)	(0.044)	(0.118)	
Medium - Low	-0.101	0.609	-0.131	0.648	
	(0.027)	(0.105)	(0.031)	(0.158)	
High - Medium (calculated)	-0.133	1.189	0.009	-0.081	
	(0.031)	(0.238)	(0.035)	(0.324)	
Averaged Value (GLS)		0.765		0.452	
		(0.080)		(0.114)	
Panel C. 1950 – 1954 (pooled)					
High - Low	-0.240	0.864	-0.208	0.658	
	(0.040)	(0.127)	(0.051)	(0.140)	
Medium - Low	-0.118	0.705	-0.183	0.904	
	(0.027)	(0.129)	(0.036)	(0.194)	
High - Medium (calculated)	-0.123	1.101	-0.025	0.220	
	(0.033)	(0.307)	(0.040)	(0.326)	
Averaged Value (GLS)		0.788		0.710	
		(0.112)		(0.135)	
Panel D. 1959 – 1969 (pooled)					
High - Low	-0.315	1.133	-0.255	0.807	
	(0.042)	(0.161)	(0.063)	(0.193)	
Medium - Low	-0.167	1.001	-0.130	0.643	
	(0.033)	(0.196)	(0.044)	(0.217)	
High - Medium (calculated)	-0.148	1.331	-0.125	1.099	
	(0.036)	(0.401)	(0.050)	(0.531)	
Averaged Value (GLS)		1.094		0.746	
		(0.150)		(0.178)	
Panel E. 1978 – 1992 (pooled)					
High - Low	-0.298	1.072	-0.323	1.021	
	(0.057)	(0.197)	(0.092)	(0.308)	
Medium - Low	-0.100	0.603	-0.056	0.276	
	(0.045)	(0.249)	(0.066)	(0.321)	
High - Medium (calculated)	-0.198	1.773	-0.267	2.347	
	(0.051)	(0.557)	(0.076)	(1.110)	
Averaged Value (GLS)		0.935		0.693	
		(0.184)		(0.272)	
R-squared	0.9496		0.8925		
Sample Counties	769		769		

Notes: Columns 1 and 3 report the results from estimating equation (2) in the text, for the value of land and revenue per-acre of farmland, weighting by county farmland in 1930. This estimates the change after 1930, controlling for the same variables as in Figure 3 as well as the values of both land values and revenue in 1930, 1925, 1920, and 1910. The indicated multiple year periods are pooled, such that the estimate is an average over those years. Columns 2 and 4 report the estimated change, as a proportion of the initial change from 1930 to 1940. Column 5 reports the change in land value as a proportion of the change in revenue from 1930 to 1940. Reported in parentheses are robust standard errors clustered by county.

Table 3. Instrumental Variables Estimate of the Change in Land Value from 1930 to 1940

	First-Stage:		2SLS:	OLS:
	% of County in High Erosion (1)	% of County in Medium Erosion (2)	Change in Land Value, 1930 – 1940 (3)	Change in Land Value, 1930 – 1940 (4)
Erosion Level:				
High - Low			-0.631**	-0.277**
			(0.219)	(0.041)
Medium - Low			-0.125	-0.165**
			(0.154)	(0.030)
Drought Instruments:				
Months in Extreme Drought, 1930s	0.0061**	0.0057		
	(0.0021)	(0.0032)		
Months in Severe Drought, 1930s	0.0031	0.0069**		
	(0.0016)	(0.0022)		
Months in Moderate Drought, 1930s	-0.0014	0.0076**		
	(0.0016)	(0.0025)		
Average Palmer Drought Severity Index, 1930s	0.0487	0.1700**		
	(0.0334)	(0.0603)		
Controls:				
Drought Vars, 1895-1929	YES	YES	YES	YES
State Fixed Effects	YES	YES	YES	YES
County Characteristics	YES	YES	YES	YES
F-stat: Instruments	4.69	4.54		
P-value: Instruments = 0	0.0010	0.0013		
Sample Counties	766	766	766	766

Notes: Column 1 reports first-stage estimates from regressing the percent of a county in high erosion areas on the drought instruments, controlling for drought from 1895 to 1929, state fixed effects, and the same pre-1930's county characteristics as in table 2. Column 2 reports analogous estimates for the percent of a county in medium erosion areas. Column 3 reports two-stage least squares estimates of the relative change in land values by erosion level, controlling for the same variables as in columns 1 and 2 and instrumenting for erosion levels using the excluded drought variables. Column 4 reports OLS estimates of the relative change in land value (as in table 2), controlling for the same variables as in column 3 and for the same sample of counties. All regressions are weighted by county farmland in 1930. Robust standard errors are reported in parentheses. ** denotes statistical significance at the 1% level, * at the 5% level.

Table 4. Changes After 1930 in Agricultural Production, by Erosion Level

		Crops vs. Animals			Wheat vs. Hay			Hansen-Libecap Adjustments	
Erosion Level:	County Share in Farmland (1)	Log Crop Productivity (2)	Log Animal Productivity (3)	Land Share in Crops (4)	Log Wheat Productivity (5)	Log Hay Productivity (6)	Land Share in Wheat (7)	Log Average Farm Size (8)	Cropland Fallowed (9)
High - Low									
1940	-0.014	-0.489**	-0.246**	0.009	-0.263*		0.027	0.052	0.019
	(0.015)	(0.118)	(0.055)	(0.010)	(0.115)		(0.022)	(0.032)	(0.016)
1945	-0.012	-0.135	-0.158*	0.003	-0.215**			0.103**	-0.032*
	(0.017)	(0.078)	(0.064)	(0.011)	(0.073)			(0.034)	(0.013)
1950 - 1954	-0.032*	-0.332**	-0.184**	-0.007	-0.233**	-0.003	0.023	0.065	-0.018
	(0.014)	(0.067)	(0.068)	(0.013)	(0.058)	(0.051)	(0.020)	(0.034)	(0.015)
1959 - 1964	-0.025	-0.303**	-0.349**	-0.030*	-0.223**	0.090	0.003	0.084*	0.005
	(0.016)	(0.076)	(0.092)	(0.015)	(0.064)	(0.048)	(0.020)	(0.040)	(0.016)
1969 - 1974	0.003				0.012	0.149**	-0.039	0.081*	
	(0.016)				(0.060)	(0.045)	(0.022)	(0.041)	
1978 - 1992	-0.022				-0.035	0.199**	-0.067**	0.054	
	(0.018)				(0.043)	(0.046)	(0.025)	(0.046)	
1997					-0.067	0.141*	-0.110**		
					(0.048)	(0.063)	(0.032)		
Medium - Low									
1940	-0.002	-0.464**	-0.099**	-0.002	-0.165		-0.016	0.025	-0.005
	(0.010)	(0.086)	(0.038)	(0.006)	(0.091)		(0.016)	(0.021)	(0.010)
1945	0.021	-0.305**	-0.103*	-0.004	-0.102			0.063*	-0.050**
	(0.012)	(0.064)	(0.042)	(0.007)	(0.054)			(0.025)	(0.010)
1950 - 1954	-0.006	-0.310**	-0.095*	-0.012	-0.214**	0.001	0.005	0.034	-0.044**
	(0.010)	(0.054)	(0.046)	(0.008)	(0.054)	(0.047)	(0.020)	(0.028)	(0.011)
1959 - 1964	-0.015	-0.272**	-0.111	-0.018	-0.300**	0.033	-0.022	0.024	-0.034**
	(0.013)	(0.066)	(0.063)	(0.011)	(0.056)	(0.050)	(0.020)	(0.033)	(0.011)
1969 - 1974	-0.001				-0.075	0.002	-0.068**	0.014	
	(0.012)				(0.046)	(0.039)	(0.020)	(0.033)	
1978 - 1992	-0.004				-0.096*	0.009	-0.097**	-0.006	
	(0.013)				(0.033)	(0.042)	(0.021)	(0.037)	
1997					-0.128	0.014	-0.134**		
					(0.047)	(0.053)	(0.026)		
R-squared	0.6133	0.7877	0.8193	0.7254	0.8915	0.8094	0.4607	0.7473	0.8133
Sample Counties	769	769	769	769	388	388	388	769	587
Weighted by 1930 Value of:	Farmland	Cropland	Pasture	Cropland plus Pasture	Wheat	Hay	Wheat plus Hay	Farmland	Farmland

Notes: For each outcome variable, the column reports estimates from equation (2) in the text and described in notes to Table 2. Where indicated, the weights are adjusted to reflect the average change for the relative unit of land. Reported in parentheses are robust standard errors clustered by county. ** denotes statistical significance at the 1% level, * at the 5% level.

Table 5. Changes in Land-use after 1930, Interacted with County Pre-Characteristics

	Relative Adjustment in Areas with More Banks			Relative Adjustment in Areas with More Tenants		
Erosion Level:	Land Share in Crops (1)	Land Share in Wheat (2)	County Share in Farmland (3)	Land Share in Crops (4)	Land Share in Wheat (5)	County Share in Farmland (6)
High - Low						
1940	-0.042**	0.015	0.008	-0.012	-0.020	0.005
	(0.014)	(0.018)	(0.014)	(0.010)	(0.018)	(0.017)
1945	-0.050**		0.025	-0.001		-0.004
	(0.017)		(0.022)	(0.012)		(0.022)
1950 – 1954	-0.071**	-0.016	-0.004	-0.013	0.048*	-0.016
	(0.019)	(0.021)	(0.015)	(0.015)	(0.020)	(0.016)
1959 – 1964	-0.074**	0.000	-0.030	-0.013	0.026	-0.013
	(0.020)	(0.021)	(0.018)	(0.017)	(0.021)	(0.018)
1969 – 1974		0.010	-0.039*		0.008	-0.035
		(0.023)	(0.018)		(0.022)	(0.018)
1978 – 1982			-0.038*			-0.023
			(0.017)			(0.020)
1987 – 1992			-0.044*			-0.025
			(0.019)			(0.022)
Medium - Low						
1940	-0.017**	-0.012	0.021*	-0.006	-0.009	0.007
	(0.006)	(0.012)	(0.009)	(0.006)	(0.014)	(0.011)
1945	-0.024**		0.007	0.008		-0.031*
	(0.007)		(0.012)	(0.008)		(0.015)
1950 – 1954	-0.013	-0.010	0.018	0.017	0.056**	0.008
	(0.008)	(0.018)	(0.010)	(0.010)	(0.018)	(0.011)
1959 – 1964	-0.021*	-0.007	0.026*	0.015	0.057**	0.027*
	(0.011)	(0.017)	(0.013)	(0.013)	(0.019)	(0.014)
1969 – 1974		0.001	0.026*		0.050*	0.024
		(0.019)	(0.012)		(0.020)	(0.013)
1978 – 1982			0.015			0.033*
			(0.011)			(0.013)
1987 – 1992			0.011			0.014
			(0.013)			(0.014)
R-squared	0.7488	0.4702	0.6180	0.7307	0.4844	0.6087
Sample Counties	696	696	696	696	696	696
Weighted by 1930 Value of:	Farmland	Wheat plus Hay	Farmland	Farmland	Wheat plus Hay	Farmland

Notes: For the indicated outcome variable, each column presents estimates from equation (3) in the text. For columns 1 – 3, the reported coefficients are the interaction term: the adjustment to erosion in areas with more banks in 1928, relative to the adjustment in areas with fewer banks. The log number of banks is normalized to have a mean of zero and a standard deviation of one. Columns 4 – 6 report the analogous coefficients, but for the normalized tenant share of farmland in 1930 (instead of the normalized log number of banks). Reported in parentheses are robust standard errors clustered by county. ** denotes statistical significance at the 1% level, * at the 5% level.

Table 6. Government Program Payments per farm acre, by Erosion Level

		Relative to Low Erosion:		
	All Counties (1)	Medium Erosion (2)	High Erosion (3)	(3) – (2) (4)
Panel A. New Deal payments (1933-1939)				
AAA payments	0.489	0.001	-0.027	-0.028
	[0.327]	(0.017)	(0.023)	(0.021)
Public works spending	0.264	0.008	-0.033	-0.041
	[0.605]	(0.058)	(0.064)	(0.066)
Relief spending	0.508	0.110	0.142	0.032
	[2.435]	(0.100)	(0.129)	(0.118)
New deal loans	0.484	-0.090	-0.087	0.003
	[1.126]	(0.094)	(0.112)	(0.087)
Mortgage loans guaranteed	0.112	0.001	-0.103	-0.104*
	[0.792]	(0.040)	(0.059)	(0.042)
Panel B. Government payments				
All payments, 1969	3.323	-0.245	-0.159	0.086
	[2.646]	(0.256)	(0.284)	(0.254)
All payments, 1974	0.364	0.088	0.257**	0.169**
	[0.390]	(0.056)	(0.066)	(0.045)
All payments, 1987	16.040	-1.838*	-5.511**	-3.673**
	[13.631]	(0.866)	(1.076)	(1.075)
All payments, 1992	7.930	-0.887	-1.323*	-0.436
	[5.580]	(0.477)	(0.549)	(0.472)
All payments, 1997	8.060	-0.222	-0.523	-0.300
	[5.843]	(0.366)	(0.467)	(0.442)
CRP payments, 1992	1.571	0.750**	1.776**	1.026**
	[1.490]	(0.162)	(0.263)	(0.245)
CRP payments, 1997	2.094	1.186**	2.562**	1.376**
	[2.061]	(0.225)	(0.350)	(0.322)
Fraction of money from CRP, 1992	0.217	0.089**	0.185**	0.096**
	[0.149]	(0.018)	(0.027)	(0.023)
Fraction of money from CRP, 1997	0.289	0.105**	0.230**	0.125**
	[0.185]	(0.022)	(0.031)	(0.027)

Notes: Panel A reports differences in 1930's New Deal spending across counties, for a constant sample of 766 counties. Panel B reports later differences in government payments, conservation reserve program payments, and the fraction of payments through the conservation reserve program. Column 1 presents the mean for each variable and the standard deviation in brackets. Column 2 reports the average difference for a county with medium erosion relative to low erosion, controlling for state fixed effects, each characteristic in Panels B – E of Table 1, and the lagged values of those characteristics. Columns 3 and 4 report the same average differences comparing high erosion vs. low erosion and high erosion vs. medium erosion, respectively. All variables and regressions are weighted by county farmland in 1930. Robust standard errors are reported in parentheses. ** denotes statistical significance at the 1% level, * at the 5% level.

Table 7. Changes After 1930 in Population Outcomes, by Erosion Level

Erosion Level:	Log Population (1)	Fraction Rural (2)	Fraction on Farm (3)
High - Low			
1940	-0.081**	-0.0118	0.0159
	(0.021)	(0.0090)	(0.0095)
1950	-0.108**	-0.0115	0.0204*
	(0.041)	(0.0157)	(0.0095)
1960	-0.148*	-0.0165	
	(0.064)	(0.0199)	
1970	-0.157	0.0084	0.0334**
	(0.080)	(0.0225)	(0.0121)
1990	-0.127		0.0120
	(0.109)		(0.0081)
Medium - Low			
1940	-0.065**	-0.0102	0.0000
	(0.020)	(0.0075)	(0.0049)
1950	-0.108**	-0.0065	0.0077
	(0.035)	(0.0126)	(0.0070)
1960	-0.170**	-0.0144	
	(0.055)	(0.0157)	
1970	-0.180*	-0.0094	-0.0027
	(0.072)	(0.0181)	(0.0096)
1990	-0.160		0.0004
	(0.103)		(0.0065)
R-squared	0.5954	0.4558	0.9451
Sample Counties	769	769	769

Notes: For each indicated outcome variable, the column reports estimates from equation (2) in the text and described in the notes to Table 2. All regressions are weighted by county population in 1930. Reported in parentheses are robust standard errors clustered by county. ** denotes statistical significance at the 1% level, * at the 5% level.

Table 8. Changes After 1930 in Unemployment and Manufacturing

		Manufacturing			
	Unemployment Rate (1)	Workers per-labor force (2)	Workers per-capita (3)	Log Establishments (4)	Log Value Added (5)
1930 Mean:	0.0410	0.0628	0.0262		
	[0.0255]	[0.0629]	[0.0276]		
Erosion Level:					
High - Low					
1940	0.0071*	0.0072	0.0040	0.022	0.108
	(0.0034)	(0.0041)	(0.0021)	(0.069)	(0.145)
1945				-0.078	0.003
				(0.066)	(0.183)
1950	-0.0031				
	(0.0020)				
1954				-0.137	-0.138
				(0.078)	(0.247)
1958 (-1964)			0.0036	-0.126	-0.058
			(0.0061)	(0.080)	(0.261)
1967 (-1974)			0.0024	-0.014	-0.226
			(0.0093)	(0.099)	(0.259)
1978 (-1982)				-0.036	-0.443
				(0.110)	(0.494)
1987 (-1992)			0.0050	-0.044	0.157
			(0.0109)	(0.129)	(0.362)
Medium - Low					
1940	-0.0016	0.0000	0.0005	-0.080	-0.127
	(0.0025)	(0.0035)	(0.0020)	(0.049)	(0.111)
1945				-0.011	0.003
				(0.047)	(0.128)
1950	-0.0007				
	(0.0017)				
1954				-0.037	-0.041
				(0.058)	(0.182)
1958 (-1964)			0.0063	-0.100	-0.111
			(0.0047)	(0.062)	(0.178)
1967 (-1974)			0.0131	-0.060	-0.153
			(0.0076)	(0.079)	(0.187)
1978 (-1982)				-0.061	0.022
				(0.093)	(0.338)
1987 (-1992)			0.0250**	-0.045	-0.087
			(0.0075)	(0.109)	(0.263)
R-squared	0.8550	0.5803	0.9203	0.5754	0.9456
Sample Counties	769	550	336	551	275

Notes: For each outcome variable, the 1930 mean is reported at the top of the column and the standard deviation in brackets. The column reports estimates from equation (2) in the text and described in the notes to Table 2. In column 3, manufacturing worker data is only available for 1958, 1967, and 1987, and population data is taken from the nearest decennial. Reported in parentheses are robust standard errors clustered by county. ** denotes statistical significance at the 1% level, * at the 5% level.

Table 9. Log Changes After 1930 in Wage Proxies

		Payroll divided by Workers		
Erosion Level:	Retail Sales per-capita (1)	Retail Sector (2)	Wholesale Sector (3)	Manufacturing Sector (4)
High - Low				
1940	-0.098**	-0.013	-0.035	-0.050
	(0.029)	(0.018)	(0.048)	(0.067)
1945				-0.049
				(0.062)
1954		-0.060**	-0.044	-0.095
		(0.022)	(0.051)	(0.069)
1958 (-1964)	-0.058	-0.030*	-0.050	-0.130*
	(0.031)	(0.015)	(0.032)	(0.060)
1967 (-1974)	-0.034	-0.041*	-0.048	-0.209*
	(0.039)	(0.016)	(0.035)	(0.104)
1978 (-1982)		-0.085*	-0.202**	-0.043
		(0.038)	(0.067)	(0.110)
1987 (-1992)	-0.059	-0.035	-0.104**	-0.077
	(0.053)	(0.018)	(0.039)	(0.065)
Medium - Low				
1940	-0.063**	-0.012	-0.048	-0.131**
	(0.022)	(0.014)	(0.034)	(0.049)
1945				-0.116**
				(0.040)
1954		-0.041**	-0.096**	-0.092*
		(0.013)	(0.032)	(0.045)
1958 (-1964)	0.005	-0.017	-0.079**	-0.106*
	(0.021)	(0.011)	(0.026)	(0.044)
1967 (-1974)	0.033	-0.016	-0.076**	-0.073
	(0.028)	(0.011)	(0.024)	(0.046)
1978 (-1982)		-0.041	-0.178**	-0.080
		(0.024)	(0.040)	(0.072)
1987 (-1992)	0.011	-0.014	-0.126**	-0.065
	(0.042)	(0.013)	(0.027)	(0.045)
R-squared	0.9749	0.9970	0.9921	0.9906
Sample Counties	758	748	495	257

Notes: For each outcome variable, the column reports estimates from equation (2) in the text and described in the notes to Table 2. In column 1, retail sales data is only available for 1958, 1967, and 1987, and population data is taken from the nearest decennial. Reported in parentheses are robust standard errors clustered by county. ** denotes statistical significance at the 1% level, * at the 5% level.

Good Fences Make Good Neighbors: Evidence on the Effects of Property Rights

Richard A. Hornbeck*

Abstract

The second essay examines the impact on agricultural development from the introduction of barbed wire fencing to the American Plains in the late 19th century. Farmers were required to construct fences to be entitled to compensation for damage by others' livestock. From 1880 to 1900, the introduction and universal adoption of barbed wire greatly reduced the cost of fences, relative to predominant wooden fences, most in counties with the least woodland. Over that period, counties with the least woodland experienced substantial relative increases in settlement, land improvement, land values, and the productivity and production share of crops most in need of protection. This increase in agricultural development appears partly to reflect farmers' increased ability to protect their land from encroachment. States' inability to protect this full bundle of property rights on the frontier, beyond providing formal land titles, might have otherwise restricted agricultural development.

---

*I thank Michael Greenstone, as well as Daron Acemoglu, Lee Alston, Josh Angrist, David Autor, Abhijit Banerjee, Dora Costa, Esther Duflo, Claudia Goldin, Tal Gross, Raymond Guiteras, Jeanne Lafortune, Philip Oreopoulos, Paul Rhode, Peter Temin, and numerous seminar participants for their comments. I also thank Daniel Sheehan for assistance with GIS software. hornbeck@mit.edu

The assignment and enforcement of property rights are thought to encourage economic efficiency [Coase 1960] and to be an important determinant of product choices, investment levels, production methods [Cheung 1970], and macroeconomic performance [North 1981; Acemoglu and Johnson 2005]. Coase presents the example of a farmer and a cattle-raiser, in which the costless assignment and tradeability of property rights leads to the efficient allocation of resources. Much attention has focused on the role of land rights in agricultural development, where more secure rights might improve economic efficiency by increasing farmers' investments.[1] The typical positive correlation between investment and property rights, however, may reflect the endogenous determination of property rights [Demsetz 1967; Miceli et al. 2001].[2]

Valid causal estimation is dependent upon plausibly exogenous variation in property rights, mainly either from instrumental variables [Besley 1995; Brasselle et al. 2002] or land rights policies [Jacoby et al. 2002; Banerjee et al. 2002]. These exogeneity assumptions are typically most plausible for micro-level variation in property rights, but this variation may overstate the effects of larger reforms or macro-level differences.[3]

This paper attempts to estimate the effects of property rights on farmers' production decisions, using a widespread and plausibly exogenous decrease in the cost of establishing and enforcing property rights: the introduction of barbed wire fencing to the American Plains in the late 19th century. At that time, farmers were required to build fences to protect their land and crops from others' livestock. Without fences, farmers were apt to suffer substantial damages and typically had no formal right to compensation from the livestock's owner. Fences may also have decreased farmers' susceptibility to land expropriation by providing an

---

[1] See Besley [1995] for a discussion of three mechanisms: decreased expropriation raises the expected return on investment; an improved ability to collateralize land increases access to credit; and lower costs of trading land raise the expected return on investment.

[2] Unobserved factors, such as land quality, may influence both investment and property rights. Investment may increase the legitimacy of land claims or otherwise make the land more worthwhile to secure. Goldstein and Udry [2005] discuss an opposite example where fallowing land may jeopardize ownership claims.

[3] When property rights vary across land plots within small areas, observed effects may reflect sorting of producers or substitution in production that would not occur if property rights were made more widely available.

informal mechanism for delineating and substantiating land claims. Both effects of fencing are tantamount to an increase in farmers' property rights in the standard security model of property rights [Besley 1995].[4]

Before barbed wire, fence construction on the Plains had been severely restricted by high costs in areas that lacked locally available fencing materials. Due to high transportation costs from Eastern wooded areas, small sections of local woodland provided a vital source of timber for fencing on the Plains. The introduction and universal adoption of barbed wire from 1880 to 1900 had its greatest effect in areas with the least woodland that had been most costly to fence.[5]

Using decennial data from the Census of Agriculture, I estimate that counties with the least woodland experienced large increases in agricultural development from 1880 to 1900, relative to counties with sufficient woodland for farmers to have accommodated previous fencing material shortages. Estimated increases were particularly substantial along intensive margins of agricultural production. Controlling for constant differences among counties and state-wide shocks to all counties, the fraction of county farmland that was improved increased by 19 percentage points in counties with the least woodland.

From 1880 to 1890, average crop productivity increased by 23% in counties with the least woodland, controlling for crop-specific differences among counties and crop-specific state-wide shocks to all counties. The increased productivity was entirely in crops most susceptible to damage from roaming livestock, as opposed to hay. In addition, farmers shifted cropland allocations toward these more at-risk crops. Crop production changed little from 1890 to 1900, indicating that farmers adjusted more rapidly along these margins.

Agricultural development increased along these intensive margins, even as counties with the least woodland expanded along the extensive margin of total farmland. The estimates

---

[4] In contrast to other sources of variation in property rights, fencing did not affect farmers' formal land ownership and, thus, had little effect on farmers' ability to trade the land or use it as collateral.

[5] Barbed wire fencing need not affect property rights in other contexts with different legal structures or geographical characteristics; rather, these factors combined to make barbed wire particularly important for property rights protection in the United States' Plains in the late 19th century. See, for example, Anderson and Hill [1975].

are robust to controlling for changes over time that are correlated with counties' distance West and distance from St. Louis; counties' region, subregion, or soil group; or counties' initial intensity of land improvement. By contrast, estimated increases along the extensive margin of total farmland are more sensitive to these robustness checks.

Improvements in production were reflected in increased land values in the six Plains states analyzed, totaling roughly 0.9% of national GDP. In all, these estimates lend support to historical accounts that "without barbed wire the Plains homestead could never have been protected from the grazing herds and therefore could not have been possible as an agricultural unit" [Webb 1931, p.317].

The historical setting provides a rare opportunity to compare the effects of the *de facto* improvement in property rights enforcement provided by barbed wire with the effects of contemporaneous *de jure* efforts by some state legislatures. In the early 1870's, some states and counties adopted "herd laws" that shifted formal fencing responsibilities to livestock owners. In contrast to the large estimated effects of barbed wire, these reforms appear to have had little effect. In this historical context when settlement was sparse, it appears that the herd laws were weakly enforced and farmers continued to require fencing for protection.

Attributing the estimated effects of barbed wire to the increased security of property rights requires that barbed wire affected agricultural development only by increasing property rights enforcement. I examine the potential for two other effects of cheaper fencing: the direct encouragement of cattle production and/or the joint production of crops and cattle. There is little evidence of a positive effect of barbed wire on cattle production, however, and joint production is found to have decreased. This fails to contradict the hypothesis that barbed wire's effects operated mainly through an increase in property rights enforcement.

The paper is organized as follows. Section I reviews historical accounts of the need for alternative fencing materials in timber-scarce areas and the introduction of barbed wire fencing. Section II nests these historical accounts in a more stylized theoretical framework. Section III describes the data and presents summary statistics. Section IV develops the empirical methodology. Section V presents the main estimation results and explores their robustness. Section VI examines other potential confounding factors. Section VII analyzes states' reforms to fencing requirements. Section VIII concludes.

I History of Barbed Wire and the Great Plains

I.A Timber Shortages Constrained Property Rights Enforcement

Breaking with English common law, legal codes in the United States in the 19th century had long required the construction of a "lawful fence"[6] to secure farmland [Kawashima 1994]. Without such a fence, a farmer had no formal entitlement to compensation for damages inflicted by roaming livestock. In practice, fences were necessary to protect crops and they required substantial investment. In 1872, fencing capital stocks in the United States were roughly equal to the value of all livestock, the national debt [U.S. House 1872], or the railroads. Annual fencing repair costs were greater than the combined annual tax receipts at all levels of government [Galveston News 1873, as cited in Webb, pp.288–289].

Fencing became increasingly costly as settlement moved into areas with little woodland. As early as 1841, agricultural journals discussed the lack of timber for fencing as the major barrier to the settlement of Western regions [Bogue 1963a, p.74]. An 1871 report, prepared as a guide for immigrants, focused on three main characteristics of farmland in Plains counties: its price, the amount of timber, and the amount fenced [U.S. House 1871].

> When the frontier line left the timbered region and came onto the Prairie Plains, the pioneers found there neither timber nor stone. There was nothing with which to fence the land. ... Without fences [they] could have no crops; yet the expense of fencing was prohibitive, especially in the Plains proper. [Webb, p.281, p.287]
>
> When he sought to fence his crops against marauding livestock, the prairie farmer faced the timber problem at its most acute. [Bogue 1963b, p.7]

---

[6] The exact requirements for a lawful fence varied by state, but a fence was generally required to be strong enough to prevent the intrusion of a roaming animal. Some laws also required that the animal's owner be shown to have contributed to the damages through particular negligence or maliciousness.

During the early period of settlement there was little conflict between those growing crops and those raising cattle, but that was because farmers were mostly forced out of necessity to remain elsewhere. Much potential conflict existed [Alston et al. 1998], which is the relevant factor in farmers' location and production decisions.[7]

High transportation costs made it impractical to supply areas without woodland with the required amounts of timber for fencing [Kraenzel 1955, p.129; Hayter 1939; Bogue 1963b, p.7], though systematic data on transportation costs are unavailable. In low-woodland areas, "the scarcity of wood and its consequent high cost encouraged experimentation" [Primack 1969, p.287]. There was some use of hedge fences, yet these were found to be too difficult and costly to grow and control. Smooth iron wire fences were available, yet they were easily broken by animals, prone to rust, and broke (sagged) in cold (hot) weather.

Writing about central Iowa, Bogue [1963b, p.6] further describes which counties were most affected by timber shortages:

> Where timber and prairie alternated, locations in or near wooded areas were relatively much more attractive.... [T]here developed a landholding pattern of which the timber lot was an intricate part. Settlers on the prairie purchased five or ten acres along the stream bottoms or in the prairie groves and drove five, ten or fifteen miles to cut building timber or to split rails during the winter months.

Smaller counties were roughly 30 miles on each side, so farmers traveling 5–15 miles for timber would still have been mostly within their home county. For simplicity and transparency, I assume that farmers' access to wooden fence materials depended on the amount of woodland in their home county.

The empirical analysis will examine continuous variation in county woodland levels, but Bogue's description provides some useful benchmark values for the level of woodland in a county. For a standard homestead farm size of 160 acres, a county would need to be

---

[7] Differences in the ability to prevent encroachment are key, rather than differences in the observed initial equilibrium levels of encroachment. The empirical analysis does not define counties with a "property rights problem" as those with both livestock and cropland, as this is endogenously determined by many unobserved local factors. Instead, the paper focuses on local woodland as an exogenous determinant of fencing availability and the cost of securing farmland.

roughly 4% woodland for each farm to acquire 5–10 acres of woodland. Based on this calculation, counties are assigned to the following three woodland categories: low (0–4%), medium (4–8%), and high (8–12%). The designated "low" counties roughly represent those most constrained by timber scarcities, while "medium" counties could have partially adjusted with this landholding pattern and then have been less affected along with "high" counties. The exact cutoffs for these categories are not relevant for the results; rather the continuous estimates will be evaluated at three corresponding benchmark levels (0%, 6%, 12%) in order to assist in interpreting the estimated magnitudes.

I.B The Arrival of Cheap Barbed Wire Fences, 1880 to 1900

Crude versions of enhanced wires were patented as early as 1867, but the most practical and ultimately successful design for barbed wire was patented in 1874 by Glidden, a farmer in DeKalb, Illinois. Glidden's design, which featured two twisted strands of wire that held metal barbs in place, had three important characteristics: barbs prevented cattle from breaking the fence; twisted wires tolerated temperature changes; and the design was easy to manufacture. Glidden sold a half-stake in the patent for only a few hundred dollars to Ellwood, a hardware merchant in DeKalb, and the two started the first commercial production of barbed wire, producing a few thousand pounds per year by hand [Hayter].

Barbed wire was cheaper than wooden fencing, particularly in timber-scarce areas, and it had lower labor requirements.[8] Correspondence from Glidden and Ellwood in 1875 reflected the close relationship between local woodland and the demand for barbed wire: "where lumber is reportedly dearer, the wire would probably sell for more" [Webb p.310]. Given current prices, they did "not expect the wire to be much in demand where farmers can build brush and pole fences out of the growth on their own land" [Hayter].

---

[8] In Iowa, wooden fences varied in total construction costs per rod from $0.91–$1.31 in 1871, while barbed wire fences cost $0.60 in 1874 and below $0.30 in 1885 [Bogue 1963b, p.8]. Other reports quote barbed wire fences as costing $0.75 per rod in Indiana in 1880, while hedge fences cost $0.90 per rod and were wasteful of the land [Primack 1977, p.73]. Primack [1977, p.82] estimates that a rod of barbed wire took 0.08, 0.06, and 0.04 days to construct in 1880, 1900, and 1910. The labor requirements for constructing wooden fences were constant throughout this period: 0.20, 0.34, and 0.40 days for board, post and rail, and Virginia rail.

In 1876, the country's largest plain wire manufacturer (Washburn & Moen) bought half of Glidden's business for $60,000 cash, plus royalties, and began the first large-scale production of barbed wire. By 1879, they had produced nine million pounds.[9] In contrast to Glidden's sale to Ellwood, Washburn & Moen's purchase showed an awareness of barbed wire's potential. Even the initial level of demand was sufficient to ensure "enormous profits" [Webb, p.309].[10]

Newspaper advertisements began to appear in Kansas and Nebraska in 1878 and 1879 [Davis 1973, pp.133–134] and there was a series of public demonstrations in the early 1880s. Once the effectiveness of barbed wire was proved, "Glidden himself could hardly realize the magnitude of his business. One day he received an order for a hundred tons; 'he was dumbfounded and telegraphed to the purchaser asking if his order should not read one hundred pounds' " [Webb, p.312].

Table 1 provides statistics on the timing of barbed wire's introduction and its universal adoption. Panel A shows a sharp increase around 1880 in the annual production of barbed wire. Panel B shows the resulting transformation in regional fence stocks. Before 1880, fences were predominately made of wood. From 1870 to 1880, there were some small increases in wire fencing, including both smooth wire and barbed wire. After 1880, there were rapid increases in barbed wire fencing. Total fencing increased most in the Plains and Southwest regions where there were more timber-scarce areas. Wood fencing also initially increased, however, highlighting that it would be inappropriate to attribute all regional increases in fencing and economic activity to the introduction of barbed wire.

Even as the quality of barbed wire improved and consumers became increasingly aware of its effectiveness in the early 1880s, falling input costs and manufacturing improvements

---

[9] This process began in 1875 when Washburn & Moen, headquartered in Massachusetts, sent an agent to investigate unusually large orders from DeKalb, Illinois. They acquired barbed wire samples and contracted an expert machinist to design automatic machines for its production.

[10] McFadden [1978] provides details on the development of these businesses, with the 1899 incorporation of the American Steel and Wire Company of New Jersey leading to the monopolization of the barbed wire industry.

drove down prices: \$20 (1874), \$10 (1880), \$4.20 (1885), \$3.45 (1890), and \$1.80 (1897).[11] Panel C of Table 1 shows that all new fences constructed after 1900 were made of barbed wire, so further price declines or quality improvements would not have had a differential effect across counties with varying access to wooden fences. Thus, the differential effects of barbed wire on farmers' fencing costs were predominately from 1880 to 1900.[12]

The empirical approach presented here requires that the introduction of barbed wire fencing was exogenous, i.e., that its rapid rise around 1880 was not caused by the anticipated development of low-woodland areas. This assumption appears plausible for two main reasons. First, from a microeconomic perspective, the demand for fencing alternatives had already been high for decades and Glidden-Ellwood appear not to have anticipated the tremendous market demand for barbed wire. Second, from a more macroeconomic perspective, the necessary cheap steel was only becoming available around 1880.

Barbed wire's widespread commercial success was made possible by unrelated developments in the industrial steel sector. The production of strong rust-free barbed wire required steel, which became dramatically cheaper as the Bessemer-steel process became widely used (patented in England in 1855). Figure 1 shows prices for iron, steel, and barbed wire. Barbed wire's introduction and mass-production follows the sharp decline in steel prices in the 1870s.[13] Primack [1969] summarizes:

> Outcries about the burdens of fencing by agriculturists in the 1850 to 1880 period seem amply justified. A need was revealed and the problem was resolved, not by changing laws and institutions but rather by technological change. This solution had to wait for the development of cheap steel in the industrial sector. Then a solution was found in wire fencing, cheap in both money and labor costs. [p.289]

---

[11] Prices are per hundred pounds [Webb p.310]. Hayter reports similar prices for 1874 and 1893.

[12] Even by 1890, new fence construction in much of the country was entirely of barbed wire, yet there was still some wooden fence construction in the Prairie and the Southwest. This likely reflects less-developed distribution networks, as well as ranchers' opposition to enclosure and fence-cutting wars that were resolved in the late 1880s [McCallum and McCallum 1965, pp.159–166; Webb, pp.312–316].

[13] Previously discussed historical accounts provided additional price observations for barbed wire, indicated with X's, though pre-mass production prices are off the chart. Note that the 1899 increase in barbed wire prices coincides with the monopolization of the barbed wire industry, though the Spanish-American War may have contributed to the observed increases in prices for barbed wire, steel, and iron.

II Theoretical Framework

To provide these historical accounts a more formal theoretical structure for clarity and estimation purposes, consider a farmer in each county c and time period t choosing a level of investment I_{ct} and property rights enforcement R_{ct} to maximize profits. The farmer produces output $F(I_{ct}, q_{ct})$, where q_{ct} denotes the quality of the land or a composite of other characteristics. $F(\cdot,\cdot)$ is increasing in both arguments and $F_{12}(\cdot,\cdot) > 0$. Following Besley [1995], an expected percentage of output is lost in each period, $\tau(R_{ct}) \in [0,1]$, and it is decreasing in the level of property rights enforcement, $\tau'(R_{ct}) < 0$. Investment and property rights are each produced at some cost, $C_{ct}(I_{ct})$ and $C_{ct}(R_{ct})$. Thus, farmers choose I_{ct} and R_{ct} to maximize:

$$(1 - \tau(R_{ct}))F(I_{ct}, q_{ct}) - C_{ct}(I_{ct}) - C_{ct}(R_{ct}). \tag{1}$$

An optimal interior solution satisfies two first-order constraints:

$$C'_{ct}(I_{ct}) = F_1(I_{ct}, q_{ct})(1 - \tau(R_{ct})) \tag{2}$$

$$C'_{ct}(R_{ct}) = -\tau'(R_{ct})F(I_{ct}, q_{ct}). \tag{3}$$

In equation (2), the marginal cost of investment is set equal to the amount of marginal return that the farmer expects to retain. In equation (3), the marginal cost of property rights is set equal to the marginal increase in retained total output. These equations generate three relationships of interest. First, the optimal choice of investment is increasing in the level of property rights because a greater proportion of the returns would be kept. Second, the optimal choice of property rights is increasing in the level of investment because there is an added incentive to keep the output. Third, higher land quality would directly increase both property rights and investment by raising total production and the marginal return to investment. The empirical identification problem is that an observed correlation between I_{ct} and R_{ct} could reflect more than one of these three effects.

It may be possible, however, to identify the direct effect of property rights on investment by examining the effect of changes in the marginal cost of producing property rights. Equation (2) defines the optimal choice of investment, $I^*_{ct}(R^*_{ct}, q_{ct})$, and inserting that function into equation (3) defines $R^*_{ct}(q_{ct}, C'_{ct}(R^*_{ct}))$. Given some equilibrium marginal cost P_{ct} of establishing property rights, it follows that $\frac{dI^*}{dP} = \frac{\partial I^*}{\partial R^*} \cdot \frac{\partial R^*}{\partial P}$. Thus, an estimate of $\frac{dI^*}{dP}$ would identify $\frac{\partial I^*}{\partial R^*}$ up to some negative term, $\frac{\partial R^*}{\partial P}$. Empirical analyses rarely have a natural or comparable unit of measurement for property rights, in which case a "reduced form" estimate of $\frac{dI^*}{dP}$ is as informative about the underlying causal relationship $\frac{\partial I^*}{\partial R^*}$ as if divided by a "first-stage" estimate of $\frac{\partial R^*}{\partial P}$.

To relate the marginal cost of producing property rights to the introduction of barbed wire, assume further that property rights are produced through the use of some combination of timber and barbed wire, $R_{ct} = R(T_{ct}, B_{ct})$. The price of barbed wire (p^B_t) is assumed to be decreasing over time, but constant across counties.[14] The price of timber in each county (p^T_c) is assumed to be constant over time, but decreasing in the percentage of the county that is woodland, i.e., $p^T_c = g(W_c)$ and $g'(W_c) < 0$.[15] The cost of producing property rights is the minimum obtainable value from choosing B_{ct} and T_{ct} to minimize

$$(4) \qquad p^B_t \cdot B_{ct} + p^T_c \cdot T_{ct}, \text{ subject to: } R_{ct}(T_{ct}, B_{ct}) = \overline{R}.$$

If timber fences are initially constructed and they are not a perfect complement to barbed wire, a decrease in the price of barbed wire that results in its construction would decrease the marginal cost of producing property rights more in counties with less woodland and higher timber prices, i.e., $\frac{\partial^3 C}{\partial p^T \partial p^B \partial R} > 0$. Once timber fences were no longer constructed, further price declines would have no differential effect across counties with different woodland

---

[14] While demand for barbed wire was higher in timber-scarce regions, the market for barbed wire became sufficiently competitive and spread throughout the country that differences in local prices would have been minimal.

[15] Local timber prices may have changed over time, at least in part due to barbed wire, but the subsequent universal adoption of barbed wire implies that changes in timber prices had little differential effect on fencing costs.

levels.[16] Thus, barbed wire would encourage higher levels of property rights protection in timber-scarce areas in the period from its widespread introduction to its universal adoption (1880–1900).

III Data and Summary Statistics

III.A Data Construction

The dataset is drawn from the United States Census of Agriculture and includes county-level data by decade from 1870 to 1920 [Gutmann 2005]. The analysis is restricted to Plains states for which data are available in every census from 1870 to 1920: counties in Minnesota, Iowa, Nebraska, Kansas, Texas, and Colorado. Data in these counties are first available in 1870 and data on land improvement are available through 1920. Some county boundaries changed over this period, and the data are adjusted to hold the 1870 geographical units constant (see Data Appendix).[17]

Figure 2 shows the sample counties, shaded to represent initial woodland levels that are held fixed in the empirical analysis.[18] Because the empirical analysis will include state-by-decade fixed effects, the relevant variation is mostly in Iowa, southern Minnesota, and the eastern parts of Kansas and Nebraska. Non-sample counties in these states are mainly excluded because of unavailable data in 1870 and 1880.

---

[16] This case represents a corner solution to equation (4).

[17] Data starting in 1870 are not available for other Plains states, such as Oklahoma, South Dakota, and North Dakota. Partial data from states further East are available from other sources, but local variation in woodland had less influence on the availability of fencing materials and the remaining woodland would have been more likely to reflect differences in previous land development.

[18] Data are available for the number of acres of woodland in farms, so county woodland levels are defined as initial acres of woodland in farms per county acre. This is not ideal because low woodland may reflect low initial settlement. For this reason, initial woodland levels are calculated using 1880 data when settlement was higher. Because wooded areas in these counties were initially most desirable, it seems plausible that most county woodland would be reflected in this measure. The empirical results are similar when using 1870 woodland levels. Later woodland levels are potentially endogenous to changes in agricultural development, but these results are also similar.

The empirical analysis focuses initially on two measures of farmers' land-use: the percent of county land that is in farms, including land for crops and livestock;[19] and the percent of land in farms that is improved.[20] The percent of county land in farms represents the extensive margin, which mainly reflects farmers' expected returns to converting land from the public domain. The percent of farmland improved represents the intensive margin, which may reflect farmers' willingness to sink investments into the future productivity of that land. Improved land could be plowed for crops or otherwise prepared for livestock, but the definition appears to exclude land that is simply fenced.

In focusing on the improvement intensity of farmland, this paper is similar to most farm-level studies of property rights that analyze whether a given plot of farmland has been improved in some way. An increase in this measure may reflect that farmers substituted toward improvement-intensive activities, where profitability is typically more sensitive to expropriation risk.[21]

The data are available by decade, so there is limited flexibility in analyzing responses to the exact timing of barbed wire's introduction. Because the mass distribution of barbed wire was just beginning by 1880 and fencing stocks had yet to respond substantially, all 1880 county outcomes represent the end of the pre-barbed wire period.[22] Because new fence construction was entirely barbed wire by 1900, at the latest, 1900 marks the moment when barbed wire no longer had a differential effect.

---

[19] The definition of land in farms differs very slightly in some periods, but generally "describes the number of acres of land devoted to considerable nurseries, orchards and market-gardens, which are owned by separate parties, which are cultivated for pecuniary profit, and employ as much as the labor of one able-bodied workman during the year. To be included are wood-lots, sheep-pastures, and cleared land used for grazing, grass or tillage, or lying fallow. Those lands not included in this variable are cabbage and potato patches, family vegetable-gardens, ornamental lawns, irreclaimable marshes, and considerable bodies of water" [Gutmann 2005].

[20] The clearest definition of improved land is from the 1920 Census: "All land regularly tilled or mowed, land in pasture which has been cleared or tilled, land lying fallow, land in nurseries, gardens, vineyards, and orchards, and land occupied by farm buildings" [Gutmann 2005].

[21] If it were true that investment increased property rights but property rights did not encourage investment [Brasselle et al.], then the intensity of land improvement should decrease when barbed wire offered a cheaper alternative for protecting property rights.

[22] Land-use measures were reported for the Census year and productivity is imputed from production and acreage in the previous year.

III.B Summary Statistics

The average amount of woodland in the sample is 10%, but most counties have lower woodland levels: 39% have 0–4% woodland, 15% have 4–8% woodland, 11% have 8–12% woodland, and 35% have more than 12% woodland. At the lower, middle, and upper end of the first three categories, three benchmarks (0%, 6%, 12%) represent meaningful points in the distribution of woodland levels: 20% of the sample has less than 1% woodland; the median woodland level is 6%; and 12% is among the higher typical levels of woodland.

Over the entire sample region, counties with different woodland levels are not evenly balanced geographically. Thus, the main results all focus on specifications that include a full set of state-by-decade fixed effects. To account for geographic differences within states, I present results from specifications that also control for distance West, distance from St. Louis, or finer regional groupings.

Table 2 reports average county characteristics in 1880 for all sample counties and within the first three woodland categories. Prior to barbed wire's introduction, low-woodland counties tended to be larger but lagged behind in settlement, land value, and agricultural production. Less of the farmland in low-woodland counties was improved or used for crops. Low-woodland counties had somewhat less cropland in corn and more in hay, which is less susceptible to damage from livestock. Overall, crop choices are fairly balanced across the three categories, however, which somewhat mitigates concerns that crop-specific technology or price shocks could have large differential effects by woodland percentage.

Total fencing expenditures were lower in low-woodland counties, but were roughly 4% of the total output value across each woodland percentage group. Given higher per unit costs in low-woodland areas, this would suggest a lower intensity of fencing in those areas.[23] Medium-woodland county averages generally fall between those for low-woodland and high-woodland counties, but are much closer to the latter.

---

[23] Data on fencing expenditures was only collected in 1880.

IV Measurement Framework

IV.A Estimation Setup: A Discrete Example

The estimation strategy is most easily described in a discrete example with two county types and two time periods. Assume that farmers in low-woodland counties with $W_c \in (0, 0.04)$ have a high timber price p^T_H, while farmers in medium-woodland counties with $W_c \in (0.04, 0.08)$ have a medium timber price p^T_M. Furthermore, the price of barbed wire is infinite in the first time period and is p^B in the second time period, with $p^B < p^T_H$. A potential estimator for the effect on the production outcome, Y, from a decrease in the cost of property rights protection is then:

$$(5) \qquad (\widehat{Y}_{c=L,t=2} - \widehat{Y}_{c=L,t=1}) - (\widehat{Y}_{c=M,t=2} - \widehat{Y}_{c=M,t=1}).$$

For this difference-in-difference estimator to be unbiased, farmers' production decisions must depend additively on unobserved factors, such that:

$$(6) \qquad Y^*_{ct}(R^*_{ct}, q_{ct}) = F(R^*_{ct}) + \gamma_t + \mu_c + \epsilon_{ct}$$

and $\mathrm{E}[\epsilon_{ct}] = 0$. The assumption is that farmers in counties with different amounts of woodland would have, on average, made the same production changes if not for an increase in property rights protection due to barbed wire's introduction. The analysis in section II was substantially simplified by excluding some variables from each function, such as in separating the costs of property rights and investment, but any potential for estimation bias can be thought of as a violation of this assumption.

It is impossible to test this identification assumption directly, but additional time periods and greater variation in county types can be used to form indirect tests. First, the same estimator for two periods before the introduction of barbed wire tests whether investment decisions in these two types of counties had been trending similarly. Second, the same estimator for two periods after the universal adoption of barbed wire would test for other sources of differences, given that further price declines would not have differential effects. Any differential trends before barbed wire's introduction or after its universal adoption may or may not have occurred between those periods, but the results can be tested for robustness to each scenario.

A third specification test is based on the presumption that local wooden fencing costs depended on local woodland in a non-linear and convex manner. Once settlers were able to acquire the needed acres of woodland, as described by Bogue, one would expect further woodland to have made less difference in the cost of fencing materials.[24] If high-woodland counties with $W_c \in (0.08, 0.12)$ had a low timber price p_L^T that was much closer to p_M^T than was p_H^T, then the estimate from equation (5) should be greater than a similar estimate comparing medium- and high-woodland counties. This test is performed by the difference-in-difference-in-difference estimator:

$$\begin{aligned}(7)\qquad [(\widehat{I}_{c=L,t=2} - \widehat{I}_{c=L,t=1}) &- (\widehat{I}_{c=M,t=2} - \widehat{I}_{c=M,t=1})] \\ &- [(\widehat{I}_{c=M,t=2} - \widehat{I}_{c=M,t=1}) - (\widehat{I}_{c=H,t=2} - \widehat{I}_{c=H,t=1})].\end{aligned}$$

The intuition for this analysis can be seen in a plot of the average improvement intensity of farmland, broken out by decade and county woodland group. In Figure 3, medium- and high-woodland counties are shown to have changed similarly over the entire period. Low-woodland counties also changed similarly, except for large relative increases during the period of adjustment to barbed wire (1880–1900). This analysis of woodland categories is intended only to illustrate the intuition for the methodology, whereas the later empirical analysis examines continuous variation in woodland levels and provides a more thorough treatment of potential confounding factors.

---

[24] In particular, as more wooded areas become scattered throughout a county, the distance to the nearest plot would tend to fall at a decreasing rate. Also, it was increasingly difficult to adjust the type of wooden fence to substitute away from using as much timber [Primack 1977, p.70].

IV.B Main Estimating Equation

For the main empirical analysis, county-level production outcomes are first-differenced to control for any county characteristics that are constant over time.[25] State-by-decade fixed effects α_{st} are included to control for state-specific shocks that have an equal effect on all counties in the state. To allow flexibly for changes over time to be correlated with county woodland levels, included for each decade is a fourth-degree polynomial function of a county's woodland level in 1880. For each production outcome, the estimated equation is:

$$(8) \qquad Y_{ct} - Y_{c(t-1)} = \alpha_{st} + \sum_{t=2}^{T}(\beta_{1t}W_c + \beta_{2t}W_c^2 + \beta_{3t}W_c^3 + \beta_{4t}W_c^4) + \epsilon_{ct}.$$

The estimated β's summarize how changes over each decade in production outcome Y varied by county woodland level W.

V Estimation Results

V.A Land Improvement and Land Settlement

Equation (8) is estimated for the fraction of farmland improved in each county. The full set of estimated β's is prohibitively difficult to interpret numerically, but the results can be seen in Figure 4.[26] The solid line reports the estimated change over the indicated time period for

---

[25] Estimating equation (8) with county fixed effects instead yields identical estimates, but their standard errors tend to be higher by 8–10% because the untransformed error terms are closer to being a random walk than serially uncorrelated. To see this, consider that if equation (8) were estimated for each time period without first-differencing Y, changes over time in the estimated coefficients would still not depend on fixed county characteristics. That is, whether some county with high woodland always had a high outcome level would not influence estimated changes in the dependence of that outcome on woodland. Thus, adding county fixed effects to the regression does not affect differences in the estimated coefficients across periods, though it does increase their precision by reducing unexplained variation. Those exact differences in estimated coefficients are obtained by first-differencing Y and, in this case, it further reduces the unexplained variation in Y.

[26] For conciseness, the displayed results are limited to woodland levels less than 0.12 or 12%, though equation (8) is estimated for the entire distribution of woodland levels.

a county with that woodland level, relative to the estimated change for a county with 0% woodland.[27] The two dashed lines report 95% confidence intervals around the estimates.

From 1880 to 1890 and 1890 to 1900, counties with the least woodland made large relative gains in the improvement intensity of farmland. By contrast, there were no substantial relative changes at low woodland levels before 1880, after 1900, or at higher woodland levels from 1880 to 1900.

In order to display and interpret these results numerically, the estimated changes are evaluated at representative woodland levels: the most affected low-woodland county, with 0% woodland; the average medium-woodland county, with 6% woodland; and the least affected high-woodland county, with 12% woodland. The predicted change for a county with 0% woodland relative to the predicted change for a county with 6% woodland is analagous to a difference-in-difference estimate for counties with those exact woodland levels, but the parameterized regression uses available data from counties with similar woodland levels.

Columns 1 and 2 of Table 3 report the evaluated numerical results from estimating equation (8) for the fraction of farmland improved. In each decade, the coefficient in column 1 corresponds exactly to the difference in the graphed solid line at 0 and 0.06 in Figure 4.

The estimated magnitude is interpreted as follows: the top coefficient in column 1 reports that acres of improved land per acre of farmland increased from 1870 to 1880, on average, 1.5 percentage points more in a county with 0% woodland than in a county with 6% woodland.[28] From the same regression, column 2 reports the predicted change for a county with 6% woodland relative to a county with 12% woodland. In parentheses is the standard error for each coefficient, corrected for heteroskedasticity and clustered at the county level. In brackets is the t-statistic of the absolute difference between the coefficients comparing 0% vs. 6% and 6% vs. 12%. For example, the coefficients in the first row of columns 1 and 2 are not statistically different with a t-statistic of 0.22.

---

[27] Due to the inclusion of state-decade fixed effects, the estimated results are only interpretable relative to some defined benchmark woodland level.

[28] That is, a county with 0% woodland that had 50% of its farmland improved would have, in expectation, caught up to a county with 6% woodland that initially had 51.5% of its farmland improved.

From 1880 to 1900, the improvement intensity of farmland increased by a statistically significant and substantial 19 percentage points in counties with 0% woodland relative to counties with 6% woodland (Table 3, column 1). Recall that from its widespread introduction in 1880 to its universal adoption by 1900, barbed wire most increased the availability of fencing in counties with the least woodland. By contrast, there are few substantial estimated changes before 1880, after 1900, or between higher woodland levels from 1880 to 1900. This result is clear in Figure 5, which plots the estimated cumulative changes after 1870.

It lends credibility to the results that there are few substantial changes apart from the estimated increases at low woodland levels from 1880 to 1900. Depending on assumptions about other agricultural changes, the reported relative changes from 1880 to 1900 could themselves be differenced across and/or within time periods. If the relative trends from 1870 to 1880 would have continued, the adjusted increase from 1880 to 1900 would be 16 percentage points for a county with 0% woodland relative to a county with 6% woodland.[29] If estimated changes at higher woodland levels (6% vs. 12%) reflect other contemporaneous changes that were linearly correlated with woodland levels, the adjusted change from 1880 to 1900 would be 17 percentage points.[30] Under both assumptions, the adjusted change from 1880 to 1900 would again be 19 percentage points.[31]

The increase in the improvement intensity of farmland came despite substantial amounts of land being converted from the public domain to private farmland. Columns 3 and 4 of Table 3 report the results from estimating equation (8) for the fraction of county land in farms. In these baseline estimates, settlement increased by 26 percentage points from 1880 to 1900 in counties with 0% woodland relative to counties with 6% woodland. There were

---

[29] This 16 percentage point change reflects the increase from 1880 to 1890, plus the increase from 1890 to 1900, minus two times the increase from 1870 to 1880.

[30] This 17 percentage point change reflects the increase from 1880 to 1900 for a county with 0% woodland relative to a county with 6% woodland, minus the change from 1880 to 1900 for a county with 6% woodland relative to a county with 12% woodland.

[31] This 19 percentage point change reflects the increase from 1880 to 1900 for a county with 0% woodland relative to a county with 6% woodland (minus two times the increase from 1870 to 1880), minus the change from 1880 to 1900 for a county with 6% woodland relative to a county with 12% woodland (minus two times the increase from 1870 to 1880).

also some relative increases from 1870 to 1880, however, and from 1880 to 1900 counties with 6% woodland made substantial relative gains on counties with 12% woodland. These estimated changes in settlement are quite large, perhaps implausibly so, and are indeed more sensitive to robustness checks than are the estimated changes in improvement intensity.

Robustness Tests

Table 4 presents the results from specification tests that include control variables meant to capture other changes in Western agricultural development. The results are condensed to show only the estimated changes from 1880 to 1890 and from 1890 to 1900 in counties with 0% woodland relative to counties with 6% woodland. For completeness, the full results for Table 4 are shown in Appendix Tables 1, 2, and 3. As a basis for comparison, column 1 of Table 4 reports the corresponding estimates from Table 3.

(i) Westward Development. One concern is that the baseline estimates could be confounded with an independent push toward increased westward development. From Figure 2, counties with less woodland tend to be further West than counties with more woodland, even within states.

To address this concern, equation (8) is estimated with a control variable for the distance West of each county center, interacted with each decade.[32] Column 2 of Table 4 reports the primary results from this specification, while the full results are reported in columns 2a–2d of Appendix Table 1.

For the improvement intensity of farmland (Panel A), the baseline results are quite robust to controlling for Westward development in each time period. For the amount of farmland (Panel B), the estimated magnitudes fall by roughly half. Indeed, there are substantial Westward settlement effects in each decade from 1870 to 1900 (not shown in the tables). By contrast, there is very little Westward change in the improvement intensity of farmland from 1880 to 1900, though some moderate insignificant increases from 1870 to 1880.

---

[32] To be precise, "distance West" refers to the x-coordinate of the county centroid on an equal area map projection of the United States.

The "Westward development" of the United States may not have been just a move West, but also an expansion from the middle. To allow for this possibility, an additional control variable is included for the distance of each county center to St. Louis (the "Gateway to the West"), interacted with each decade. Column 3 of Table 4 reports these results, which are similar to controlling only for distance West. Note that for both specifications, Appendix Table 1 reports little substantive change in the improvement intensity of farmland before 1880, after 1900, or between higher woodland levels from 1880 to 1900.

Linear distance coefficients are easiest to interpret, but these robustness results are similar when controlling for non-linear changes in Westward development (distance West squared, distance West cubed, distance from St. Louis squared, distance from St. Louis cubed).[33] In these specifications, Colorado and Texas are contributing to the estimation of the distance controls even though these states contribute little to the estimation of the woodland variables. These robustness results are also similar when excluding sample counties in Colorado and Texas, so that the distance controls and woodland variables are estimated primarily on the same sample of counties.

Overall, it does not seem that Westward development accounted for the estimated increases in land improvement from 1880 to 1900 in counties with the least woodland. By contrast, Westward development may account for a substantial portion of the baseline estimated increase in total farmland.

(ii) Regional Development. Another concern is that low-woodland counties may be in different resource regions or soil groups, where they could be affected differently by changes in product prices and technologies [see, e.g., Olmstead and Rhode 2002]. To explore this possibility, the county boundary files were merged with two maps indicating land resource regions and subregions, and great soil groups.[34] Separate variables are defined for the percent

---

[33] Note that whenever I refer to results being similar, that refers both to the coefficients and their precision.

[34] Land resource regions and subregions were mapped in the 1966 U.S. Department of Agriculture Handbook 296. Soil groups were mapped by the U.S. Soil Conservation Service in 1951 and were retained in the National Archives Record Group 114, item 148. These maps were scanned, traced in GIS software, and digitally merged to 1870 county boundaries.

of each county area falling into each region, subregion, or soil group. Within the sample of counties, there are 11 regions, 43 subregions, and 19 soil groups. The soil groups appear also to capture general differences in climate and other environmental factors; in particular, one soil group effectively corresponds to the presence of a major river.

To allow for differential growth patterns across each region, equation (8) is estimated with an additional quadratic time trend for each of the 11 regions.[35] The primary results are reported in column 4 of Table 4. Similarly, columns 5 and 6 of Table 4 report the primary results when controlling for quadratic time trends for each of the 43 subregions and 19 soil groups, respectively.

For the improvement intensity of farmland, the results are fairly robust to each of these specifications. The estimated magnitudes decline most when controlling for quadratic time trends for each of the 43 subregions, but by less than one standard error. The estimates are substantively similar to the baseline, remain statistically significant, and are statistically greater than the estimates at higher woodland levels. Appendix Table 2 reports the full results, which are also similar. For the total amount of farmland, the estimated increases are less robust but still substantial and statistically significant.

Instead of including quadratic time trends, it is possible to include a decade fixed effect for each region, subregion, or soil group. Although demanding of the data, especially for the 43 subregions, the results are mostly consistent with the baseline results.[36]

In summary, it appears that the baseline results cannot be explained by broad differences in the regional balance among counties with different levels of woodland. Estimated changes in improvement intensity are robust to finer regional controls, while changes in total farmland appear to be more sensitive.

---

[35] To be precise, the first-differenced analog of a quadratic time trend is included.

[36] When including region-by-decade controls, the results are similar. When including soil-by-decade controls, the changes from 1880 to 1900 are similar but there is also a substantial increase in improvement intensity from 1870 to 1880. In contrast to the 1880 to 1900 increases, however, this 1870 to 1880 increase is not statistically greater than the increase at higher woodland levels. When including subregion-by-decade controls, improvement intensity has similar increases from 1890 to 1900 but only a small insignificant increase from 1880 to 1890. There is a moderate increase from 1870 to 1880, but not statistically greater than at higher woodland levels.

(iii) Convergence in Agricultural Development. Another concern is that the baseline estimates could be confounded with convergence in agricultural development, whereby counties with lower initial levels of land improvement or settlement may have naturally experienced higher subsequent growth. Counties with little woodland initially lagged behind counties with more woodland, though there were only small relative increases from 1870 to 1880 when these initial differences were greatest.

To address this concern, equation (8) is estimated with an additional fourth-degree polynomial function of the county's fixed 1870 outcome level, interacted with each decade.[37] This effectively restricts the analysis to counties with different woodland levels but similar outcome levels in 1870.[38] The primary results are reported in column 7 of Table 4, while the full results are in Appendix Table 3.

For the improvement intensity of farmland, the estimated increases from 1880 to 1900 are approximately three-fourths their original size, statistically significant, and statistically greater than changes at higher woodland levels (Table 4, column 7, Panel A). These increases reverse the substantial relative decline before barbed wire's introduction and smaller but statistically significant declines after barbed wire's universal adoption (Appendix Table 3, column 1).[39] For the amount of total farmland, by contrast, the estimates are substantially smaller and now statistically insignificant. These increases would only seem substantial if the negative pre-trend from 1870 to 1880 were subtracted from the subsequent changes.

---

[37] The estimates are similar for alternative polynomial functions of the initial outcome level, but it is appealing to use the same function for both the initial outcome level and woodland level.

[38] This robustness check has some statistical appeal, but it is less motivated by economic theory. Two counties with the same initial outcome level and different woodland levels might be expected to differ along important dimensions, in order for farmers in the lower woodland county to be compensated for the lack of woodland. For example, prior to barbed wire, farmers might only improve as much land in a county with little woodland as in a county with more woodland if the county with little woodland had higher quality soil to justify the higher costs of fencing. The ideal empirical analysis would control for all important differences in county characteristics, but this partial statistical correction may be counterproductive in identifying counties with different woodland levels that were otherwise similar.

[39] For the improvement intensity of farmland, it is worth noting the results when controlling for the initial outcome level and implementing the other specification changes discussed above. The results are generally similar to those reported in Appendix Table 3, with the notable exception that the estimated increase from 1880 to 1890 varies between zero and its original baseline estimate.

Overall, these estimates are consistent with previous findings that counties with the least woodland experienced relative increases in land improvement from 1880 to 1900. However, it casts some doubt on whether previously estimated increases in farmland settlement should indeed be attributed to these counties' lower levels of woodland.

(iv) Functional Form and the Exclusion of State-Decade F.E. The baseline estimates are not sensitive to using different polynomial functional forms to model the dependence of agricultural outcomes on woodland. As long as the functional form is sufficiently flexible to capture the basic non-linearity appearing in figure 4, the evaluated effects at benchmark woodland levels are similar.

Figure 6 displays the raw data underlying the main results, stripping away the regression methodology. For each sample county, the change from 1880 to 1900 in the improvement intensity of farmland is plotted against that county's 1880 woodland level. Counties with the least woodland tended to experience large increases from 1880 to 1900, while counties with more woodland were relatively unchanged.

The estimated changes from 1880 to 1900 decrease by roughly 20–40% when the state-by-decade fixed effects are replaced with decade fixed effects, which might be surprising if one expected movement to low-woodland states in response to lower fencing costs. This is entirely driven by Colorado, however, which is far from most previous settlement and has low average woodland. When Colorado is dropped from the sample, the baseline results are similar and the estimates increase 5–20% when the state-by-decade fixed effects are replaced with decade fixed effects.

Summary of Robustness Tests

Throughout the different specifications discussed in these subsections, the baseline results are fairly robust for the estimated changes in the improvement intensity of farmland. Counties with the least woodland generally experienced large relative increases in land improvement from 1880 to 1900. Similar increases did not occur before 1880, after 1900, and did not occur at higher woodland levels from 1880 to 1900. Changes in total farmland are less robust to alternative specifications, in particular, controlling for Westward settlement and convergence in settlement.

V.B Crop Productivity and Crop Choice

All cropland is "improved" by definition, so the previous estimates do not address whether farmers adjusted crop production along the intensive margin. Given the secure protection of property rights, the optimal allocation of land need not favor crops over livestock, as Coase discusses. In the absence of secure property rights, however, the returns to certain crops would likely be more sensitive to the threat of uncompensated damage by others' livestock. Farmers might reduce investments in crop growing, harvest produce early, or otherwise adjust production in ways that contribute to sub-optimal productivity. This section presents results on changes in crop productivity, as a summary measure of all unobserved adjustments in crop production.

Beginning in 1880, county-level data are available on the total production and acreage for each of the six main crops on the Plains (corn, wheat, hay, oats, barley, rye).[40] Productivity for each crop p in each county c is defined as its total production divided by the total acreage devoted to that crop. For ease of interpretation, productivity in each decade is normalized by its value in 1880.[41] Measurement error creates some extreme outliers, so the upper and lower centiles of the normalized productivity distribution are dropped.[42]

To implement the analysis, equation (8) is slightly modified: state-decade fixed effects are replaced with crop-state-decade fixed effects, and the equation is first-differenced by crop-county. For this baseline analysis, data from all crops are stacked so that the change

---

[40] Cotton is excluded from the analysis, as data are only available for Texas and the boll weevil blight severely impacted cotton productivity. Using the same technique as before, the data are adjusted to maintain the 1880 geographical boundaries.

[41] This gives similar results to analyzing changes in the log of productivity, but the estimates are always relative to productivity in 1880.

[42] This drops values less than 0.36 or greater than 6.4. The results are not sensitive to these cutoffs, as long as the clearly extreme observations are dropped.

in productivity across woodland levels is constrained to be the same for all crops:

$$\frac{Y_{pct} - Y_{pc(t-1)}}{Y_{pc1880}} = \alpha_{pst} + \sum_{t=2}^{T}(\beta_{1t}W_c + \beta_{2t}W_c^2 + \beta_{3t}W_c^3 + \beta_{4t}W_c^4) + \epsilon_{pct}. \tag{9}$$

Columns 1 and 2 of Table 5 present baseline results, evaluated at the same representative woodland levels as before. From 1880 to 1890, average productivity across all six crops increased 23.4% more in a county with 0% woodland than in a county with 6% woodland.[43] Over this same period, crop productivity increased by only 1.8% more in a county with 6% woodland than in a county with 12% woodland; a statistically smaller change with a t-statistic of 3.31. To give a sense of the large magnitude of this 23.4% increase, total US crop production increased by roughly 1.7% annually from 1880 to 1920. US crop yields increased by only 0.23% annually from 1880 to 1920, though this is likely to understate productivity growth due to compositional changes in cropland.[44]

In contrast to previous land-use results, these estimated increases did not continue through the 1890s: crop productivity decreased by a statistically insignificant 4.6% from 1890 to 1900, leaving it 18.8% higher than in 1880. One interpretation is that the more-immediate reaction to barbed wire's availability was to secure cropland. Even as barbed wire became more widely used from 1890 to 1900, productivity would only increase if cropland became more secure which, in turn, could only happen if farmers had continued to cultivate substantial amounts of land without fences after 1890.

To explore the timing of these effects, equation (8) is estimated for the fraction of farmland planted in crops. Columns 3 and 4 of Table 5 report the results. From 1880 to 1890, the allocation of farmland to crops increased by 12 percentage points in a county with 0%

---

[43] An alternative specification would weight observations by the number of county acres devoted to that crop in 1880. This specification yields similar estimates, with the coefficient increasing from 23.4% to 26.9%. The standard error changes only minimally, from 5.7% to 5.8%. Without differences in efficiency motivating the weighting, it seems more transparent to focus simply on the unweighted results.

[44] These numbers were estimated using indexes of total US crop production and yield per acre harvested for twelve major crops (NBER Macrohistory Database, files a01005aa and a01297). The production index was computed by weighting the production of each commodity by average farm prices from 1910-1914, and both indexes are then defined relative to a base year. To obtain the average annual increase, the natural log of each index is regressed on a time trend from 1880 to 1920.

woodland relative to a county with 6% woodland, with little change between higher woodland levels. From 1890 to 1900, there was little change in the allocation of farmland to crops.

It appears that farmers' initial reaction (1880–1890) was mainly to expand crop production, both in terms of a greater intensity of cultivation and a greater fraction of farmland allocated to crops. Then, from 1890 to 1900, increases in the improvement intensity of farmland were driven mainly by improvements to pasture land or land otherwise not currently planted in crops.

Extensions and Robustness Tests

(i) At-Risk Crops vs. Hay. A natural extension of these results exploits crop-level differences in vulnerability to livestock damage. While cattle eat hay (various grasses), fields of hay are more resistant to livestock damage before being harvested. Fields of hay can even be intended for grazing at certain times of the year. The other crops (corn, wheat, oats, barley, rye) would yield substantially less grain if they were trampled, so it would be more critical to protect them from others' livestock.[45]

Restricting the analysis to these five crops more at-risk of damage, columns 5–6 report the results from estimating equation (9). From 1880 to 1890, productivity increased by 29% in counties with the least woodland and the results are otherwise similar. By contrast, there was no change in hay productivity from 1880 to 1890 (not shown). Columns 7–8 report the results from estimating equation (8) for the fraction of cropland allocated to these more at-risk crops. Indeed, from 1880 to 1890, more cropland became allocated to these five crops instead of to hay.

(ii) Cropland Composition. Estimated changes in productivity could be confounded by compositional changes in cropland, to the extent that farmland differs in productivity within a county. This seems unlikely, however, to explain the observed large increases in productivity from 1880 to 1890. The typical response would be to expand production into

---

[45] Even controlled grazing would typically lower grain yields by 25–79% [Smith et al. 2004]. Wheat, oats, barley, and rye could be grown for hay instead of for grain, but would then be managed differently. The Census defines data for these crops as that which is grown for grain.

otherwise unprofitable and lower quality land. This would lead the empirical analysis to understate the increase in productivity for a fixed plot of land. If production expanded into less-wooded areas of counties that were inherently more productive, it could contribute to the observed increases in productivity. This is likely of little substantive importance, though, as very little of the land in farms was actually wooded: in 1880, less than 6% of the land in farms was wooded in 76% of the counties with less than 6% woodland. The settlement of new farmland is sometimes associated with short-term productivity benefits, but these increases in productivity mostly persisted.

(iii) Regional Development. In an attempt to account for potential region-specific technological improvements, this subsection briefly reports on robustness checks using the previously defined land resource regions, subregions, and soil groups. The estimated increases in productivity from 1880 to 1890 are robust to the inclusion of crop-specific quadratic time trends by region, subregion, or soil group.[46] The estimated increase in cropland intensity from 1880 to 1890 is robust to including quadratic time trends by region or soil group; by subregion, the increase is reduced by one-third but is statistically significant.[47]

(iv) Westward Development. Another potential concern is that there could be a Westward trend associated with cropland management. The estimated increase in cropland intensity from 1880 to 1890 is robust to controlling for distance West and distance from St. Louis in each decade, with or without the inclusion of Colorado and Texas. Recall that excluding Colorado and Texas forces the distance control variables to be estimated on the same set of counties that mostly identify the changes associated with woodland levels.

When excluding Colorado and Texas, the crop productivity results are robust to including the distance controls in each decade. The productivity results are sensitive, however, to these control variables when including Colorado and Texas: there is no 1880 to 1890 increase and

---

[46] When allowing for crop-decade fixed effects by region or soil group, the 1880 to 1890 increase is roughly half its original size and marginally statistically significant. Productivity across woodland levels is also less stable in subsequent periods. Notably, the results are not robust to controlling for crop-decade-subregion fixed effects, though that is quite demanding of the data.

[47] These results are also similar when replacing the quadratic trends with decade fixed effects by region, soil group, or subregion.

productivity is quite variable in later periods. Given this difference in results, it seems appropriate to place more emphasis on the results when the woodland variables of interest and the distance control variables are primarily estimated on the same sample of observations.

Summary of Extensions and Robustness Tests

Counties with the least woodland appear to experience a substantial increase in crop productivity following the introduction of barbed wire. These productivity increases are concentrated among those crops most at-risk without fencing. Farmland was increasingly allocated to crops and, in particular, more at-risk crops. These results are mostly robust to the inclusion of regional control variables, but there are some notable exceptions: the inclusion of crop-subregion-decade fixed effects, and controlling for distance West and from St. Louis when the sample includes Colorado and Texas.

V.C Land Value

One can gauge the monetary value of barbed wire to farmers, by estimating the resulting change in land values. While land improvement and crop productivity may have increased substantially, these changes would have come at considerable cost to farmers. Changes in land values may capitalize the net benefit to agricultural production, if land markets were well-functioning. In this era, there was a large amount of land speculation and taxes were paid on the value of land [Gates 1973], so it appears plausible that farmers would have some sense of their lands' value when asked by Census enumerators.[48]

In each period, "land value" data are available for the combined value of farmland, buildings, and fences.[49] However, land appears to have been the largest component of this measure. In 1900 and 1910, the value of buildings averaged between 13% and 17% of the total for low-, medium-, and high-woodland counties. Fence stock values are unknown, but

---

[48] If farmers were very quick to update reported valuations, they may have partly anticipated the arrival of barbed wire by the 1880 Census. Unsettled land would still have been valued at zero, though, so much of the response would still not occur until newly desirable lands were settled.

[49] Additional data are obtained from Haines and ICPSR [2005].

the costs of building and repairing fences in 1879 averaged 1% of the total value of land, buildings, and fences.

The natural log of land value is analyzed because technology, prices, and property rights are typically modeled to have multiplicative effects on the total value of production. Additive shocks would have a larger percentage effect in areas with low initial levels, so equation (8) is estimated controlling for differences in initial land values.[50] For example, given similar large increases in settlement across all counties from 1870 to 1880, land values would increase by a larger percentage in low-woodland counties that averaged lower initial land values.

Table 6 presents the results. Land values increased substantially from 1880 to 1900 in counties with 0% woodland relative 6% woodland, though only the change from 1880 to 1890 was statistically greater than the relative increases at higher woodland levels. By contrast, before 1880 and after 1900, there were either relative declines or small changes.

Focusing on the more robust increase from 1880 to 1890, this represents an economically substantial change. The 1880 to 1890 change of 0.406 log points is an increase of 50% above 1880 levels. This is approximately 1.7 times the 1880 value of all agricultural products in low-woodland counties. If we assume that barbed wire had no effect on counties with more than 6% woodland, these estimates imply that the overall monetary benefit of barbed wire to farmers in sample counties was $103 million with a standard error of $32 million, in 1880 US dollars.[51] This is approximately 0.9% of total US GDP in 1880 [Historical Statistics 2006]. Throughout various robustness checks, this total estimated value ranges from $67 million to $139 million and is always statistically significant.[52]

---

[50] Included is an additional fourth-degree polynomial function of log land value in 1870, interacted with each decade. Without these additional controls, the estimates fit a clear pattern of economic convergence for both low- and medium-woodland counties: there are large relative increases from 1870 to 1880 that then decline over time. This specification is subject to the same concerns as before, but it adjusts for this otherwise dominant pattern of convergence.

[51] Sixty-four counties had between 0% and 1% woodland, with an average of 0.42% woodland and $1.7 million of land value in 1880. For an average county with 0.42% woodland, land values increased by an estimated 43% relative to a county with 6% woodland. This gives an overall effect of approximately $47 million (64 × $1.7m × 43%). Summing across the woodland bins (1–2%, 2–3%, 3–4%, 4–5%, 5–6%) yields an estimate of $103 million with a standard error of $32 million.

[52] These include controlling for quadratic time trends or decade fixed effects by region, subregion, or soil group; distance West by decade; distance West and distance from St. Louis by decade.

VI Analysis of Potential Confounding Factors

VI.A Other Shocks Correlated with Woodland Levels

The empirical results would not identify the effects of barbed wire to the extent that other factors caused counties with less woodland to change differently from 1880 to 1900. As a complement to the previous robustness checks, this section provides qualitative evidence on other potential sources of differential changes.

One concern is that a shock to the timber market might have differentially affected farming in counties with different amounts of woodland.[53] However, the timber market in all sample counties was a very small sector of the economy: in 1870, forest products averaged only 0.6% of the total value of all farm products. These counties had such little woodland that any changes in its use would have had limited effects on overall land-use. Land improvement in the Plains did not generally involve clearing woodland because there was already so much open land.

Two other potential confounding factors were the various land acts and the railroads. Due to the lack of fencing materials, however, the 1862 "homestead law was a snare and a delusion" [Webb, p.286]. Subsequent revisionary land acts were generally ill-suited to the region, poorly enforced, and are thought to have played little role in the successful settlement of the Plains [Webb, pp.424–431].[54] Furthermore, most railroad construction in the sample region had been completed by 1873, when the railroad industry faced a panic due to excessive construction and low usage on the Plains.[55] Webb summarizes:

> We have noted that the agricultural frontier came to a standstill about 1850, and that for a generation it made but little advance into the sub-humid region of the Great Plains. It was barbed wire and not the railroads or the homestead law

---

[53] Indeed, barbed wire may have lowered demand for timber by replacing wooden fences. Barbed wire could also have increased demand for timber because of increases in settlement. Farmers would still have required some timber for buildings and fence posts.

[54] See Libecap [2006] for an analysis of land allocation policies in these and more western regions.

[55] The construction of spur lines is quite endogenous to settlement, so it would be difficult to interpret a formal analysis controlling for spur construction even if that data were available.

> that made it possible for farmers to resume, or at least accelerate, their march across the prairies and onto the Plains. [pp.316–317]

Other potential confounding factors are changes in technology. The post-Civil War period saw many technological improvements and mechanization began to increase in the 1900s, but the rise of barbed wire from 1880 to 1900 appears to coincide with a notable lack of other major advances [Rasmussen 1962, Primack 1977]. Olmstead and Rhode discuss biological innovations in wheat, but the 1880 to 1890 period does not stand out. Windmills were used increasingly during this period to obtain water in areas that were often also low in woodland. Knowledge of windmills had existed for a long time, however, and modern industrialized production was widespread by the 1860s. Webb concludes that the increased use of windmills was caused by the introduction of barbed wire and increased cultivation of water-scarce areas [Webb, p.341, p.348]. For the sample region, note also that most conflict with Native Americans had ended by this time.

VI.B Evidence on Potential Other Effects of Barbed Wire

The empirical results would not identify only the effects of property rights if barbed wire had important effects through other channels. Even if land were perfectly secure and there was no threat of uncompensated damage by others' livestock, cheaper fencing might have some effect on agricultural production. If fencing allowed greater control over cattle feeding and breeding, it might increase the productivity of land for raising livestock. By enabling the joint production of cattle and crops, fencing might also improve production possibilities.[56]

If these effects were substantively important, then barbed wire should lead to an increase in cattle production and/or the joint production of cattle and crops. The security of property rights might also affect these outcomes, but these tests are designed to address whether non-property right effects could be driving the earlier results.[57] Empirical estimates suggest

---

[56] For example, if nearby lands varied in their suitability for cattle and crops.

[57] If so, it seems reasonable to expect that non-property right effects should also dominate along these dimensions.

that barbed wire did not increase cattle production and decreased the joint production of cattle and crops. This finding is consistent with the notion that barbed wire mainly affected agricultural outcomes through greater enforcement of property rights.

To test for increases in cattle production, equation (8) is estimated for the number of cattle per five county acres.[58] The estimated magnitudes can be roughly compared to previously estimated changes in land-use, as it requires approximately five acres of land for a cow to graze in this region.

Columns 1 and 2 of Table 7 present the results. From 1880 to 1890, there was no substantial increase in cattle production for a county with 0% woodland relative to a county with 6% woodland. The estimated coefficient implies that approximately 0.39% less of the county land became devoted to cattle production. By contrast, cattle production increased moderately in all subsequent periods and at higher woodland levels from 1880 to 1890. In all, there is little indication that barbed wire specifically led to increased cattle production, especially of a magnitude that could account for the previously estimated effects.[59]

It is less clear how to estimate changes in the joint production of cattle and crops, because the exact amount of land in each county devoted to cattle is unobserved.[60] As a possible solution, I assume that all farmland not used for crops is either used for cattle or other purposes that are not systematically changing over time within states. Based on this assumption, I define a specialization index equal to the squared difference between the percentage of county farmland devoted to crops and the average over all counties in that decade and state: $I_{ct} = (M_{ct} - \overline{M_{st}})^2$. This index increases when a county with an above-average percentage of farmland devoted to crops increases that proportion, and vice versa.

---

[58] Data on cattle are first available in 1880, so county regions are held constant at their 1880 boundaries.

[59] These results are similar when controlling for quadratic time trends by region, region-decade fixed effects, distance West by decade, or distance West and distance from St. Louis by decade. When controlling for quadratic time trends or decade fixed effects by soil group or subregion, there are some increases in cattle production from 1880 to 1890. These increases are never substantively or statistically greater than at higher woodland levels, and the magnitudes are not greater than in later periods.

[60] The assumption of five acres per cow is too rough for these purposes, as acreage requirements vary with the environment, production methods, desired cattle quality, and desired sustainability.

Changes in this index are estimated using equation (8), which controls for average county deviations from the mean and state-by-decade shocks.

Columns 3 and 4 of Table 7 present the results. From 1880 to 1890, counties with 0% woodland became increasingly specialized, relative to counties with 6% woodland. The estimated magnitude is difficult to interpret, but it is large relative to the index mean and standard devation. Throughout various robustness checks, this increase is either similar or larger and always statistically significant.[61] Thus, it does not appear that barbed wire encouraged the joint production of cattle and crops.

Overall, the empirical estimates fail to contradict the hypothesis that barbed wire's effects operated mainly through an increase in property rights enforcement.

VII Evidence on States' Legal Reforms

Prior to barbed wire's introduction, there were contentious political debates over whether settlers should be granted full land protection without having to build a fence. In contrast to the example of a farmer and cattle-raiser considered by Coase, individual cattle-raisers would mainly have been unable to provide effective guarantees to the farmer.[62] Plains state legislatures considered passing "herd laws" that would make livestock owners formally responsible for damages to farmers' unfenced crops.

These legislative debates highlight the property rights problem, but contentious debates do not imply that legislative responses could be effective. "Herd laws" would only be effective if enforced sufficiently to provide the credible promise of compensation if farmers' crops were damaged. Webb argues that formal laws often had little practical influence on the Plains.

Table 8 reports the results from separately estimating equation (8) in Kansas and Nebraska, two states to implement versions of these reforms described by Kawashima and Davis

---

[61] These include controlling for quadratic time trends or decade fixed effects by region, subregion, or soil group; distance West by decade; distance West and distance from St. Louis by decade.

[62] Because of the large number of current and potential cattle-raisers, transaction costs were prohibitive.

[1973].[63] The empirical estimates are less precise within individual states, but the bulk of the evidence indicates that the legal reforms were not effective.

Nebraska adopted a state-wide herd law in 1871, which should have benefited farmers more in counties with less woodland and higher fencing costs. From 1870 to 1880, however, the improvement intensity of farmland declined substantially in a county with 0% woodland relative to a county with 6% woodland. It was not until 1890, after the introduction of barbed wire, that counties with the least woodland showed substantial increases in land improvement. Note that the coefficients are imprecisely estimated at higher woodland levels because 90% of the sample counties in Nebraska had less than 6% woodland.

Kansas gave counties the option of adopting a herd law in 1872. I do not know exactly which counties adopted the herd law, but it was generally not adopted in counties with little woodland where the small proportion of crop growers had weak political influence [Webb, p.500; Sanchez and Nugent 2000; Kawashima]. Settlement increased more in lower-woodland counties from 1870 to 1880, however, and there was no change in land improvement. It was not until 1880 to 1890, after the introduction of barbed wire, that land improvement increased in counties with the least woodland. Kansas then adopted a state-wide herd law in 1889, which should have benefited farmers in lower-woodland counties, but land improvement declined from 1890 to 1900.

Davis argues that herd laws provided many of the benefits that were subsequently reinforced by the introduction of barbed wire, but most evidence comes from newspaper advocacies for legal reform. There is little evidence that the laws were strictly enforced and, as one would expect, cattle owners resisted making compensating payments for damages. Once barbed wire was introduced, those same newspapers wrote that "every farm needs *some* fencing" and as "soon as a farmer is able, he fences his farm. There must be an apparent benefit" [Davis, p.134]. Farmers made substantial investments in fencing before the legal reforms, after the legal reforms, and after the introduction of barbed wire. These laws

---

[63] Texas also adopted a county-option policy in 1870, but the results are mostly noise because the sample has only a small number of counties in Texas with little woodland.

may have had some small influence or they might not have been so hotly debated, but their passage appears to reflect changes in local political power rather than a major influence on farmers' decisions. In the absence of physical barriers, formal laws appear to have provided farmers little refuge from roaming livestock.

VIII Conclusion

In the 19th century American Plains, fences were required to secure farmland. Initially, fences were prohibitively expensive in areas with little local woodland. In this historical context, the introduction of barbed wire fencing encouraged property rights protection most in areas with the least woodland. This paper tests the hypothesis that this increase in property rights protection played an important role in the agricultural development of the American Plains.

Between barbed wire's introduction and its universal adoption, counties with the least woodland are estimated to have experienced substantial relative increases in agricultural development. Increases along *intensive* margins appear to have been dominant: farmland became more intensely improved, a greater percentage of farmland was allocated to crops, and crop productivity increased. Estimated increases in land values suggest that these effects were associated with substantial economic gains. On balance, these findings appear to be robust to a range of specification tests.

Overall, the results here suggest that secure property rights over land may have an important role in facilitating agricultural development. In contrast to the typical policy reform, however, these results reflect changes in property rights that largely fall outside formal institutional channels. Indeed, the apparent ineffectiveness of contemporaneous legal reforms may indicate a difficulty protecting property rights in sparsely settled areas through legislation and the court system.

Data Appendix

The data was converted into time-invariant geographical units using ArcView (GIS) software and the Historical United States Boundary Files [Carville et al. 1999].[64] The conversion process is designed to obtain data for each 1870 county when the data are subsequently reported for different county definitions. This turns out not to be particularly important for the empirical results, because most counties that change borders are excluded from the sample for other reasons (Figure 2). Still, this process is qualitatively important for the data to make sense and to give each original 1870 county equal weight.

This process is now described for converting 1920 data into the 1870 county definitions. The 1870 map is intersected with the 1920 map and each 1920 county is assigned the identification number of its 1870 county origin. When the 1920 county falls within more than one 1870 county, each piece of the 1920 county is assigned its unique 1870 origin along with the area of that piece. This information is then matched to the Census data, where the adjusted data for each 1920 piece are found by multiplying the original 1920 Census value by that piece's relative size. Finally, the 1920 data for each 1870 county region are found by summing across all 1920 pieces that make up the region.[65]

When counties in later periods fall within more than one original 1870 region, it is impossible to assign the exact measure of the variable in that period to each original region. This procedure assumes that each variable is evenly distributed across the county area. This introduces no additional measurement error for the typical changes in county borders over this time, whereby a single county split into subsections. In practice, the digital map projections are imperfect, resulting in small sections of counties being assigned to different original regions, but the effect on the total variable measures is very small. For the six states of interest, roughly 85% of counties in every period have less than 1% of their area in a second original region.

---

[64] The map files are available at the beginning of each decade. In each period, 3–9 census counties are dropped because they do not match to the map files.

[65] The summed measure is considered to be missing when data for any piece is missing.

In the final data, counties will not have the exact same number of acres in different periods. Slight changes in the map projections or data collection procedures would create small differences. There are some very large differences, by hundreds of thousands or millions of acres, that must be due to data coding errors or mismatches regarding when large border changes took place. Counties for which the final standard deviation of the number of acres is greater than 50,000 are dropped (3% of the sample).

References

Acemoglu, Daron and Johnson, Simon. 2005. "Unbundling Institutions." *Journal of Political Economy*. Vol. 113, No. 5, pp. 949-995.

Alston, Lee J., Gary D. Libecap, and Bernardo Mueller. 1998. "Property Rights and Land Conflict: A Comparison of Settlement of the U.S. Western and Brazilian Amazon Frontiers," in Latin America and the World Economy Since 1800, ed. John H. Coatsworth and Alan M. Taylor. Cambridge: Harvard University Press: pp. 55-84.

Anderson, Terry L. and Hill, Peter J. 1975. "The Evolution of Property Rights: A Study of the American West." *Journal of Law and Economics*. Vol. 18, No. 1, pp. 163-179.

Banerjee, Abhijit V.; Gertler, Paul J.; Ghatak, Maitreesh. 2002. "Empowerment and Efficiency: Tenancy Reform in West Bengal." *Journal of Political Economy*. Vol. 110, No. 2, pp. 239-280.

Besley, Timothy. 1995. "Property Rights and Investment Incentives: Theory and Evidence from Ghana." *Journal of Political Economy*. Vol. 103, No. 5, pp. 903-937.

Bogue, Allan G. 1963a. From Prairie to Corn Belt: Farming on the Illinois and Iowa Prairies in the Nineteenth Century. University of Chicago Press.

Bogue, Allan G. 1963b. "Farming in the Prairie Peninsula, 1830-1890." *Journal of Economic History*. Vol. 23, No. 1, pp. 3-29.

Brasselle, Anne-Sophie; Gaspart, Frederic; Platteau, Jean-Philippe. 2002. "Land Tenure Security and Investment Incentives: Puzzling Evidence from Burkina Faso." *Journal of Development Economics*. Vol. 67, No. 2, pp. 373-418.

Carville, Earle; Heppen, John; Otterstrom, Samuel. 1999. HUSCO 1790-1999: historical United States county boundary files [electronic resource]. Louisiana State University (Baton Rouge, LA): Geoscience. http://www.ga.lsu.edu/husco.html

Cheung, Steven N.S. 1970. "The Structure of a Contract and the Theory of a Non-Exclusive Resource." *Journal of Law and Economics*. Vol. 13, No. 1, pp. 49-70.

Coase, Ronald H. 1960. "The Problem of Social Cost." *Journal of Law and Economics*. Vol. 3, No. 1, pp. 1-44.

Davis, Rodney O. 1973. "Before Barbed Wire: Herd Law Agitations in Early Kansas and Nebraska," in Essays in American History in Honor of James C. Malin, ed. Burton J. Williams. Coronado Press.

Demsetz, Harold. 1967. "Toward a Theory of Property Rights." *American Economic Review: Papers and Proceedings*. Vol. 57, No. 2, pp. 347-359.

Encyclopedia Britannica, eleventh edition. 1911. "Barbed Wire." http://1911encyclopedia.org/B/BA/BARBED_WIRE.htm

Gates, Paul W. 1973. "The Role of the Land Speculator in Western Development," in Landlords and Tenants on the Prairie Frontier, pp. 48-71. Cornell University Press.

Goldstein, Markus and Udry, Christopher. 2005. "The Profits of Power: Land Rights and Agricultural Investment in Ghana." Mimeo, Yale University, November.

Gutmann, Myron P. 2005. "Great Plains Population and Environment Data: Agricultural Data [Computer file]." Ann Arbor, MI: University of Michigan / Ann Arbor, MI: ICPSR.

Hayter, Earl W. 1939. "Barbed Wire Fencing – A Prairie Invention." *Agricultural History*. Vol. 13, pp. 189-207.

Haines, Michael R., and the Inter-university Consortium for Political and Social Research. 2005. Historical, Demographic, Economic, and Social Data: The United States, 1790-2000 [Computer file]. ICPSR02896-v2. Hamilton, NY: Colgate University / Ann Arbor, MI: ICPSR.

Jacoby, Hanan G.; Li, Guo; Rozelle, Scott. 2002. "Hazards of Expropriation: Tenure Insecurity and Investment in Rural China." *American Economic Review*. Vo. 92, No. 5, pp. 1420-1447.

Kawashima, Yasuhide. 1994. "Fence Laws on the Great Plains, 1865-1900," in Essays on English Law and the American Experience, eds. Elizabeth A. Cawthorn and David E. Narrett. Walter Prescott Webb Memorial Lectures, No. 27, pp. 100-119. Texas A&M University Press.

Kraenzel, Carl Frederick. 1955. The Great Plains in Transition. University of Oklahoma Press.

Libecap, Gary D. 2006. "The Assignment of Property Rights on the Western Frontier: Lessons for Contemporary Environmental and Resource Policy." NBER Working Paper 12598.

McFadden, Joseph M. 1978. "Monopoly in Barbed Wire: The Formation of the American Steel and Wire Company." *Business History Review*, Vol. 52, No. 4, pp. 465-489.

McCallum, Henry D. and McCallum, Frances T. 1965. The Wire that Fenced the West. University of Oklahoma Press: Norman. 1972 printing.

Miceli, Thomas J.; Sirmans, C.F.; Kieyah, Joseph. 2001. "The Demand for Land Title Registration: Theory with Evidence from Kenya." *American Law and Economics Review*. Vol. 3, No. 2, pp. 275-287.

North, Douglass C. 1981. Structure and Change in Economic History. W.W. Norton & Company.

Olmstead, Alan L. and Paul W. Rhode. 2002. "The Red Queen and the Hard Reds: Productivity Growth in American Wheat, 1880-1940." *Journal of Economic History*. Vol. 62, No. 4, pp. 929-966.

Primack, Martin L. 1966. "Farm Capital Formation as a Use of Farm Labor in the United States, 1850-1910." *Journal of Economic History*. Vol. 26, No. 3, pp. 348-362.

Primack, Martin L. 1969. "Farm Fencing in the Nineteenth Century." *Journal of Economic History*. Vol. 29, No. 2, pp. 287-291.

Primack, Martin L. 1977. Farm Formed Capital in American Agriculture 1850-1910. 1962 Dissertation. Arno Press: New York.

Rasmussen, Wayne D. 1962. "The Impact of Technological Change on American Agriculture, 1862-1962." *Journal of Economic History*. Vol. 22, No. 4, pp. 578-591.

Sanchez Nicolas and Nugent, Jeffrey B. 2000. "Fence Laws vs. Herd Laws: A Nineteenth-Century Kansas Paradox." *Land Economics*. Vol. 76, No. 4, pp. 518-533.

Smith, S. Ray; Benson, Brinkley; Thomason, Wade. 2004. "Growing Small Grains for Forage in Virginia." Virginia Cooperative Extension, Publication No. 424-006.

U.S. House of Representatives, 42nd Congress, 1st session. 1871. "Special Report on Immigration; accompanying Information for Immigrants…" by Edward Young, Chief Bureau of Statistics. Washington: Government Printing Office.

U.S. House of Representatives, 42nd Congress, 2nd session. 1872. "Statistics of fences in the United States." *House Executive Documents*. Vol. 17, No. 327 (Report of the Commissioner of Agriculture), pp. 497-512. Washington: Government Printing Office. (US Congressional Serial Set Vol. 1522).

Webb, Walter Prescott. 1931. The Great Plains. Grosset & Dunlap: New York. 1978 printing.

Table 1. Statistics on Barbed Wire Production, Fence Stocks, and New Fence Construction

A. Annual Production. thousands of tons										
1911 Encyclopedia	1874	1875	1876	1877	1878	1879	1880	1890	1900	1907
Britannica	0.005	0.3	1.5	7	13	25	40	125	200	250
Webb 1931, p.309	1874	1875	1876	1877	1878	1879	1880		1901	
	0.005	0.3	1.3	6	12	23	37		~ 250	
Hayter 1939							1880-1884	1888	1895	
							80 - 100	150	157	

B. Fence Stocks, millions of rods (1 rod ≈ 5 meters)							
North Central	1850	1860	1870	1880	1890	1900	1910
Total	228	303	359	427	443	493	483
Wood	226	285	320	369	279	192	75
Stone	2	3	3	3	0	0	0
Hedge	0	9	22	26	27	30	27
Wire	0	6	14	30	137	271	382
South Central							
Total	175	230	245	344	531	685	701
Wood	171	219	235	330	425	411	280
Stone	3	5	4	7	0	0	0
Hedge	0	5	2	3	0	0	0
Wire	0	2	3	3	106	274	420
Prairie							
Total	5	22	41	80	255	607	718
Wood	4	17	23	40	130	176	7
Stone	0	1	2	2	0	0	0
Hedge	0	3	13	26	3	18	22
Wire	0	1	4	12	122	413	689
Southwest							
Total	39	78	94	162	280	710	749
Wood	38	71	80	123	174	312	187
Stone	1	2	2	2	0	0	0
Hedge	0	2	4	5	0	0	0
Wire	0	4	9	32	106	398	562

C. New Fence Construction, percentage						
North Central	1850-59	1860-69	1870-79	1880-89	1890-99	1900-09
Wood	79	66	73	3	0	0
Stone	1	0	0	0	0	0
Hedge	12	21	6	1	0	0
Wire	8	13	22	96	100	100
South Central						
Wood	90	94	100	50	0	0
Stone	2	0	0	0	0	0
Hedge	4	1	1	0	0	0
Wire	3	5	0	50	100	100
Prairie						
Wood	71	39	38	45	18	0
Stone	5	4	0	0	0	0
Hedge	18	45	38	0	0	0
Wire	6	12	24	55	82	100
Southwest						
Wood	84	56	63	42	32	0
Stone	2	3	0	0	0	0
Hedge	5	11	2	0	0	0
Wire	9	29	35	58	68	100

Notes: “Wood” fences include three types: Virginia worm, post and rail, and board. “Wire” fences are the smooth iron variety from 1850-1870 and include barbed wire from 1880 on. Panel B is excerpted from Primack 1977, table 23, pp. 206-208. Panel C is excerpted from Primack 1977, table 26, pp. 83-84.

Table 2. Mean County Characteristics in 1880, by County Woodland Group

	All Counties (1)	Low Woodland, 0% – 4% (2)	Medium Woodland, 4% – 8% (3)	High Woodland, 8% – 12% (4)	P-value (2) vs. (3) (5)	P-value (3) vs. (4) (6)
1870 county boundaries						
# counties	377	147	57	43	--	--
Land-use outcomes:						
Acres of land in farms,	0.53	0.42	0.59	0.65	0.000	0.257
per county acre	(0.26)	(0.26)	(0.28)	(0.26)		
Acres of improved land,	0.54	0.55	0.64	0.65	0.032	0.835
per acre in farms	(0.23)	(0.20)	(0.28)	(0.24)		
Other characteristics:						
Acres in county	550,718	645,898	470,219	430,631	0.018	0.348
	(526,638)	(801,941)	(239,127)	(180,333)		
Acres of land in farms	237,407	188,967	242,318	254,210	0.002	0.548
	(113,987)	(125,574)	(104,289)	(92,210)		
Acres of improved land	133,995	108,404	170,818	177,500	0.000	0.751
	(91,561)	(82,723)	(110,841)	(98,161)		
Value of land, buildings,	3,192,401	2,294,290	4,271,744	4,776,354	0.000	0.467
and fences	(2,851,257)	(1,750,834)	(3,635,498)	(3,245,953)		
Value of all products	838,986	593,556	1,060,637	1,158,761	0.000	0.548
	(659,730)	(452,311)	(867,683)	(756,448)		
Cost of building and	33,514	23,267	41,589	40,162	0.000	0.782
repairing fences	(24,120)	(18,173)	(31,547)	(19,616)		
1880 county boundaries						
# counties	490	246	61	44	--	--
Acres of cropland,	0.31	0.29	0.38	0.42	0.009	0.415
per acre in farms	(0.20)	(0.19)	(0.25)	(0.19)		
Acres of cropland	71,436	54,350	102,594	114,472	0.000	0.448
	(65,041)	(54,606)	(85,283)	(73,867)		
% acreage for each crop:						
Corn	40.2	34.3	51.1	42.6	0.000	0.060
	(22.4)	(23.6)	(24.0)	(21.3)		
Wheat	23.2	28.5	22.6	24.1	0.033	0.711
	(19.6)	(19.0)	(19.5)	(19.7)		
Hay	18.3	26.2	13.7	14.6	0.000	0.662
	(20.5)	(24.5)	(9.2)	(10.8)		
Oats	7.6	8.2	6.7	8.6	0.050	0.037
	(6.5)	(7.9)	(4.3)	(4.6)		
Barley	1.0	1.4	1.2	1.0	0.589	0.664
	(1.7)	(1.8)	(2.6)	(1.6)		
Rye	0.5	0.6	0.6	0.5	0.773	0.499
	(0.9)	(1.0)	(0.9)	(0.7)		

Notes: In the top panel, the sample is the same as in Figure 2 and Table 3 (377 counties in a balanced panel). In the bottom panel, the sample is the same as in Tables 5 and 7 (490 counties in a balanced panel). Missing data for crop acreage is treated as a zero. P-values are calculated based on standard errors that are adjusted for heteroskedasticity.

Table 3. Estimated Changes in the Improvement Intensity of Farmland and Settlement, Evaluated at Representative Woodland Levels

		Acres of Improved Land, per acre in farms		Acres of Land in Farms, per county acre	
	Woodland:	0% vs. 6% (1)	6% vs. 12% (2)	0% vs. 6% (3)	6% vs. 12% (4)
	Decade:				
Before Barbed Wire	1870 - 1880	0.015	0.023	0.039	0.028**
		(0.040)	(0.013)	(0.026)	(0.010)
		[0.22]		[0.41]	
After Barbed Wire's Introduction	1880 - 1890	0.100**	- 0.004	0.129**	0.048**
		(0.030)	(0.012)	(0.023)	(0.009)
		[4.12]		[3.81]	
	1890 - 1900	0.086**	0.020*	0.128**	0.057**
		(0.020)	(0.009)	(0.026)	(0.009)
		[3.40]		[3.13]	
After Barbed Wire's Universal Adoption	1900 - 1910	- 0.019	0.004	- 0.022	- 0.014
		(0.011)	(0.006)	(0.024)	(0.009)
		[2.25]		[0.36]	
	1910 - 1920	0.003	0.006	0.024	- 0.003
		(0.010)	(0.004)	(0.016)	(0.007)
		[0.31]		[1.59]	
	R^2	0.4432		0.5012	
	Observations	1885		1885	

Notes: For the indicated outcome variable, this table reports estimates from equation 8 in the text. The estimated polynomial function of woodland is evaluated at three woodland levels (0%, 6%, 12%) and the reported estimates represent the difference in predicted changes between the two indicated woodland levels. Each cell can be interpreted as a difference-in-difference coefficient: the change over that period for an average county with 0% woodland relative to the change for a county with 6% woodland (or a county with 6% woodland relative to a county with 12% woodland). Robust standard errors clustered by county are reported in parentheses: ** denotes statistical significance at 1% and * at 5%. In brackets is the t-statistic of the difference between the two contemporaneous estimates comparing changes at different woodland levels.

Table 4. Estimated 1880-1890 and 1890-1900 Changes for 0% Woodland Counties vs. 6% Woodland Counties, Robustness to Different Specifications

		Additional Controls for:					
	Baseline Specification (1)	Distance West (2)	Distance West and Distance From St. Louis (3)	Quadratic Time Trends by Region (4)	Quadratic Time Trends by Subregion (5)	Quadratic Time Trends by Soil Group (6)	1870 Outcome Differences (7)
Panel A: Acres of Improved Land, per acre in farms							
1880 - 1890	0.100**	0.094**	0.119**	0.100**	0.087**	0.120**	0.065*
	(0.030)	(0.034)	(0.031)	(0.029)	(0.031)	(0.031)	(0.033)
	[4.12]	[3.76]	[4.54]	[3.76]	[3.29]	[4.36]	[2.90]
1890 - 1900	0.086**	0.094**	0.106**	0.087**	0.070**	0.094**	0.076**
	(0.020)	(0.026)	(0.026)	(0.020)	(0.021)	(0.021)	(0.021)
	[3.40]	[3.33]	[3.64]	[3.22]	[2.61]	[3.50]	[2.95]
R^2	0.4432	0.4458	0.4589	0.4773	0.5095	0.4864	0.5735
Observations	1885	1885	1885	1885	1885	1885	1885
Panel B: Acres of Land in Farms, per county acre							
1880 - 1890	0.129**	0.068*	0.083**	0.106**	0.078**	0.121**	0.031
	(0.023)	(0.027)	(0.026)	(0.025)	(0.024)	(0.025)	(0.029)
	[3.81]	[1.95]	[2.33]	[2.48]	[1.87]	[3.54]	[1.26]
1890 - 1900	0.128**	0.079*	0.066*	0.116**	0.097**	0.122**	0.062
	(0.026)	(0.031)	(0.032)	(0.026)	(0.027)	(0.026)	(0.039)
	[3.13]	[1.56]	[1.27]	[2.48]	[1.92]	[2.91]	[0.99]
R^2	0.5012	0.5583	0.5716	0.5418	0.5737	0.5352	0.5606
Observations	1885	1885	1885	1885	1885	1885	1885

Notes: This table reports on a series of specification tests for the results in Table 3. Only the primary results are shown here: the estimated changes from 1880-1890 and 1890-1900 for counties with 0% woodland relative to counties with 6% woodland. The full results are shown in Appendix Tables 1–3. For comparison, the results from Table 3 are reported in column 1. Column 2 controls for a county's distance West, interacted with each decade. Column 3 controls for a county's distance West and distance from St. Louis, interacted with each decade. Column 4 controls for a quadratic time trend for each of 11 Regions. Column 5 controls for a quadratic time trend for each of 43 Subregions. Column 6 controls for a quadratic time trend for each of 19 Soil Groups. Column 7 controls for a fourth-degree polynomial of a county's 1870 outcome level, interacted with each decade.

Table 5. Estimated Changes in Crop Productivity and Crop Intensity

		All Crops				At-Risk Crops			
		Productivity		Acres of Cropland, per acre in farms		Productivity		Acres of At-Risk Crops, per acre of cropland	
	Woodland:	0% vs. 6% (1)	6% vs. 12% (2)	0% vs. 6% (3)	6% vs. 12% (4)	0% vs. 6% (5)	6% vs. 12% (6)	0% vs. 6% (7)	6% vs. 12% (8)
	Decade:								
After Barbed Wire's Introduction	1880 - 1890	0.234**	0.018	0.121**	0.015**	0.292**	0.036	0.058*	0.000
		(0.057)	(0.025)	(0.015)	(0.005)	(0.067)	(0.027)	(0.023)	(0.007)
		[3.31]		[8.01]		[3.47]		[2.84]	
	1890 - 1900	– 0.046	– 0.003	– 0.013	– 0.011**	– 0.057	– 0.012	0.007	– 0.005
		(0.039)	(0.018)	(0.009)	(0.004)	(0.046)	(0.021)	(0.015)	(0.006)
		[0.97]		[0.24]		[0.86]		[0.91]	
After Barbed Wire's Universal Adoption	1900 - 1910	0.054	– 0.019	0.039**	0.004	0.058	– 0.027	0.020	– 0.015
		(0.035)	(0.018)	(0.010)	(0.004)	(0.039)	(0.019)	(0.019)	(0.009)
		[1.70]		[3.44]		[1.81]		[1.75]	
	1910 - 1920	– 0.036	0.014	0.010	0.001	0.011	0.038	0.016	0.005
		(0.042)	(0.025)	(0.007)	(0.004)	(0.049)	(0.029)	(0.015)	(0.009)
		[0.95]		[1.06]		[0.44]		[0.63]	
	R^2	0.3949		0.3922		0.4000		0.2777	
	Observations	9104		1960		7320		1960	

Notes: For changes in productivity, estimates are from equation 9 in the text. For changes in cropland, estimates are from equation 8 in the text. Estimated coefficients are presented in the same form as in Table 3.

Table 6. Estimated Changes in Land Value

		Value of Land in Farms, per county acre	
	Woodland:	0% vs. 6% (1)	6% vs. 12% (2)
	Decade:		
Before Barbed Wire	1870 - 1880	– 0.364*	– 0.224**
		(0.151)	(0.053)
		[1.11]	
After Barbed Wire's Introduction	1880 - 1890	0.406**	0.074*
		(0.105)	(0.037)
		[3.99]	
	1890 - 1900	0.213**	0.126**
		(0.072)	(0.027)
		[1.45]	
After Barbed Wire's Universal Adoption	1900 - 1910	– 0.101	0.013
		(0.073)	(0.027)
		[2.07]	
	1910 - 1920	0.044	0.018
		(0.063)	(0.024)
		[0.50]	
	R^2	0.7569	
	Observations	1880	

Notes: Estimates are from equation 8 in the text and are reported in the same form as in Table 3.

Table 7. Estimated Changes in Cattle Production and County Specialization

		Number of Cattle, per five county acres		Degree of Specialization in Crops or Cattle	
	In 1880:				
	Mean	0.2034		0.0175	
	Std. deviation	0.1570		0.0248	
	Woodland levels:	0% vs. 6% (1)	6% vs. 12% (2)	0% vs. 6% (3)	6% vs. 12% (4)
	Decade:				
After Barbed Wire's Introduction	1880 - 1890	−0.0039	0.0247**	0.0117**	−0.0007
		(0.0137)	(0.0066)	(0.0039)	(0.0012)
		[2.01]		[3.50]	
	1890 - 1900	0.0567**	0.0355**	0.0007	−0.0070**
		(0.0150)	(0.0071)	(0.0032)	(0.0020)
		[1.29]		[2.05]	
After Barbed Wire's Universal Adoption	1900 - 1910	0.0437**	−0.0122	−0.0023	0.0021
		(0.0140)	(0.0065)	(0.0025)	(0.0014)
		[3.50]		[1.48]	
	1910 - 1920	0.0348**	−0.0064	−0.0015	−0.0023
		(0.0102)	(0.0052)	(0.0020)	(0.0014)
		[3.32]		[0.28]	
	R^2	0.5823		0.1681	
	Observations	1960		1960	

Notes: For the indicated outcome variable, estimates are from equation 8 in the text and are reported in the same form as in Table 3.

Table 8. Estimated Changes in Improvement Intensity and Settlement, Nebraska and Kansas

	Acres of Improved Land, per acre in farms		Acres of Land in Farms, per county acre	
Woodland:	0% vs. 6% (1)	6% vs. 12% (2)	0% vs. 6% (3)	6% vs. 12% (4)
Panel A: Nebraska				
1870 - 1880	– 0.216*	0.301	0.034	0.640
	(0.104)	(0.367)	(0.174)	(0.742)
1880 - 1890	0.294**	– 0.350	0.060	0.206
	(0.062)	(0.299)	(0.063)	(0.189)
1890 - 1900	– 0.069*	0.335*	0.175**	– 0.120
	(0.033)	(0.161)	(0.057)	(0.239)
1900 - 1910	0.063	– 0.339	– 0.008	0.220
	(0.035)	(0.181)	(0.043)	(0.148)
1910 - 1920	– 0.001	0.014	0.033	– 0.011
	(0.026)	(0.075)	(0.019)	(0.107)
R^2	0.8129		0.7534	
Observations	155		155	
Panel B: Kansas				
1870 - 1880	– 0.027	0.011	0.479**	0.045
	(0.163)	(0.045)	(0.102)	(0.035)
1880 - 1890	0.307**	0.053*	– 0.135*	0.116**
	(0.060)	(0.026)	(0.066)	(0.030)
1890 - 1900	– 0.091	0.008	0.017	0.027
	(0.051)	(0.028)	(0.031)	(0.024)
1900 - 1910	0.017	– 0.018	0.080*	– 0.023
	(0.041)	(0.026)	(0.037)	(0.030)
1910 - 1920	0.028	– 0.018	0.037	– 0.029*
	(0.030)	(0.015)	(0.019)	(0.013)
R^2	0.5283		0.7959	
Observations	255		255	

Notes: Panels A and B report the estimated effects from Table 3 for Nebraska only and Kansas only, respectively.

Figure 1. Declining Steel Prices and the Introduction of Barbed Wire

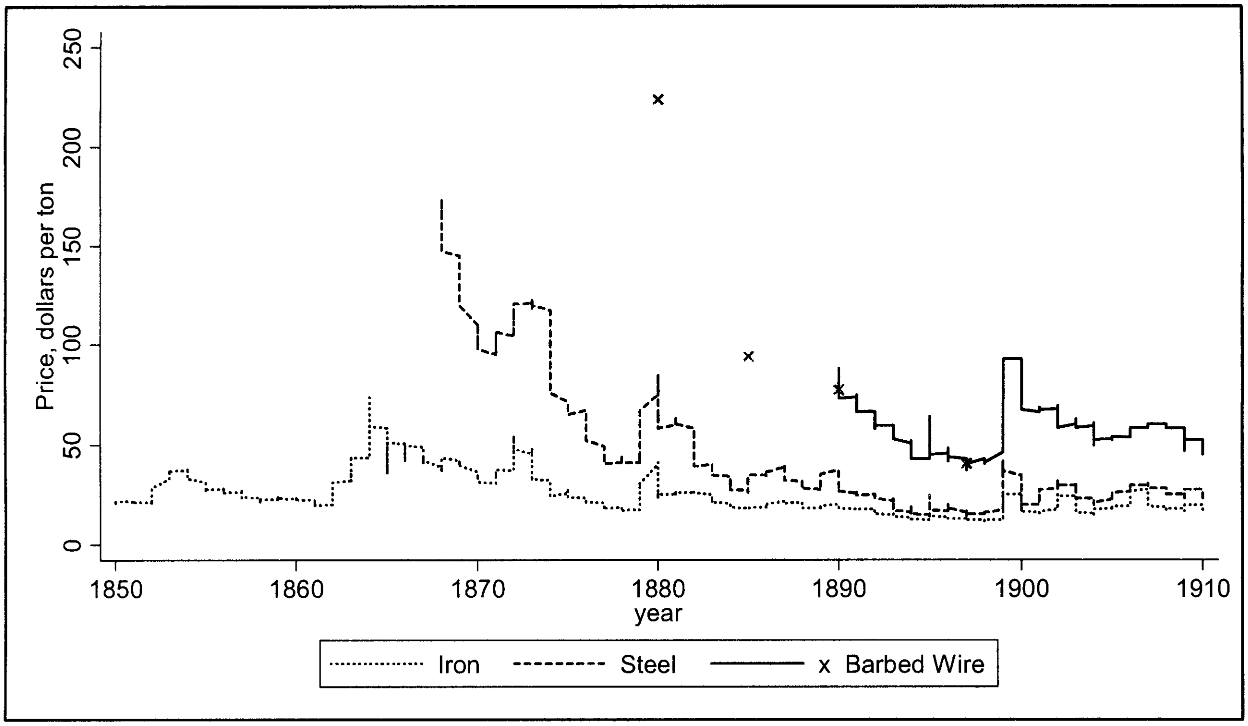

Notes: Nominal price data was obtained from the NBER Macrohistory Database. "Barbed Wire" is the price of galvanized barbed wire in Chicago, with four additional prices denoted by X's [Webb; Hayter]. In 1874, the price of barbed wire was approximately $450 per ton. Before 1890, "Steel" is the price of Bessemer steel rails in PA. After 1890, "Steel" is the price of Bessemer steel billets in PA. "Iron" is the price of pig iron in PA.

Figure 2. Sample Counties by Defined Woodland Levels

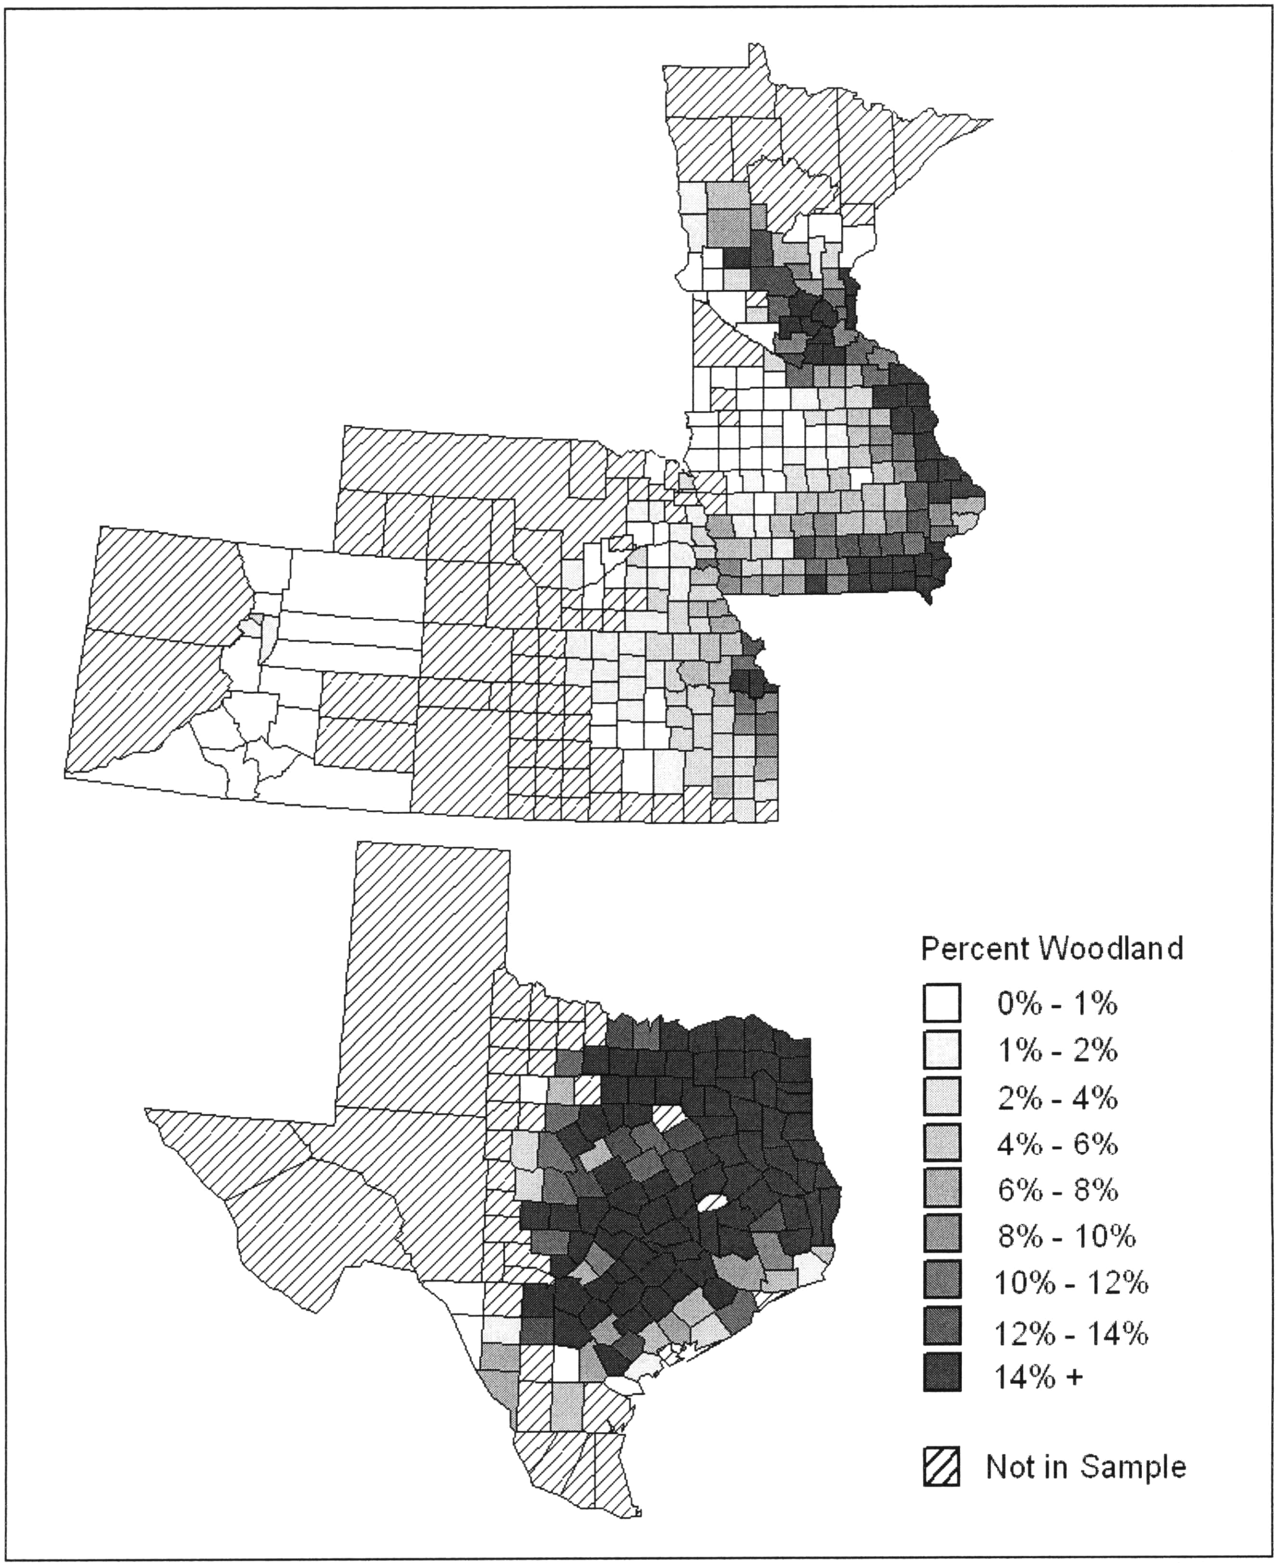

Notes: This figure displays the 377 sample counties based on 1870 geographical boundaries, shaded to represent their level of woodland.

Figure 3. Acres of Improved Land (per farm acre), by County Woodland Group and Decade

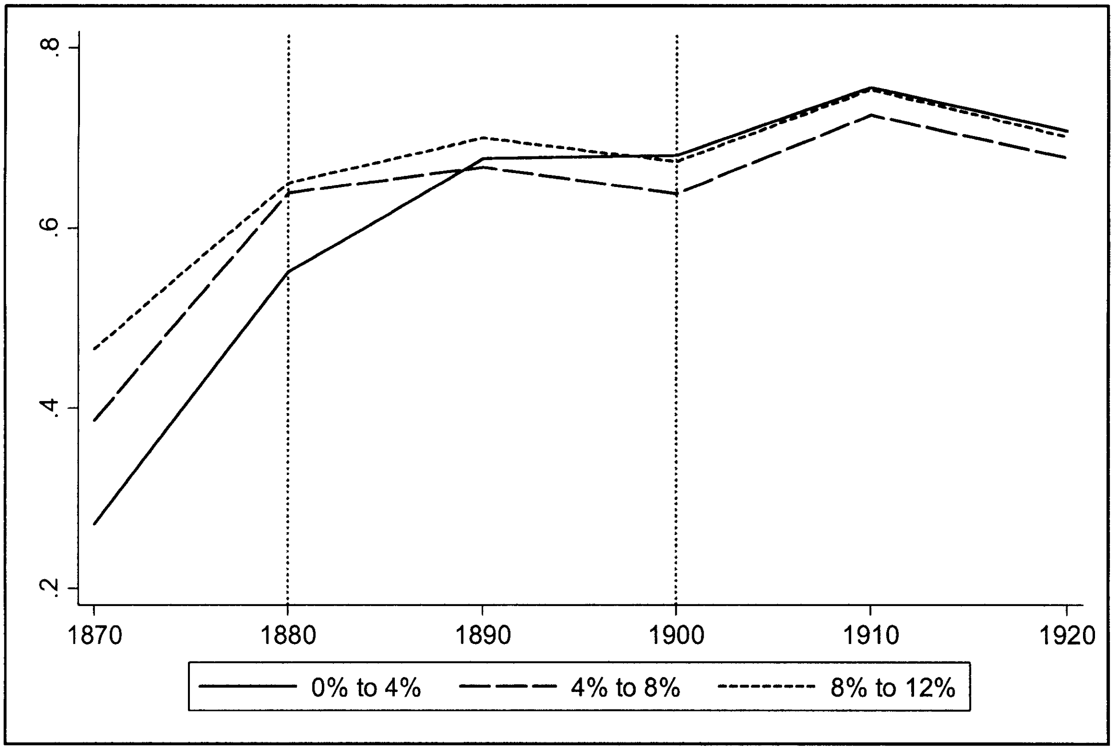

Notes: For counties in these three woodland groups, this figure displays the average number of improved acres of land per acre of farmland in each time period. The two vertical dotted lines represent the approximate date of barbed wire's introduction to farmers (1880) and barbed wire's universal adoption by farmers (1900).

Figure 4. Estimated Changes in the Improvement Intensity of Farmland, +/- 2 Standard Errors, Relative to a County with 0% Woodland

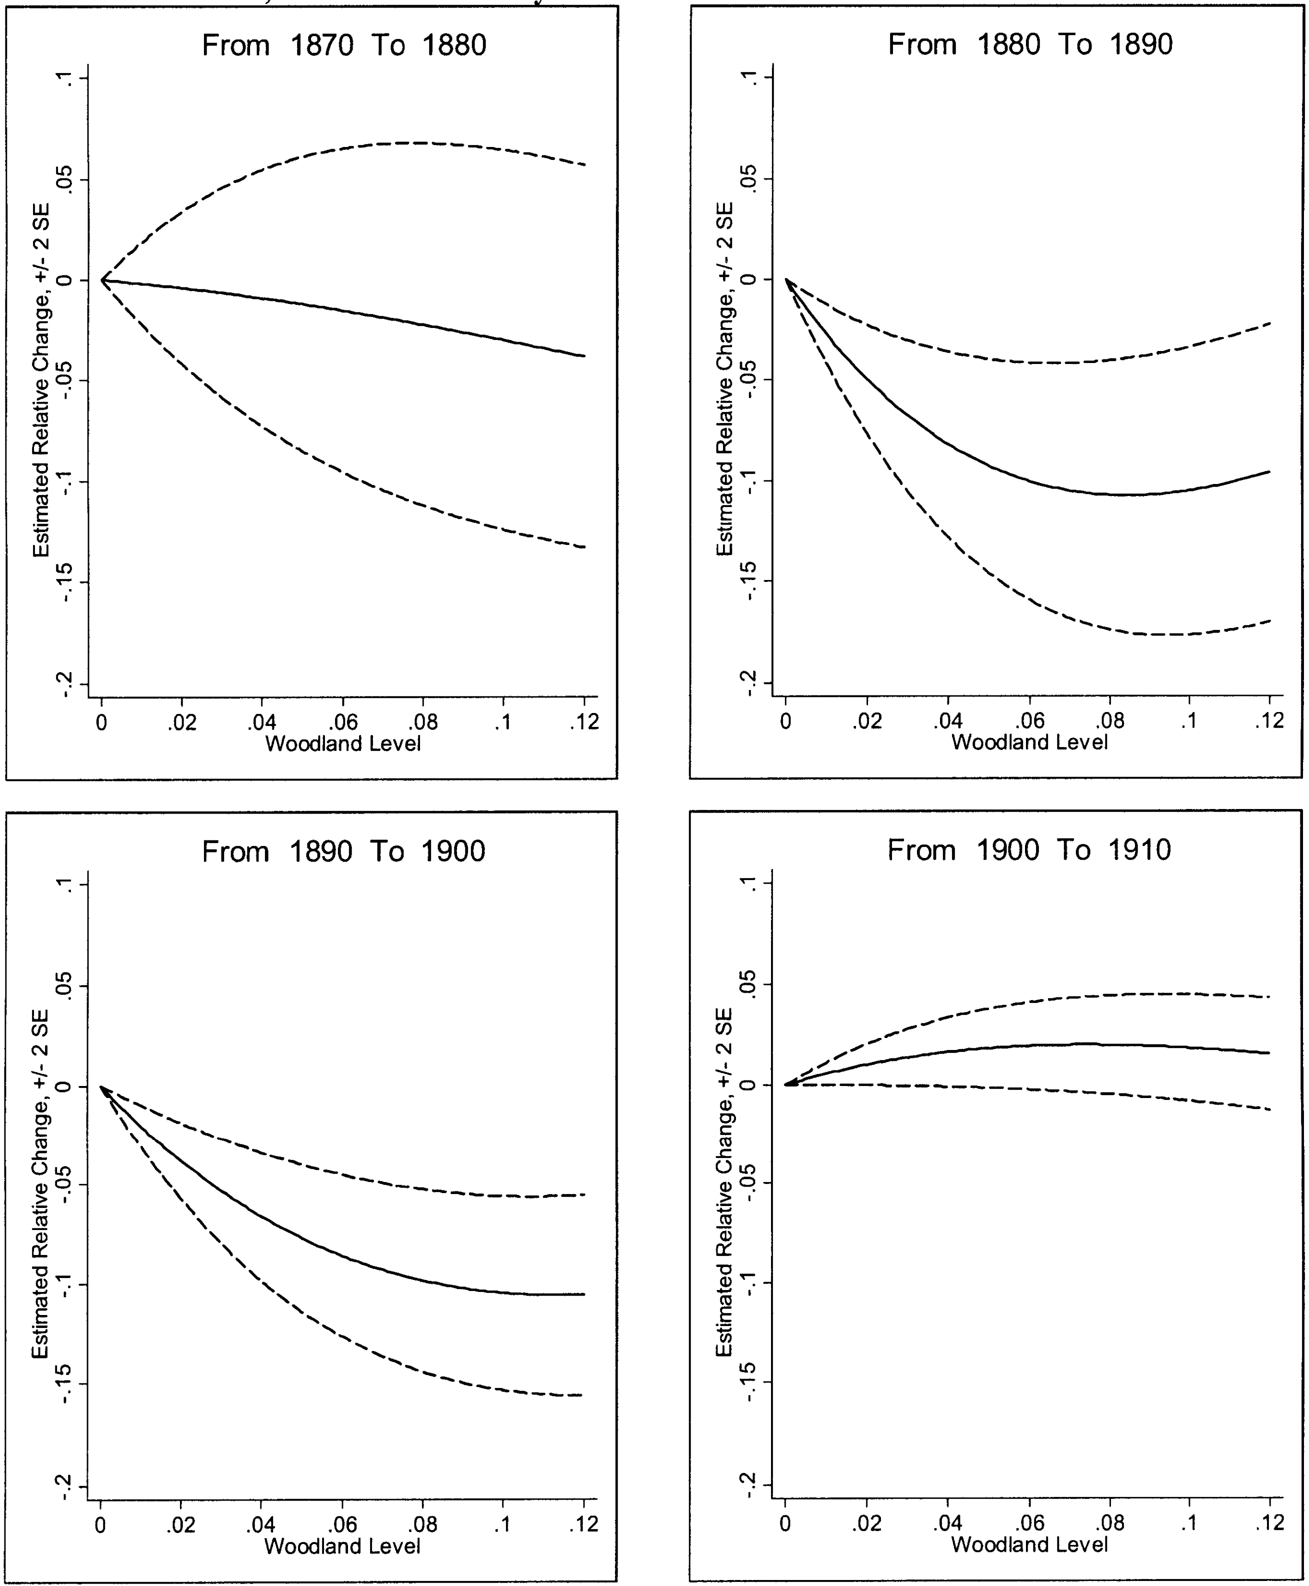

Notes: From estimating equation 8 for the acres of improved land per acre of farmland, these figures report the change in a county with each woodland level relative to the change for a county with 0% woodland, plus or minus 2 standard errors.

Figure 5. Estimated Cumulative Change in the Improvement Intensity of Farmland

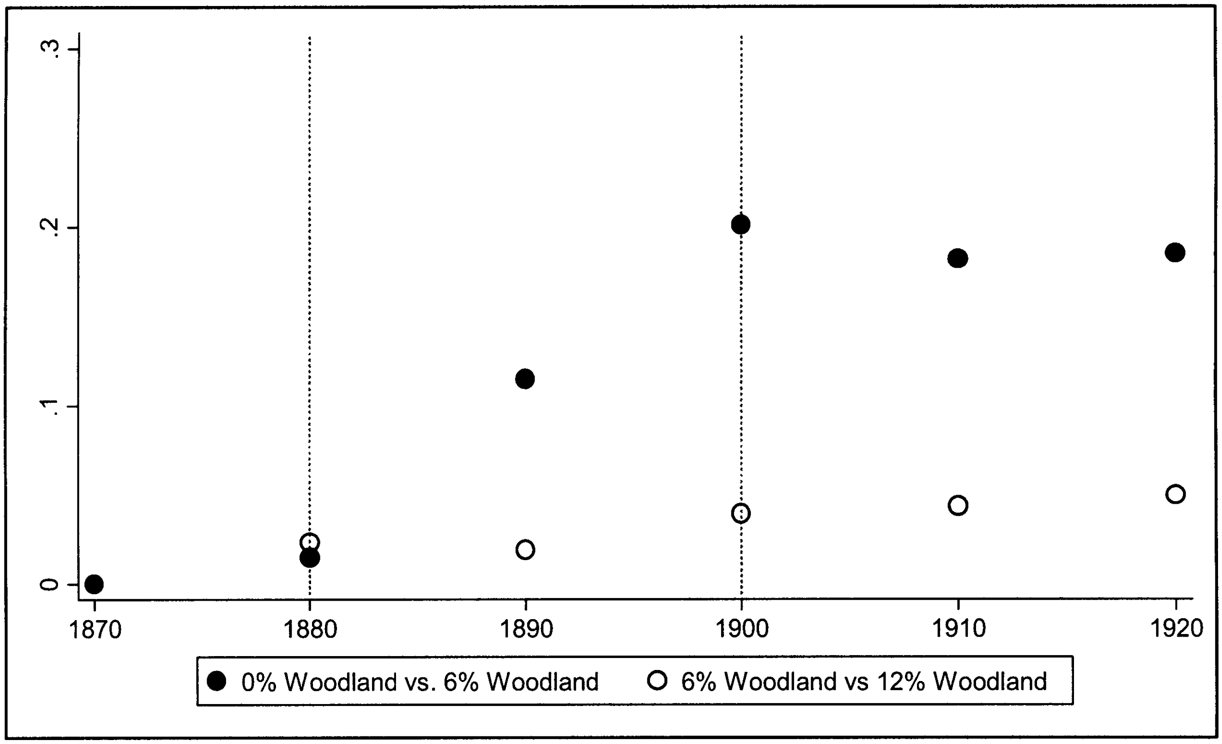

Notes: This figure displays the cumulative estimated changes in acres of improved land per acre of farmland, as reported in Table 3, columns 1 and 2. The two vertical dotted lines represent the approximate date of barbed wire's introduction to farmers (1880) and barbed wire's universal adoption by farmers (1900).

Figure 6. County Changes in the Improvement Intensity of Farmland, 1880 to 1900

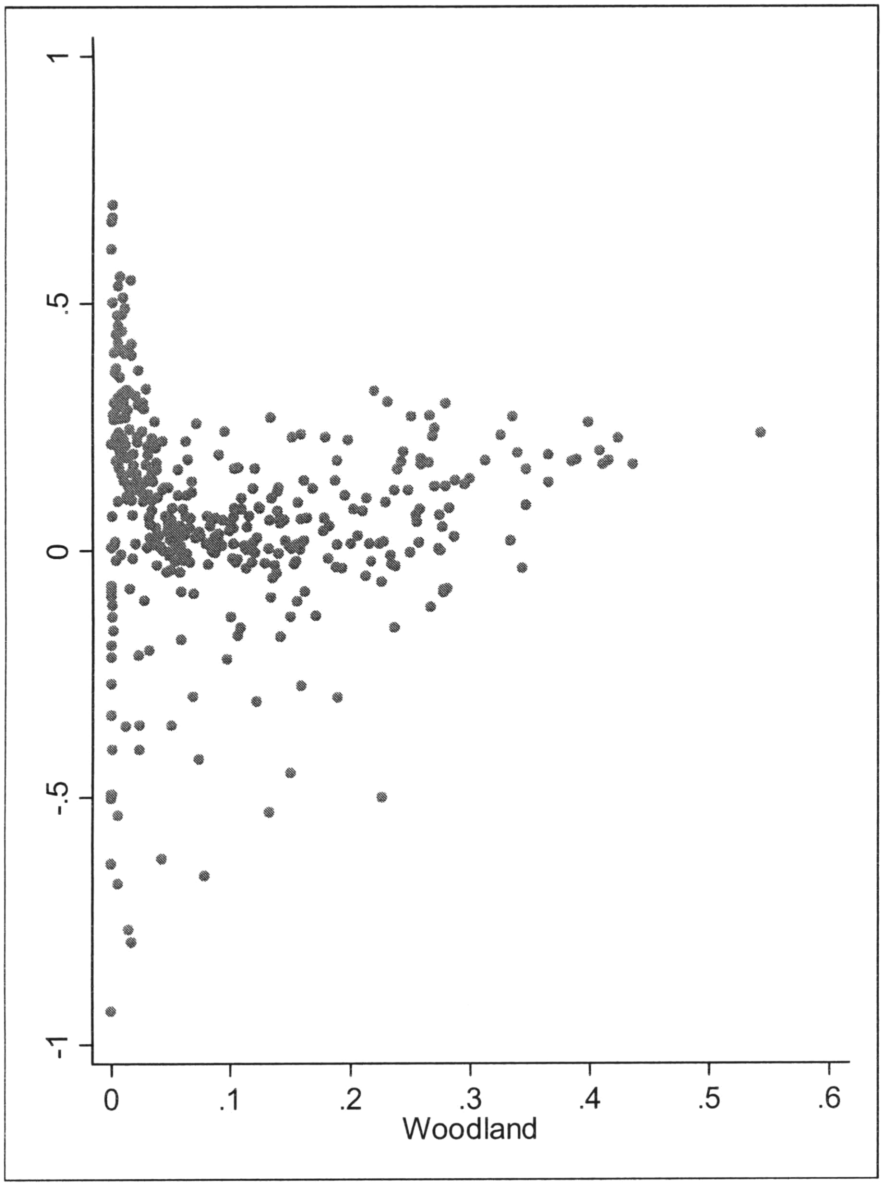

Notes: This figure shows the change from 1880 to 1900 in acres of improved land per acre of farmland, where each point represents a single county. This change is plotted against that county's 1880 woodland level.

Appendix Table 1. Estimated Changes in the Improvement Intensity of Farmland and Settlement, Controlling for Westward Development

	Distance West				Distance West and from St. Louis			
	Improved Land, per acre in farms		Land in Farms, per county acre		Improved Land, per acre in farms		Land in Farms, per county acre	
Woodland:	0% - 6% (2a)	6% - 12% (2b)	0% - 6% (2c)	6% - 12% (2d)	0% - 6% (3a)	6% - 12% (3b)	0% - 6% (3c)	6% - 12% (3d)
Decade:								
1870 - 1880	– 0.006 (0.044) [0.55]	0.016 (0.013)	– 0.050 (0.024) [2.02]	– 0.003 (0.009)	– 0.005 (0.046) [0.52]	0.016 (0.013)	– 0.042 (0.025) [1.82]	0.001 (0.009)
1880 - 1890	0.094** (0.034) [3.76]	– 0.007 (0.014)	0.068* (0.027) [1.95]	0.026** (0.010)	0.119** (0.031) [4.54]	0.006 (0.013)	0.083** (0.026) [2.33]	0.033** (0.009)
1890 - 1900	0.094** (0.026) [3.33]	0.023 (0.012)	0.079* (0.031) [1.56]	0.040** (0.011)	0.106** (0.026) [3.64]	0.029* (0.011)	0.066* (0.032) [1.27]	0.034** (0.011)
1900 - 1910	– 0.009 (0.016) [1.42]	0.008 (0.008)	– 0.009 (0.023) [0.01]	– 0.009 (0.010)	– 0.015 (0.015) [1.75]	0.005 (0.008)	0.007 (0.024) [0.34]	– 0.001 (0.009)
1910 - 1920	0.006 (0.013) [0.08]	0.007 (0.005)	0.033 (0.017) [1.95]	0.000 (0.008)	0.003 (0.014) [0.22]	0.006 (0.004)	0.023 (0.019) [1.53]	– 0.005 (0.007)
R^2	0.4458		0.5583		0.4589		0.5716	
Obs.	1885		1885		1885		1885	

Notes: This Table presents the full results from Table 4, columns 2 and 3.

Appendix Table 2. Estimated Changes in the Improvement Intensity of Farmland and Settlement, Controlling for Regional Changes

	Quadratic Time Trends by Region				Quadratic Time Trends by Subregion				Quadratic Time Trends by Soil Group			
	Improved Land, per acre in farms		Land in Farms, per county acre		Improved Land, per acre in farms		Land in Farms, per county acre		Improved Land, per acre in farms		Land in Farms, per county acre	
Woodland:	0% - 6% (4a)	6% - 12% (4b)	0% - 6% (4c)	6% - 12% (4d)	0% - 6% (5a)	6% - 12% (5b)	0% - 6% (5c)	6% - 12% (5d)	0% - 6% (6a)	6% - 12% (6b)	0% - 6% (6c)	6% - 12% (6d)
Decade:												
1870 - 1880	0.013	0.039**	0.006	0.034**	0.004	0.031**	– 0.032	0.016	0.047	0.033**	0.029	0.022*
	(0.040)	(0.013)	(0.024)	(0.010)	(0.035)	(0.012)	(0.026)	(0.010)	(0.035)	(0.011)	(0.025)	(0.010)
	[0.73]		[1.22]		[0.82]		[1.94]		[0.45]		[0.29]	
1880 - 1890	0.100**	0.006	0.106**	0.052**	0.087**	0.000	0.078**	0.039**	0.120**	0.003	0.121**	0.043**
	(0.029)	(0.011)	(0.025)	(0.009)	(0.031)	(0.012)	(0.024)	(0.009)	(0.031)	(0.012)	(0.025)	(0.010)
	[3.76]		[2.48]		[3.29]		[1.87]		[4.36]		[3.54]	
1890 - 1900	0.087**	0.025**	0.116**	0.060**	0.070**	0.019*	0.097**	0.053**	0.094**	0.024*	0.122**	0.055**
	(0.020)	(0.009)	(0.026)	(0.010)	(0.021)	(0.010)	(0.027)	(0.010)	(0.021)	(0.010)	(0.026)	(0.010)
	[3.22]		[2.48]		[2.61]		[1.92]		[3.50]		[2.91]	
1900 - 1910	– 0.016	0.005	– 0.023	– 0.012	– 0.038**	– 0.001	– 0.033	– 0.014	– 0.024*	0.006	– 0.026	– 0.013
	(0.011)	(0.006)	(0.024)	(0.010)	(0.011)	(0.006)	(0.027)	(0.010)	(0.011)	(0.006)	(0.023)	(0.009)
	[1.98]		[0.44]		[3.36]		[0.75]		[2.76]		[0.58]	
1910 - 1920	0.008	0.001	0.034*	– 0.004	– 0.019	– 0.004	0.033	0.000	– 0.014	0.005	0.022	0.000
	(0.011)	(0.005)	(0.017)	(0.006)	(0.014)	(0.005)	(0.021)	(0.007)	(0.013)	(0.005)	(0.016)	(0.007)
	[0.58]		[2.17]		[1.14]		[1.63]		[1.55]		[1.26]	
R^2	0.4773		0.5418		0.5095		0.5737		0.4864		0.5352	
Obs.	1885		1885		1885		1885		1885		1885	

Notes: This Table presents the full results from Table 4, columns 4 – 6.

Appendix Table 3. Estimated Changes in the Improvement Intensity of Farmland and Settlement, Controlling for Initial Outcome Differences

	Improved Land, per acre in farms		Land in Farms, per county acre	
Woodland:	0% - 6% (7a)	6% - 12% (7b)	0% - 6% (7c)	6% - 12% (7d)
Decade:				
1870 - 1880	– 0.098**	– 0.003	– 0.086**	– 0.032**
	(0.034)	(0.013)	(0.029)	(0.011)
	[3.44]		[2.15]	
1880 - 1890	0.065*	– 0.015	0.031	0.003
	(0.033)	(0.012)	(0.029)	(0.013)
	[2.90]		[1.26]	
1890 - 1900	0.076**	0.017	0.062	0.034*
	(0.021)	(0.009)	(0.039)	(0.015)
	[2.95]		[0.99]	
1900 - 1910	– 0.026*	0.002	– 0.045	– 0.026
	(0.012)	(0.006)	(0.037)	(0.015)
	[2.62]		[0.64]	
1910 - 1920	– 0.005	0.004	0.012	– 0.009
	(0.012)	(0.005)	(0.022)	(0.010)
	[0.71]		[0.98]	
R^2	0.5735		0.5606	
Obs.	1885		1885	

Notes: This Table presents the full results from Table 4, column 7.

Identifying Agglomeration Spillovers:

Evidence from Million Dollar Plants*

Michael Greenstone

Richard A. Hornbeck

Enrico Moretti

Abstract

The third essay quantifies agglomeration spillovers by comparing the growth of total factor productivity (TFP) among incumbent plants in "winning" counties that attracted a large manufacturing plant to "losing" counties that were the plant's second choice. Five years after the opening, incumbent plants' TFP is 12% higher in winning counties. This effect is larger for plants with similar labor and technology pools as the new plant. We find evidence of increased wages in winning counties, indicating that profits increase by less than productivity.

---

*We thank Daron Acemoglu, Jim Davis, Vernon Henderson, William Kerr, Jeffrey Kling, Jonathan Levin, Stuart Rosenthal, Christopher Rohlfs, Chad Syverson, and seminar participants at Berkeley, the Brookings Institution, MIT, NBER Summer Institute, San Francisco Federal Reserve, Stanford, and Syracuse for insightful comments. Elizabeth Greenwood provided valuable research assistance. The research in this paper was conducted while the authors were Special Sworn Status researchers of the U.S. Census Bureau at the Boston Census Research Data Center (BRDC). Support of the Census Research Data Center network from NSF grant no. 0427889 is gratefully acknowledged. Research results and conclusions expressed are those of the authors and do not necessarily reflect the views of the Census Bureau. This paper has been screened to insure that no confidential data are revealed.

In most countries, economic activity is spatially concentrated. While some of this concentration is explained by the presence of natural advantages that constrain specific productions to specific locations, Ellison and Glaeser (1999) and others argue that natural advantages alone cannot account for the observed degree of agglomeration. Spatial concentration is particularly remarkable for industries that produce nationally traded goods, because the areas where economic activity is concentrated are typically characterized by high costs of labor and land. Since at least Marshall (1890), economists have speculated that this concentration of economic activity may be explained by cost or productivity advantages enjoyed by firms when they locate near other firms. The list of potential sources of these agglomerations advantages includes: cheaper and faster supply of intermediate goods and services; proximity to workers or consumers; better quality of the worker-firm match in thicker labor markets; lower risk of unemployment for workers and lower risk of unfilled vacancies for firms following idiosyncratic shocks; and knowledge spillovers.[1]

The possibility of documenting that productivity advantages through agglomeration are real is tantalizing, because it could provide insights into a series of important questions. Why are firms that produce nationally traded goods willing to locate in cities like New York, San Francisco or London, characterized by extraordinary production costs? In general, why do cities exist and what explains their historical development? Why do income differences persist across regions and countries?

Beside its obvious interest for urban and growth economists, the existence of agglomeration spillovers has tremendous practical relevance. Increasingly, local governments compete by offering substantial subsidies to industrial plants to locate within their jurisdictions. The main economic rationale for these incentives depends on whether the attraction of new businesses generates some form of agglomeration externalities. In the absence of positive externalities, it is difficult to justify the use of taxpayer money for subsidies based on efficiency grounds. The optimal magnitude of these incentives depends on the magnitude of agglomerations

---

[1] The literature on this topic is enormous, and can not be fully summarized here. Examples include, but are not limited to, Lucas (1988), Krugman (1991a, 1991b), Henderson (2001a, 2001b, 2003), Davis and Henderson (2004), Davis and Weinstein (2002), Henderson and Black (1999), Rosenthal and Stange (2001, 2004), Duranton and Puga (2004), Audretsch and Feldman, (1996, 2004), Moretti (2004a, 2004b, 2004c), Dumais, Ellison and Glaeser (2002), Glaeser (1999), Ottaviano and Thisse (2004).

spillovers, if they exist.[2]

Despite their enormous theoretical and practical relevance, the existence and exact magnitude of agglomeration spillovers are considered open questions by many. To date, there are two primary approaches for testing for spillovers. The first tests for an unequal distribution of firms across the country. These "dartboard" style tests reveal that firms are spread unevenly across the country and that coagglomeration rates are higher between industries that are economically similar (Ellison, Glaeser and Kerr, 2007). This approach is based on equilibrium location decisions and does not provide a direct measure of spillovers.

The second approach uses micro data to assess whether firms' total factor productivity (TFP) is higher when similar firms are located nearby. A notable example is Henderson (2003), which estimates plant level production functions for machinery and high-tech industries as a function of the scale of other plants in the same and different industries.[3] The challenge for both approaches is that firms base their location decisions on where their profits will be highest, and this could be due to spillovers, natural advantage, or other cost shifters. A causal estimate of the magnitude of spillovers requires a solution to this problem of identification.

This paper tests for and quantifies agglomeration spillovers by estimating how the productivity of incumbent manufacturing plants changes when a new, large plant opens in their county. We estimate augmented Cobb-Douglass production functions that allow TFP to depend on the presence of the new plant, using plant-level data from the Annual Survey of Manufacturers. Because the location decision is made to maximize profits, the chosen county is likely to differ substantially from an average or randomly chosen county both at the time of opening and in future periods. Valid estimates of the plant opening's spillover effect require the identification of a county that is identical to the county where the plant decided to locate in the determinants of incumbent plants' TFP. These determinants are likely to include factors that affect the new plant's TFP.

This paper's solution is to rely on the reported location rankings of profit-maximizing firms to identify a valid counterfactual for what would have happened to incumbent plants' TFP in winning counties in the absence of the plant opening. These rankings come from the

---

[2] We discuss in more detail the policy implications of local subsidies in Greenstone and Moretti (2004). See also Card, Hallock, and Moretti (2007) and Glaeser (2001).

[3] Moretti (2004b) takes a similar approach to estimate agglomeration externalities generated by human capital spillovers.

corporate real estate journal *Site Selection*, which includes a regular feature titled "Million Dollar Plants" that describes how a large plant decided where to locate. When firms are considering where to open a large plant, they typically begin by considering dozens of possible locations. They subsequently narrow the list to roughly 10 sites, among which 2 or 3 finalists are selected. The "Million Dollar Plants" articles report the county that the plant ultimately chose (i.e., the 'winner'), as well as the one or two runner-up counties (i.e., the 'losers'). The losers are counties that have survived a long selection process, but narrowly lost the competition.

The identifying assumption is that the incumbent plants in the losing county form a valid counterfactual for the incumbents in the winning counties, after conditioning on differences in pre-existing trends, plant fixed effects, industry-by-year fixed effects, and other control variables. Compared to the rest of the country, winning counties have higher rates of growth in income, population, and labor force participation. But compared to losing counties in the years before the opening of the new plant, winning counties have similar trends in most economic variables. This finding is consistent with both our presumption that the average county is not a credible counterfactual and our identifying assumption that the losers form a valid counterfactual for the winners.

We first measure the effect of the new Million Dollar Plant (MDP) on total factor productivity (TFP) of all incumbent manufacturing plants in winning counties. In the 7 years before the MDP opened, we find statistically equivalent trends in TFP for incumbent plants in winning and losing counties, which supports the validity of the identifying assumption. After the MDP opened, incumbent plants in winning counties experienced a sharp relative increase in TFP. Five years later, the MDP opening is associated with a 12% relative increase in incumbent plants' TFP.[4] This effect is statistically significant and economically substantial; on average incumbent plants' output in winning counties is $430 million higher 5 years later (relative to incumbents in losing counties), holding constant inputs. We interpret this finding as evidence of meaningful productivity spillovers generated by increased agglomeration.

Having found evidence in favor of the existence of agglomeration spillovers, we then try to shed some light on the possible mechanisms. We follow Moretti (2004b) and Ellison, Glaeser, and Kerr (2007) and investigate whether the magnitude of the spillovers depends on economic

---

[4] Notably, naïve estimates that control for observables but do not use the MDP research design find negative productivity effects.

linkages between the incumbent plant and the MDP. Specifically, we test whether incumbents that are geographically and economically linked to the MDP experience larger spillovers, relative to incumbents that are geographically close but economically distant from the MDP. We use several measures of economic links including input and output flows, measures of the degree of sharing of labor pools, and measures of technological linkages.[5]

We find that spillovers are larger for incumbent plants in industries that share worker flows with the MDP industry. A one standard deviation increase in our measure of worker transitions is associated with a 7 percentage point increase in the magnitude of the spillover. Similarly, the measures of technological linkages indicate statistically meaningful increases in the spillover effect. Surprisingly, we find little support for the importance of input and output flows in determining the magnitude of the spillover. Overall, this evidence provides some support for the notion that spillovers occur between firms that share workers and use similar technologies.

To guide the analysis and interpret the results, we set out a straightforward Roback (1982) style model that incorporates spillovers between producers and derives an equilibrium allocation of firms and workers across locations. In the model, the entry of a new firm produces spillovers. This leads new firms who are interested in gaining access to the spillover to enter. This process of entry leads to competition for inputs so that incumbent firms face higher prices for labor, land, and other local inputs. In the model, firms produce nationally traded goods, so they cannot raise the price of their output in response to the higher input prices. Thus, the long run equilibrium is obtained when the value of the increase in output due to spillovers is equal to the increased costs of production due to the higher input prices

The data support two predictions derived from this model. First, we find positive net entry in winning counties, which the model predicts will occur if there are sufficiently large positive spillovers. Second, we find increases in quality-adjusted wages following MDP openings. These higher wages are consistent with the documented increase in economic activity in the winning counties (which presumably is a response to the spillovers). Furthermore, the higher wages support the model's prediction that the productivity gains from agglomeration do not necessarily translate into higher profits for incumbents in the long run.

---

[5] We are deeply indebted to Glenn Ellison, Edward Glaeser, and William Kerr for providing their data for five of these measures of economic distance.

The remainder of the paper is organized as follows. Section I presents a simple model. Section II discusses the identification strategy. Section III presents the data sources. Section IV presents the econometric model. Sections V and VI presents the empirical findings and interpret them. Section VII concludes.

I. Theories of Agglomeration and Theoretical Framework

We are interested in identifying how the opening of a new plant in a county affects the productivity, profits, and input use of existing plants in the same county. We begin by reviewing the theories of agglomeration. We then present a simple theoretical framework that guides the subsequent empirical exercise and aids in interpreting the results.

A. Theories of Agglomeration

Economic activity is geographically concentrated (Ellison and Glaeser, 1997). What are the forces that can explain such agglomeration of economic activity? Here we summarize five possible reasons for agglomeration, and briefly discuss what each of them implies for the relationship between productivity and the density of economic activity.

(1) First, it is possible that firms (and workers) are attracted to areas with a high concentration of other firms (and other workers) by the <u>size of the labor market</u>. There are at least two different reasons why larger labor markets may be attractive. First, a thick labor market is beneficial in the presence of search frictions, if jobs and workers are heterogeneous. In the presence of frictions, a worker-firm match will be on average more productive in areas where there are many firms offering jobs and many workers looking for jobs.[6]

Alternatively, it is possible that large labor markets are more desirable because they provide insurance against idiosyncratic shocks, either on the firm side or on the worker side (Krugman 1991a). If moving is costly for workers and firms are subject to idiosyncratic and unpredictable demand shocks that lead to lay-offs, workers will prefer to be in areas with thick labor markets to reduce the probability of being unemployed. Similarly, if finding new workers is costly, firms will prefer to be in areas with thick labor markets to reduce the probability of

---

[6] For a related point in a different context, see Petrongolo and Pissarides (2005).

having unfilled vacancies.[7]

These two hypotheses have different implications for the relationship between concentration of economic activity and productivity. If the size of the labor market leads only to better worker-firm matches, we should see that firms located in denser areas are more productive than otherwise identical firms located in less dense areas. The exact form of this productivity gain depends on the shape of the production function.[8]

On the other hand, if the only effect of thickness in labor market is a lower risk of unemployment for workers and a lower risk of unfilled vacancies for firms, there should not be differences in productivity between dense and less dense areas. While productivity would not vary, wages would vary across areas depending on the thickness of the labor market, although the exact effect of density on wages is *a priori* ambiguous.[9] This change in relative factor prices will change the relative amount of labor and capital used. Unlike the case of improved matching described above, the production function does not change: for the same set of labor and capital inputs, the output of firms in denser areas should be similar to the output of firms in less dense areas.

(2) A second reason why the concentration of economic activity may be beneficial has to do with transportation costs (Krugman 1991a and 1991b, Glaeser and Kohlhase, 2003). Because in this paper we focus on firms that produce nationally traded goods, transportation costs of finished products are unlikely to be the relevant cost in this paper's setting. Only a small fraction of buyers of the final product is likely to be located in the same area as our manufacturing plants. The relevant costs are the transportation costs of suppliers of local services and local intermediate goods. Firms located in denser areas are likely to enjoy cheaper and faster delivery

---

[7] A third alternative hypothesis has to do with spillovers that arise because of endogenous capital accumulation. For example, in Acemoglu (1996), plants have more capital and better technology in areas where the number of skilled workers is larger. If firms and workers find each other via random matching and breaking the match is costly, externalities will arise naturally even without learning or technological externalities. The intuition is simple. The privately optimal amount of skills depends on the amount of physical capital a worker expects to use. The privately optimal amount of physical capital depends on the number of skilled workers. If the number of skilled workers in a city increases, firms in that city, expecting to employ these workers, will invest more. Because search is costly, some of the workers end up working with more physical capital and earn more than similar workers in other cities.

[8] For example, it is possible that the productivities of both capital and labor benefit from the improved match in denser areas. It is also possible that the improved match caused by a larger labor market benefits only labor productivity. This has different implications for the relative use of labor and capital, but total factor productivity will be higher regardless.

[9] Its sign depends on the relative magnitude of the compensating differential that workers are willing to pay for lower risk of unemployment (generated by an increase in labor supply in denser areas) and the cost savings that firms experience due to lower risk of unfilled vacancies (generated by an increase in labor demand in denser areas).

of local services and local intermediate goods. For example, a high-tech firm that needs a specialized technician to fix a machine is likely to get service more quickly and at lower cost if it is located in Silicon Valley than in the Nevada desert.

This type of agglomeration spillover does not imply that the production function varies as a function of density of economic activity: for the same set of labor and capital inputs, the output of firms in denser areas should be similar to the output of firms in less dense areas. However, production costs should be lower in denser areas.

(3) A third reason why the concentration of economic activity may be beneficial has to do with knowledge spillovers. There are at least two different versions of this hypothesis. First, economists and urban planners have long speculated that the sharing of knowledge and skills through formal and informal interaction may generate positive production externalities across workers.[10] Empirical evidence indicates that this type of spillover may be important in some high-tech industries. For example, patent citations are more likely to come from the same state or metropolitan area as the originating patent (Jaffe et al. 1993). Saxenian (1994) argues that geographical proximity of high-tech firms in Silicon Valley is associated with a more efficient flow of new ideas and ultimately causes faster innovation.[11] Second, it is also possible that proximity results in sharing of information on new technologies and therefore leads to faster technology adoption. This type of social learning phenomenon applied to technology adoption was first proposed by Griliches (1958).

If density of economic activity results in intellectual externalities, the implication of this type of agglomeration model is higher productivity. In particular, we should see that firms located in denser areas are more productive than otherwise identical firms located in less dense areas. As with the search model, this higher productivity could benefit both labor and capital, or only one of the two factors, depending on the form of the production function. On the other hand, if density of economic activity only results in faster technology adoption and the price of new technologies reflects their higher productivity, there should be no relationship between productivity and density, after properly controlling for quality of capital.

(4) It is possible that firms concentrate spatially not because of any technological

---

[10] See, for example, Marshall (1890), Lucas (1988), Jovanovic and Rob (1989), Grossman and Helpman (1991), Saxenian (1994), Glaeser (1999), and Moretti (2004a, 2004b and 2004c).

[11] The entry decisions of new biotechnology firms in a city depend on the stock of outstanding scientists there, as measured by the number of relevant academic publications (Zucker et al., 1998). Moretti (2004b) finds stronger human capital spillovers between pairs of firms in the same city that are economically or technologically closer.

spillover, but because local amenities valued by workers are concentrated. For example, skilled workers may prefer certain amenities more than unskilled workers. This would lead firms that employ relatively more skilled workers to concentrate in locations where these amenities are available. In this case, we should not see any difference in productivity between dense areas and less dense areas, although we should see differences in wages that reflect the compensating differential.

(5) Finally, spatial concentration of some industries may be explained by the presence of natural advantages. For example, the oil industry is concentrated in a limited number of states because those states have the most accessible oil fields. Similarly, the wine industry is concentrated in California because that is where good weather and suitable land are to be found. For some manufacturing productions, the presence of a harbor may be important. The presence of natural advantages has the implication that firms located in areas with high concentration of similar firms are more productive, but of course this correlation has nothing to do with agglomeration spillovers. Since most natural advantages are fixed over time, this explanation is not particularly relevant for our empirical estimates, which exploit variation over time in agglomeration.

B. A Simple Model

We begin by considering the case where incumbent firms are homogenous in size and technology. Later we consider what happens when incumbent firms are heterogeneous. Throughout the paper, we focus on the case of factor-neutral spillovers.

(a) Homogeneous Incumbents. We assume that all incumbent firms use a production technology that uses labor, capital, and land to produce a nationally traded good whose price is fixed and is normalized to 1. Incumbent firms choose their amount of labor, L, capital, K, and land, T, to maximize the following expression:

$$\text{Max}_{L,K,T} \{ f[A, L, K, T] - w L - r K - q T \}$$

where w, r and q are input prices and A is a productivity shifter (TFP). Specifically, A includes all factors that affect the productivity of labor, capital, and land equally, such as technology and agglomeration spillovers, if they exist. In particular, to explicitly allow for agglomeration effects, we allow A to depend on the density of economic activity in an area:

$$A = A(N) \tag{1}$$

where N is the number of firms that are active in a county, and all counties have equal size. We define factor-neutral agglomeration spillovers as the case where A increases in N:

$$\delta A / \delta N > 0$$

If instead $(\delta A / \delta N) = 0$, we say that there are no factor-neutral agglomeration spillovers.

Let $L^*(w,r,q)$ be the optimal level of labor inputs, given the prevailing wage, cost of capital, and cost of industrial land. Similarly, let $K^*(w,r,q)$ and $T^*(w,r,q)$ be the optimal level of capital and land, respectively. In equilibrium, L^*, K^*, and T^* are set so that the marginal product of each of the three factors is equal to its price.

We assume that capital is internationally traded, so its price does not depend on local demand or supply conditions. However, we allow for the price of labor and land to depend on local economic conditions. In particular, we allow the supply of labor and land to be less than infinitely elastic at the county level.

We attribute the upward sloping labor supply curve to the existence of moving costs. Like in the standard Roback (1982) model, we assume that workers' indirect utility depends on wages and cost of housing, and that in equilibrium workers are indifferent across locations. Workers are mobile across locations, but unlike the standard Roback (1982) model we allow for moving costs. For simplicity, we ignore labor supply decisions within a given location and assume that all residents provide a fixed amount of labor.

To illustrate this, consider that there are *m* workers in county *c* before the opening of the new plant. In particular, *m* is such that, given the distribution of wages and the housing costs across localities, the marginal worker in another county is indifferent between moving to county *c* and staying. When a new plant opens in county *c,* wages there start rising, and some workers find it optimal to move to *c*. The number of workers who move, and therefore the slope of the labor supply function, depend on the shape of the mobility cost function. Let w(N) be the inverse of the reduced-form labor supply function that links the number of firms, N, active in a county to the local nominal wage level, w.

Similarly, we allow the supply of industrial land to be less than infinitely elastic at the county level. For example, it is possible that the supply of land is fixed because of geography or land-use regulations. Alternatively, it may not be completely fixed, but it is possible that the best industrial land has already been developed, so that the marginal land is of decreasing quality or more expensive to develop. Irrespective of the reason, we call q(N) the (inverse of the) reduced form land supply function that links the number of firms, N, to the price of land, q. We can therefore write the equilibrium level of profits, Π^*, as

$$\Pi^* = f[\, A(N), L^*(w(N), r, q(N)), K^*(w(N), r, q(N)), T^*(w(N), r, q(N)) \,] \; - \\ - \, w(N)\, L^*(w(N), r, q(N)) - r\, K^*(w(N), r, q(N)) - q(N)\, T^*(w(N), r, q(N))$$

where we now make explicit the fact that TFP, wages, and land prices depend on the number of firms active in a county.

Consider the total derivative of incumbents' profits with respect to a change in the number of firms:

$$(2) \qquad d\Pi^*/dN = (\delta f/\delta A \;\; \delta A/\delta N) \\ + \delta w/\delta N \,\{\, [\delta L^*/\delta w\, (\delta f/\delta L - w) - L^*] + [\delta K^*/\delta w\, (\delta f/\delta K - r)] + [\delta T^*/\delta w\, (\delta f/\delta T - q)] \} \\ + \delta q/\delta N \,\{\, [\delta L^*/\delta q\, (\delta f/\delta L - w)] + [\delta K^*/\delta q\, (\delta f/\delta K - r)] + [\delta T^*/\delta q\, (\delta f/\delta T - q) - T^*] \,\}$$

If all firms are price takers and all factors are paid their marginal product, equation (2) simplifies considerably and can be written as:

$$(3) \qquad d\Pi^*/dN = (\delta f/\delta A \;\; \delta A/\delta N) \; - [\, \delta w/\delta N \;\; L^* + \delta q/\delta N \;\; T^*].$$

Equation (3) makes clear that the effect of an increase in N is the sum of two opposite effects. First, if there are positive spillovers, the productivity of all factors increases. In equation (3), this effect on TFP is represented by the first term, $(\delta f/\delta A \;\; \delta A/\delta N)$. This effect is unambiguously positive, because it allows an incumbent firm to produce more output using the same amount of inputs. Formally, $\delta f/\delta A > 0$ by assumption and, if there are positive spillovers, $\delta A/\delta N > 0$.

The second term, $- [\, \delta w/\delta N \;\; L^* + \delta q/\delta N \;\; T^*]$, represents the negative effect from increases in the cost of production, specifically the prices of labor and land. Formally, this term is negative because we have assumed that $\delta w/\delta N > 0$ and $\delta q/\delta N > 0$. (The magnitudes depend on the elasticity of the supply of labor and land.) Intuitively, an increase in N is an increase in the level of economic activity in the county and therefore an increase in the local demand for labor and land.

Unlike the beneficial effect of agglomeration spillovers, the increase in factor prices is costly for incumbent firms, because they now have to compete for locally scare resources with the new entrant. The increase in wages and land prices has two effects on incumbents. First, for a given level of input utilizations, it mechanically raises production costs. Second, it leads the firm to re-optimize and to change its use of the different production inputs. In particular, given that the price of capital is not affected by an increase in N, the firm is likely to end up using more capital than before:

$$\delta K^{*}/\delta N => 0.$$

By contrast, the effect on the use of labor and land is ambiguous. On the one hand, the productivity of all factors increases. On the other hand, the price of labor and land might increase. The net effect depends on the magnitude of the factor price increases, as well as on the exact shape of the production function (i.e., the strength of technological complementarities between labor, capital, and land).

It is instructive to apply these derivations to the case of a MDP opening that causes positive spillovers. We initially consider the case where for incumbent firms $d\Pi^{*}/dN \leq 0$. This would occur when the agglomeration spillover is smaller than the increase in production costs. In this case, the MDP's opening would not lead to any entry and could cause some existing firms to exit.[12]

The alternative case is that $d\Pi^{*}/dN > 0$, which occurs when the magnitude of the spillover due to the MDP's opening exceeds the increase in factor prices due to the MDP's demand for local inputs. In the short run, profits will be positive for new entrants. These positive profits will disappear over time as the price of local factors, like land and possibly labor, is bid up.

In the long run, there is an equilibrium such that firms and workers are indifferent between the county where the new plant has opened and other locales. Since the amount of land is fixed, the higher levels of productivity are likely to be capitalized into land prices. It is also likely that wages will increase. This may occur due to moving costs as noted above.[13] These adjustments make workers indifferent between the county with the new plant and other counties. Similarly, the changes in factor prices mean that firms earn the same profits in the county with the new plant (even in the presence of the spillovers) and in other locations. From a practical perspective, it is impossible to know when the short run ends and the long run begins.

There are two empirical predictions that apply when there are positive spillovers. First, if

---

[12] Similarly, if the spillovers are zero or negative, one might expect exit of incumbent firms and a reduction in local economic activity.

[13] Even with zero moving costs and an infinitely elastic supply of labor, wages will increase if there are land price increases as workers will demand higher wages as compensation for the higher land rents for their homes (Roback 1982).

the magnitude of the spillovers is large enough, new firms will enter the MDP's county to gain access to the spillover. This prediction of increased economic activity holds at any point after potential new entrants have had sufficient time to respond. The second prediction is that the prices of locally traded inputs will rise as the MDP and the new entrants bid for these inputs.[14]

(b) Heterogeneous Incumbents. What happens if the population of incumbent firms is non-homogeneous? Consider the case where there are two types of firms: high-tech and low-tech. Assume that for technological reasons, the type of workers employed by high-tech firms, L_H, differs to some extent from the type of workers employed by low-tech firms, L_L, although there is some overlap. Assume that the new entrant is a high-tech firm. Equations (4) and (5) characterize the effect of the new high-tech firm on high-tech and low-tech incumbents:

$$(4) \qquad d\Pi_H^*/dN_H = (\delta f_H/\delta A_H \;\; \delta A_H/\delta N_H) - [\,\delta w_H/\delta N_H \;\; L^*_H + \delta q/\delta N_H \;\; T^*]$$

$$(5) \qquad d\Pi_L^*/dN_H = (\delta f_L/\delta A_L \;\; \delta A_L/\delta N_H) - [\,\delta w_L/\delta N_H \;\; L^*_L + \delta q/\delta N_H \;\; T^*]$$

It is plausible to expect that the beneficial effect of agglomeration spillovers generated by a new high-tech entrant is larger for high-tech firms than for low-tech firms:

$$(5') \qquad (\delta f_H/\delta A_H \;\; \delta A_H/\delta N_H) > (\delta f_L/\delta A_L \;\; \delta A_L/\delta N_H)$$

At the same time, one might expect that the increase in labor costs is also higher for the high-tech incumbents, given that they are now competing for workers with an additional high-tech firm:

$$\delta w_H/\delta N_H > \delta w_L/\delta N_H$$

The effect on land prices should be similar for both firm types, since the assumption of a single land market seems reasonable.

There are two takeaways here. First, it may be reasonable to expect larger spillovers to firms

---

[14] This paper focuses on the case where the productivity benefits of the agglomeration spillovers are distributed equally across all factors. What happens when agglomeration spillovers are factor biased? Assume, for example, that agglomeration spillovers raise the productivity of labor, but not the productivity of capital. Like before, the technology is $f[A, L, K, T]$, but now L represents units of effective labor. In particular, $L = \theta H$, where H is the number of physical workers and θ is a productivity shifter. We define factor-biased agglomeration spillover as the case where the productivity shifter θ depends positively on the density of the economic activity in the county $\theta = \theta(N)$ and $\delta\theta/\delta N > 0$. If $\delta A/\delta N = 0$ and factors are paid their marginal product, then the effect of an increase in the density of the economic activity in a county on incumbent firms simplifies to $d\Pi^*/dN = (\delta f/\delta H \;\; \delta\theta/\delta N)\, H - [\,\delta w/\delta N \;\; H^* + \delta q/\delta N \;\; T^*]$. The effect on profits can be decomposed in two parts. The first term represents the increased productivity of labor. It is the product of the sensitivity of output to labor ($\delta f/\delta H > 0$), times the magnitude of the agglomeration spillover ($\delta\theta/\delta N > 0$ by definition), times the number of workers. The second term is the same as in equation (3), and represents the increase in the costs of locally supplied inputs. The increase in N changes the optimal use of the production inputs. Labor is now more productive, and its equilibrium use increases: $\delta L^*/\delta N <= 0$. Land is equally productive but its price increases. Its equilibrium use declines: $\delta T^*/\delta N <= 0$. Neither the price nor the productivity of capital is affected by an increase in N. Its equilibrium use depends on technology. Specifically, it depends on the elasticity of substitution between labor and capital.

that are economically "close" to the new plant. Second the relative impact of the new plant on profits is unclear, because the economically "closer" plants are likely to have bigger spillovers and larger increases in production costs.

C. Empirical Predictions

The simple theoretical framework above generates four predictions that we bring to the data. Specifically if there are positive spillovers, then:

1. the opening of a new plant will increase the TFP of incumbent plants.
2. the increase in TFP may be larger for firms that are economically "closer" to the new plant.
3. the density of economic activity in the county will increase as firms move in to gain access to the positive spillovers (if the spillovers are large enough).
4. the price of factors of production that are traded locally will increase. We test for changes in the price of quality-adjusted labor, which is arguably the most important locally supplied factor of production for manufacturing establishments.

II. Plant Location Decisions and Research Design

In testing the four empirical predictions outlined above, the main econometric challenge is the fact that firms do not choose their location randomly. Firms are profit maximizers and choose to locate where their expectation of the present discounted value of the stream of future profits is greatest. This net present value varies tremendously across locations, depending on many factors, including transportation infrastructure, the availability of workers with particular skills, subsidies, etc. These factors are frequently unobserved. Further, they are likely to be correlated with the TFP of existing plants.

Therefore, a naïve comparison of the TFP of incumbents in counties that experience a plant opening with the TFP of incumbents in counties that do not experience a plant opening is likely to yield biased estimates of productivity spillovers. Credible estimates of the impact of a plant opening on TFP of incumbent plants require the identification of a location that is similar to the location where the plant decided to locate in the determinants of incumbent plants' TFP.

This section provides a case study for how BMW picked the location for one of its plants.[15] The intent is to demonstrate the empirical difficulties that arise when estimating the effect of plant openings on the TFP of incumbent plants. Further, it illustrates informally how our research design may circumvent these difficulties.

After overseeing a worldwide competition and considering 250 potential sites for its new plant, BMW announced in 1991 that they had narrowed the list of potential candidates to 20 counties. Six months later, BMW announced that the two finalists in the competition were Greenville-Spartanburg, South Carolina, and Omaha, Nebraska. In 1992, BMW announced that they would site the plant in Greenville-Spartanburg and that they would receive a package of incentives worth approximately $115 million funded by the state and local governments.

Why did BMW choose Greenville-Spartanburg? Two factors were important in this decision. The first was BMW's expected future costs of production in Greenville-Spartanburg, which are presumably a function of the county's expected supply of inputs and BMW's production technology. According to BMW, the characteristics that made Greenville-Spartanburg more attractive than the other 250 sites initially considered were: low union density, a supply of qualified workers, the numerous global firms, including 58 German companies, in the area; the high quality transportation infrastructure, including air, rail, highway, and port access; and access to key local services.

For our purposes, the important point to note here is that these county characteristics are a potential source of unobserved heterogeneity. While these characteristics are well documented in the BMW case, they are generally unknown and unobserved. If these characteristics also affect the growth of TFP of existing plants, a standard regression that compares Greenville-Spartanburg with the other 3,000 United States counties will yield biased estimates of the effect of the plant opening. A standard regression will overestimate the effect of plant openings on outcomes if, for example, counties that have more attractive characteristics (e.g., improving transportation infrastructure) tend to have faster TFP growth. Conversely, a standard regression would underestimate the effect if, for example, incumbent plants' declining TFP encouraged new entrants (e.g., cheaper availability of local inputs).

A second important factor in BMW's decision was the value of the subsidy it received. Presumably Greenville-Spartanburg was willing to provide BMW with $115 million in subsidies

---

[15] This plant is in Greenstone and Moretti's (2004) set of 82 MDP plants. Due to Census confidentiality restrictions, we cannot report whether this plant is part of this paper's analysis.

because it expected economic benefits from BMW presence. According to local officials, the facility's ex-ante expected five-year economic impact on the region was $2 billion. As a part of this $2 billion, the plant was expected to create 2,000 jobs *directly* and another 2,000 jobs *indirectly*. In principle, these 2,000 additional jobs could reflect the entry of new plants or the expansion of existing plants caused by agglomeration economies. (The empirical section tests whether this is indeed the case on average). Thus, the subsidy is likely to be a function of the expected gains from agglomeration for the county.[16]

This possibility is relevant for this paper's identification strategy, because the magnitude of the spillover from a particular plant depends on the level and growth of a county's industrial structure, labor force, and a series of other unobserved variables. For this reason, the factors that determine the total size of the potential spillover (and presumably the size of the subsidy) represent a <u>second</u> potential source of unobserved heterogeneity. If this unobserved heterogeneity is correlated with incumbent plants' TFP, standard regression equations will be misspecified due to omitted variables, just as described above.

In order to make valid inferences in the presence of the heterogeneity associated with the plant's expected local production costs and the county's value of attracting the plant, knowledge of the exact form of the selection rule that determines plants' location decisions is generally necessary. As the BMW example demonstrates, the two factors that determine plant location decisions are generally unknown to researchers and, in the rare cases where they are known, are difficult to measure. Thus, the effect of a plant opening on incumbents' TFP is very likely to be confounded by differences in factors that determine the plants' profitability at the chosen location.

As a solution to this identification problem, we rely on the reported location rankings of profit-maximizing firms to identify a valid counterfactual for what would have happened to incumbent plants in winning counties in the absence of the plant opening. We implement the research design using data from the corporate real estate journal *Site Selection*. Each issue of this journal includes an article titled the "Million Dollar Plants" that describes how a large plant decided where to locate. These articles always report the county that the plant chose (i.e., the

---

[16] The fact that business organizations such as the Chambers of Commerce support these incentive plans (as was the case with BMW) suggests that incumbent firms expect such increases. Greenstone and Moretti (2004) present a model that describes the factors that determine local governments' bids for these plants and whether successfully attracting a plant will be welfare increasing or decreasing for the county.

'winner'), and usually report the runner-up county or counties (i.e., the "losers").[17] As the BMW case study indicates, the winner and losers are usually chosen from an initial sample of "semi-finalist" sites that in many cases number more than a hundred.[18] The losers are counties that have survived a long selection process, but narrowly lost the competition.

We use the losers to identify what would have happened to the productivity of incumbent plants in the winning county in the absence of the plant opening. Specifically, we assume that incumbent firms' TFP would have trended identically in the absence of the plant opening in winning and losing counties within a case. In practice, we adjust for covariates so our identifying assumption is weaker. The subsequent analysis provides evidence that supports the validity of this assumption. Even if this assumption fail to hold, we presume that this pairwise approach is more reliable than using regression adjustment to compare the TFP of incumbent plants in counties with new plants to the other 3,000 United States counties or to using a matching procedure based on observable variables.[19]

III. Data Sources and Summary Statistics

A. Data Sources

The "Million Dollar Plants" articles typically reveal the county where the new firm (the "Million Dollar Plant") ultimately chose to locate (the "winning county"), as well as the one or two runner-up counties (the "losing counties"). The articles tend to focus on large plants that are the target of local government subsidies. Important limitations of these articles are that the magnitude of the subsidy offered by the winning counties is in many cases unobserved and that the bid is almost always unobserved for losing counties.

We identify the Million Dollar Plants in the Standard Statistical Establishment List (SSEL) – which is the Census Bureau's "most complete, current, and consistent data for U.S.

---

[17] In some instances the "Million Dollar Plants" articles do not identify the runner-up county. For these cases, we did a Lexis/Nexis search for other articles discussing the plant opening and in 4 cases, among the original 82, we were able to identify the losing counties. The Lexis/Nexis searches were also used to identify the plant's industry when this was unavailable in *Site Selection*. Comprehensive data on the subsidy offered by winning and losing counties is unavailable in the *Site Selection* articles.

[18] The names of the semi-finalists are rarely reported.

[19] Propensity score matching is an alternative approach (Rosenbaum and Rubin 1983). Its principal shortcoming relative to our approach is its assumption that the treatment (i.e., winner status) is "ignorable" conditional on the observables. As it should be clear from the example, adjustment for observable variables through the propensity score is unlikely to be sufficient.

business establishment"[20] – and matched the plants to the Annual Survey of Manufactures (*ASM*) and the Census of Manufactures (*CM*) from 1973-1998.[21] Of the 82 MDP openings in Greenstone and Moretti (2004), we identified 47 genuine and useable MDP openings in the manufacturing data. In order to qualify as a genuine and useable MDP manufacturing opening, we imposed the following criterion: 1) there had to be a new plant in the manufacturing sector, owned by the reported firm, appearing in the SSEL within 2 years before and 3 years after the publication of the MDP article; 2) the plant identified in the SSEL had to be located in the county indicated in the MDP article; and 3) there had to be incumbent plants in both winning and losing counties that were there for each of the previous 8 years. Among the 35 MDP openings that did not qualify, roughly 20 were outside of the manufacturing sector. (We cannot report the exact number because of the Census Bureau's confidentiality rules).

To obtain information on incumbent establishments in winner and loser counties, we use the *ASM* and *CM*. The *ASM* and *CM* contain information on employment, capital stock, total value of shipments, plant age, and firm identifiers. The 4-digit SIC code and county of location are also reported and these play a key role in the analysis. Importantly, the manufacturing data contain a unique plant identifier, making it possible to follow individual plants over time.[22] The sample that we use includes plants that were continuously present in the *ASM* in the 8 years preceding the year of the plant opening plus the year of the opening. Additionally, we drop observations on plants that have the same owner as the MDP plants. In this period, the *ASM* sampling scheme was positively related to firm and plant size. Any establishment that was part of a company with manufacturing shipments exceeding $500 million was sampled with certainty, as were establishments with 250 or more employees.

There are a few noteworthy features of this sample of potentially affected plants. First, the focus on existing plants allows for a test of spillovers on a fixed sample of pre-existing plants, which eliminates concerns related to the endogenous opening of new plants and compositional bias. Second, it is possible to form a genuine panel of manufacturing plants. Third, a disadvantage is that the results may not be externally valid to smaller incumbent plants

---

[20] The SSEL is confidential and was accessed in a Census Data Research Center. The SSEL is updated continuously and incorporates data from all Census Bureau economic and agriculture censuses and current business surveys, quarterly and annual Federal income and payroll tax records, and other Departmental and Federal statistics and administrative records programs.

[21] The sample is cut at 1998 because sampling methods in the *ASM* changed for 1999. The sample begins in 1973 because of minor known inconsistencies with the 1972 *CM*.

[22] See the appendix in Davis, Haltiwanger, and Schuh (1996) for a more thorough description of the *ASM* and *CM*.

that are not sampled with certainty throughout this period. Nevertheless, it is relevant that this sample of plants accounts for 54% of county-wide manufacturing shipments in the last *CM* before the MDP opening.

Besides testing for an average spillover effect, we also test whether the estimated agglomeration effects are larger in industries that are more closely linked to the MDP based on some measure of economic distance. We use six measures of economic distance in three categories. First, to measure supplier and customer linkages, we use data on the fraction of each industry's manufactured inputs that come from each 3-digit industry and the fraction of each industry's outputs sold to manufacturers that are purchased by each 3-digit industry. Second, to measure the frequency of worker mobility between industries, we use data on labor market transitions from the *Current Population Survey* (*CPS*) outgoing rotation file. In particular, we measure the fraction of separating workers from each 2-digit industry that move to firms in each 2-digit industry. Third, to measure technological proximity, we use data on the fraction of patents manufactured in a 3-digit industry that cite patents manufactured in each 3-digit industry. We also use data on the amount of R&D expenditure in a 3-digit industry that is used in other 3-digit industries.

Finally, one further data issue merits attention. We have two sources of information on the date of the plant opening. The first is the MDP articles, which often are written when ground is broken on the plant but other times are written when the location decision is made or the plant begins operations. The second source is the SSEL, which in principle reports the plant's first year of operation. However, it is known that plants occasionally enter the SSEL after their opening.

Thus, there is uncertainty about the date of the plant's opening. Further, the date at which the plant could affect the operations of existing plants depends on the channel for any agglomeration economies. If the agglomeration economies are a consequence of supplier relationships, then they could occur as soon as the plant is announced. For example, the new plant's management might visit existing plants and provide suggestions on operations. Alternatively, the agglomeration spillovers may be driven by the labor market and therefore may depend on sharing labor. In this case, agglomeration economies may not be evident until the plant is operating. Based on these data and conceptual issues, there is not clear guidance on when the new plant could affect other plants. Rather than take an unsupportable stand, we emphasize results using the earlier of the year of the publication of the magazine article and the year that the new plant appears in the SSEL. We also report separate results based on using these two years as the opening date. The basic findings are robust to these alternatives.

B. Summary Statistics

Table 1 presents summary statistics on the sample of plant location decisions that forms the basis of the analysis. As discussed in the previous subsection, there are 47 manufacturing MDP openings that we can match to plant level data. There are plants in the same 2-digit SIC industry in both winning and losing counties in the 8 years preceding the opening for just 16 of these openings.

The table reveals some other facts about the plant openings.[23] We refer to the winner and accompanying loser(s) associated with each plant opening announcement as a "case." There are two or more losers in 16 of the cases, so there are a total of 73 losing counties along with 47 winning counties. Some counties appear multiple times in the sample (as either a winner or loser), and the average county in the sample appears a total of 1.09 times. The difference between the year of the MDP article's publication and the year the plant appears in the SSEL is roughly spread evenly across the categories -2 to -1 years, 0 years, and 1 to 3 years. For clarity, positive differences refer to cases where the article appears after the plant is identified in the SSEL. The date of the plant openings ranges from the early 1980s through the early 1990s.

The remainder of Table 1 provides summary statistics on the new MDP plants five years after their assigned opening date to provide a sense of their magnitude. These MDPs are quite large: they are more than twice the size of the average incumbent plant and account for roughly nine percent of the average county's total output one year prior to its opening.

Table 2 provides summary statistics on the measures of industry linkages and further descriptions of these variables. In all cases, the proximity between industries is increasing in the value of the variable. For ease of interpretation in the subsequent regressions, these variables are normalized to have a mean of zero and a standard deviation of one.

Table 3 presents the means of county-level and plant-level variables across counties. These means are reported for winners, losers, and the entire United States in columns (1), (2),

---

[23] A number of the statistics in Table 1 are reported in broad categories to comply with the Census Bureau's confidentiality restrictions and to avoid disclosing the identities of any individual plants.

and (3), respectively.[24] In the winner and loser columns, the plant-level variables are calculated among the incumbent plants present in the *ASM* in the 8 years preceding the assigned opening date and the assigned opening date. All entries in the entire United States column are weighted across years to produce statistics for the year of the average MDP opening in our sample. Further, the plant characteristics are only calculated among plants that appear in the *ASM* for at least 9 consecutive years. Column (4) presents the t-statistics from a test that the entries in (1) and (2) are equal, while Column (5) repeats this for a test of equality between columns (1) and (3). Columns (6) through (10) repeat this exercise among the cases where there are plants within the same 2-digit SIC industry as the MDP plant. In these columns, the plant characteristics are calculated among the plants in the same 2-digit industry.

This exercise provides an opportunity to assess the validity of the research design, as measured by pre-existing observable county and plant characteristics. To the extent that these observable characteristics are balanced among winning and losing counties, this should lend credibility to the analysis. The comparison between winner counties and the rest of the United States provides an opportunity to assess the validity of the type of analysis that would be undertaken in the absence of a quasi-experiment.

The top panel reports on county-level characteristics measured in the year before the assigned plant opening and the percentage change between 7 years and 1 year before the opening. It is evident that compared to the rest of the country, winning counties have higher incomes, population and population growth, labor force participation rates and growth, and a higher share of labor in manufacturing. Among the 8 variables in this panel, 6 of the 8 differences would be judged to be statistically significant at conventional levels. These differences are substantially mitigated when the winners are compared to losers, and this is reflected in the fact that 3 of the 8 variables are statistically different at the 5% level but none are at the 1% level. Notably, the raw differences between winners and losers within the subset of cases where there are plants in the same 2 digit SIC industry are generally smaller, and none of them would be judged to be statistically significant.

The second panel reports on the number of sample plants and provides information on

---

[24] The losing county entries in column (2) are weighted in the following manner. Losing counties are weighted by the inverse of their number in that case. Losing plants are weighted by the inverse of their number per-county, multiplied by the inverse of the number of losing counties in their case. The result is that each county (and each plant within each county) is given equal weight within the case and then all cases are given equal weight.

some of their characteristics. In light of our sample selection criteria, the number of plants is of special interest. On average, there are 18.8 plants in the winner counties and 25.6 in the loser ones (and just 8.0 in the United States). The covariates are well balanced between plants in winning and losing counties; in fact, there are no statistically significant differences either among all plants or among the plants within the same 2-digit industry.[25]

Overall, Table 3 has demonstrated that the MDP winner-loser research design balances many (although not all) observable county-level and plant-level covariates. In the subsequent analysis, we demonstrate that trends in TFP were similar in winning and losing counties prior to the opening of the MDP, which lends further credibility to this design. Of course, this exercise does not guarantee that unobserved variables are balanced across winner and loser counties or their plants. The next section outlines our full econometric model and highlights the exact assumptions necessary for consistent estimation.

IV. Econometric Model

Building on the model in section I, we start by assuming that incumbent plants use the following Cobb-Douglas technology:

$$(6) \qquad Y_{pijt} = A_{pijt}\, L^{\beta 1}{}_{pijt}\, K^{B\,\beta 2}{}_{pijt}\, K^{E\,\beta 3}{}_{pijt}\, M^{\beta 4}{}_{pijt}$$

where p references plant, i industry, j case, and t year; Y_{ijct} is the total value of shipments; A_{pijt} is TFP; and we allow total labor hours of production $\left(L_{pijt}\right)$, building capital stock $\left(K^{B}_{pijt}\right)$, machinery and equipment capital stock $\left(K^{E}_{pijt}\right)$, and the dollar value of materials $\left(M_{pijt}\right)$ to have separate impacts on output. In practice, the two capital stock variables are calculated with the permanent inventory method that uses earlier years of the data on book values and subsequent investment.[26]

---

[25] Roughly 20% of the winners were in the Rust Belt, compared to roughly 25% of the losers (where the Rust Belt is defined as MI, IN, OH, PA, NJ, IL, WI, NY). Roughly 65% of the winners were in the South, compared to roughly 45% of the losers.

[26] For the first date available, plants' historical capital stock book values are deflated to constant dollars using BEA data by 2-digit industry. In all periods, plants' investment is deflated to the same constant dollars using Federal Reserve data by 3-digit industry. Changes in the capital stock are constructed by depreciating the initial deflated capital stock using Federal Reserve depreciation rates and adding deflated investment. In each year, productive capital stock is defined as the average over the beginning and ending values, plus the deflated level of capital rentals. The analysis is performed separately for building capital and machinery capital. This procedure is described further by Becker et al. (2005), Chiang (2004), and Davis et al. (1996), from whose files we gratefully obtained deflators.

Recall that equation (1) in Section I allows for agglomeration spillovers by assuming that TFP is a function of the number of firms that are active in a county: $A_{pijt} = A(N_{pijt})$. Here we also allow for some additional heterogeneity in A_{pijt}. In particular, we generalize equation (1) by allowing for permanent differences in TFP across plants (α_p), cases (λ_j), industry-specific time-varying shocks to TFP (μ_{it}), and a stochastic error term (ε_{pijt}):

$$\ln(A_{pijt}) = \alpha_p + \mu_{it} + \lambda_j + \varepsilon_{pijt} + A(N_{pijt}).$$

The goal is to estimate the causal effect of winning a plant on incumbent plants' TFP. To do so, we need to impose some structure on $A(N_{pijt})$. In particular, we use a specification that allows for the new plant in winning counties to affect both the level of TFP as well as its growth over time:

$$\begin{aligned} \ln(A_{pijt}) = {} & \delta\, 1(\text{Winner})_{pj} + \psi\, \text{trend}_{jt} + \Omega\, (\text{trend} * 1(\text{Winner}))_{pjt} \\ & + \kappa\, 1(\tau >= 0)_{jt} + \gamma\, (\text{trend} * 1(\tau >= 0))_{jt} \\ & + \theta_1\, (1(\text{Winner}) * 1(\tau >= 0))_{pjt} + \theta_2\, (\text{trend} * 1(\text{Winner}) * 1(\tau >= 0))_{pjt} \\ & + \alpha_p + \mu_{it} + \lambda_j + \varepsilon_{pijt} \end{aligned} \tag{7}$$

where 1(Winner) is a dummy equal to 1 if plant p is located in a winner county; and τ denotes year, but it is normalized so that for each case the assigned year of the plant opening is $\tau = 0$. The variable trend_{jt} is a simple time trend.

Combining equations (6) and (7) and taking logs, we obtain the regression equation that forms the basis of our empirical analysis:

$$\begin{aligned} \ln\left(Y_{pijt}\right) = {} & \beta_1 \ln\left(L_{pijt}\right) + \beta_2 \ln\left(K^B_{pijt}\right) + \beta_3 \ln\left(K^E_{pijt}\right) + \beta_4 \ln\left(M_{pijt}\right) \\ & + \delta\, 1(\text{Winner})_{pj} + \psi\, \text{trend}_{jt} + \Omega\, (\text{trend} * 1(\text{Winner}))_{pjt} + \kappa\, 1(\tau >= 0))_{jt} + \gamma\, (\text{trend} * 1(\tau >= 0))_{jt} \\ & + \theta_1\, (1(\text{Winner}) * 1(\tau >= 0))_{pjt} + \theta_2\, (\text{trend} * 1(\text{Winner}) * 1(\tau >= 0))_{pjt} \\ & + \alpha_p + \mu_{it} + \lambda_j + \varepsilon_{pijt}. \end{aligned} \tag{8}$$

Equation (8) is an augmented Cobb-Douglass production function that allows labor, building capital, machinery capital, and materials to have differential impacts on output. The paper's focus is the estimation of the impact of the new plant on incumbent plants' TFP, so the parameters of interest are θ_1 and θ_2, which are the spillover effects. The former tests for a mean shift in TFP among incumbent plants in the winning county after the opening of the MDP, while the latter allows for a trend break in TFP among the same plants.

In practice, we estimate two variants of Equation (8). In some specifications, we fit a parsimonious model that simply tests for a mean shift. In this model, any productivity effect is assumed to occur immediately and to remain constant over time. Specifically, we make the restrictions that $\psi = \Omega = \gamma = \theta_2 = 0$, which assumes that differential trends are not relevant here. This specification is essentially a difference in differences estimator and we refer to it as Model 1. Formally, after adjustment for the inputs, $1(\text{Winner})_p$, and $1(\tau >= 0)_{jt}$, the consistency of θ_1 in this model requires the assumption that $E[(1(\text{Winner}) * 1(\tau >= 0))_{pjt}\, \varepsilon_{pijt} | \alpha_p, \mu_{it}, \lambda_j] = 0$.

In other specifications, we estimate the model without imposing such restrictions on the trends. In other words, we estimate the entire Equation (8). We label this Model 2. While Model 1 only allows for a mean shift in productivity, this specification allows both for a mean shift and a trend break in productivity. In other words, Model 2 allows us to investigate whether any productivity effect occurs immediately and whether the impact evolves over time. This specification is demanding of the data, because our sample is only balanced through $\tau = 5$ so there are only 6 years per case to estimate θ_1 and θ_2.[27]

Equation (8) allows for unobserved determinants of TFP that are unrelated to the plant opening but could be confounded with the spillover effects if not properly accounted for. It includes a differential intercept for all observations from winning counties, δ. It also allows for a common time trend, ψ. The parameter Ω allows the time trend to differ for winning counties prior to the plant opening. This will serve as an important way to assess the validity of this research design. Finally, κ and γ capture the level change and trend break in TFP common to plants in winning and losing counties after the MDP opening (i.e., when $\tau >= 0$).

In addition, all our models include three sets of fixed effects. First, they include separate fixed effects for each plant, α_p, so the comparisons are within a plant. Second, μ_{it} represents the parameters associated with a vector of 2-digit SIC industry by year fixed effects to account for industry-specific shocks to TFP. Third, the λ_j's are separate fixed effects for each case that ensure that the impact of the MDP's opening is identified from comparisons within a winner-loser pair; they are a way to retain the intuitive appeal of pairwise differencing in a regression framework.

A few further estimation details bear noting. First, unobserved demand shocks are likely

---

[27] This specification allows for spillovers to affect the level of TFP and to grow over time as in Glaeser and Mare (2001).

to affect input utilization, and this raises the possibility that the estimated β's are inconsistent (see, e.g., Griliches and Mairesse 1995). This has been a topic of considerable research and we are unaware of a bullet-proof solution. We implement the standard fixes including modeling the inputs with alternative functional forms (e.g., the translog), using cost shares at the plant and industry-level rather than estimating the β's, and controlling for flexible functions of investment, capital and materials (Syverson 2004; van Biesebroeck 2004; Olley and Pakes 1996; Levinsohn and Petrin 2003). Additionally, we experiment with adding fixed effects for region by year or region by industry by year, and allowing the effect of inputs to differ by industry or by winner- and post-MDP status. The basic results are unchanged by these alterations in the specification. We also note that unobserved demand shocks are only a concern for the consistent estimation of the parameters of interest (i.e., θ_1 and θ_2) if they systematically affect incumbent plants in winning counties in the years after the MDP's opening, after adjustment for the rich set of covariates in equation (8).

Second, in some cases this equation is fit on a sample of plants from the entire country, but in most specifications the sample is limited to plants from winning and losing counties in the *ASM* for every year from $\tau = -8$ through $\tau = 0$. When data from the entire country is used, the sample is limited to plants that are in the *ASM* for at least 14 consecutive years. The smaller sample of plants from the winning and losing counties allows for the impact of the inputs and the industry shocks to differ in the winning-losing county sample from the rest of the country. Finally, for most of the analysis, we further restrict the sample to observations in the years between $\tau = -7$ and $\tau = 5$. Due to the dates of the MDP openings, this is the longest period for which we have data from all cases.[28]

Third, we probe the validity and robustness of our estimates with a number of alternative specifications. For example, we investigate whether changes in the price of output, capital utilization, public investment, or attrition influence the estimates.

Fourth, all of the reported standard errors are clustered at the county level to account for the correlation in outcomes among plants in the same county.[29]

Fifth, we focus on weighted versions of equation (8). Specifically, the specifications are

---

[28] Data from all cases is also available for $\tau = -8$, but output in this period is used to weight the regressions.

[29] We experimented with clustering the standard errors at the 2-digit SIC by county level, but this occasionally produced variance-covariance matrices that weren't positive definite. In instances where they were positive definite, these standard errors were similar to those from clustering at the county level.

weighted by the square root of the total value of shipments in τ = -8 to account for heteroskedasticity associated with differences in plant size. This weighting also means that the results measure the change in productivity for the average dollar of output, which in our view is more meaningful than the impact of the MDP on the average plant.

The analysis will also explore two additional issues. It will report on the fitting of versions of equation (8) that interact the spillover variables with measures of economic distance between the MDP and the incumbent plant. These specifications assess whether the magnitude of the estimated spillovers varies with economic distance. Finally, the paper will assess whether the MDP's opening affects local skill adjusted wages and the entry and exit decisions of plants in the MDP's county.

V. Results

This section is divided into four subsections. The first reports baseline estimates of the effect of the opening of a new Million Dollar Plant on the productivity of incumbent plants in the same county through the estimation of equation (8). The second subsection discusses the validity of our design and explores the robustness of our estimates to a variety of different specifications. The third subsection explores potential channels for the agglomeration effects by testing whether the estimated spillovers vary as a function of economic distance. The final subsection discusses the implications of our estimates for the profits of local firms.

A. Baseline Estimates

Columns (1) and (2) of Table 4 report estimated parameters and their standard errors from a version of equation (8). Specifically, the natural log of output is regressed on the natural log of inputs, year by 2-digit SIC industry fixed effects, plant fixed effects, and the event time indicators in a sample that is restricted to the years τ = -7 through τ = 5. The parameters associated with the event time indicators report mean TFP in winning and losing counties, respectively, in each event year relative to the year before the MDP opened (i.e., we have subtracted off the τ = -1 parameter estimates for the winners (losers) from the winner (loser) estimates for each event year). Column (3) reports the difference between the estimated TFP levels within each year.

The top panel of Figure 1 separately plots the mean TFP levels for winner and loser counties (taken from columns (1) and (2) of Table 4) against τ. The bottom panel of Figure 1 plots the difference in the estimated winner and loser coefficients against τ. Thus, it is a graphical version of column (3) of Table 4.

Two important findings are apparent in these figures. First, the trends in TFP among incumbent plants were very similar in the winning and losing counties in the years before the MDP opening. In fact, a statistical test fails to reject that the trends were equal. This finding supports the validity of our identifying assumption that incumbent plants in losing counties provide a valid counterfactual for incumbents in winning counties.

Second, there is a sharp upward break in the difference in TFP between the winning and losing counties beginning with the year that the plant opened. The top graph reveals that this relative improvement is due to the continued decline in TFP in losing counties and a flattening out of the TFP trend in winning counties. The figures also serve to underscore the importance of the availability of losing counties as a counterfactual. For example, a naïve comparison of mean TFP in winning counties before and after the opening would lead to the conclusion that the opening had a small negative impact on incumbents' TFP. Overall, these graphs reveal much of the paper's primary finding. This relative increase in TFP among incumbent plants in winning counties will be confirmed throughout the battery of tests in the remainder of the paper.

Before proceeding, it is worth noting that the TFP of incumbent plants was declining in both sets of counties in advance of the MDP plant opening. This finding may appear surprising, because productivity increases over time in the overall economy. However, this downward trend in TFP among large manufacturing plants appears to be a general phenomenon and is not specific to plants in winning and losing counties. For example, we estimated augmented Cobb-Douglas production functions where the constant dollar total value of shipments is the dependent variable and the covariates are the capital stock, labor, materials (also in constant dollars), 2-digit industry fixed effects, and plant fixed effects. The equations are weighted by the total value of shipments and the sample is restricted to plants in the *ASM* for at least 14 consecutive years in the period 1973-1998. Thus, the approach and sample selection rule are similar in spirit to the ones used in the MDP analyses.

The average 6 year change in TFP calculated over the years preceding the MDP openings was a statistically significant decline of 4.7%.[30] This decline is similar to the decline among incumbents in winning and losing counties in Figure 1. To the best of our knowledge, this finding of declining TFP among large manufacturing plants has not been noted previously.[31]

Now turning to the statistical models, the first four columns of Table 5 present the results from fitting different versions of equation (8). Models 1 and 2 are in Panels A and B, respectively. Panel A reports the estimated mean shift parameter, θ_1, and its standard error (in parentheses) in the "Mean Shift" row. Panel B reports the estimated impact of the MDP on incumbent plants' TFP evaluated at $\tau = 5$ in the "Effect after 5 years" row, which is determined by θ_1 and θ_2 that are also both reported.[32] The row "Pre-Trend" contains the coefficient measuring the difference in the pre-existing trends between plants in the winning and losing counties. In all of these specifications, the estimated impact of the MDP's opening is determined during the period where $-7 \leq \tau \leq 5$, as the sample is balanced during these years.

In columns (1) and (2), the sample includes all manufacturing plants in the *ASM* that report data for at least 14 consecutive years, excluding all plants owned by the MDP firm. In column (3), the sample is restricted to include only plants in counties that won or lost a MDP. This restriction means that the impact of the inputs and the industry-year fixed effects are estimated solely from plants in these counties. Incumbent plants are now required to be in the data only for $-8 \leq \tau \leq 0$ (not also for 14 consecutive years, though this does not change the results). Finally in column (4), the sample is restricted further to include only plant-year observations within the period of interest (where τ ranges from -7 through 5). This forces the input parameters and industry-year fixed effects to be estimated solely on plant by year observations that identify the spillover parameters. This sample is used throughout the remainder of the paper. Estimation details are noted at the bottom of the table and apply to both Models 1 and 2.

The entries in Table 5 confirm the visual impression from Figure 1 that the opening of the

---

[30] The 6 year average is a weighted averaged calculated over the 6 year periods before each of the MDP openings where the weights are the number of plant openings associated with each 6 year period. For example, if there are 2 plant openings in 1987 and 1 in 1988, then the average change between 1980 and 1986 receives twice the weight as the average change calculated between 1981 and 1987.

[31] Since we are looking at large plants that have been active for a large number of years, we speculate that this decline may have to do with aging. Additionally, many of the years preceding the MDP openings were in the late 1970s and early 1980s, which was a period of poor economic performance. Foster, Haltiwanger, and Krizan (2000) have documented that within plant productivity growth is positively correlated with the economic cycle.

[32] This is calculated as $\theta_1 + 6\theta_2$, because we allow the MDP to affect outcomes from $\tau = 0$ through $\tau = 5$.

MDP is associated with a substantial increase in TFP among incumbent plants in winning counties. Specifically, Model 1 implies an increase in TFP of roughly 4.8%. As the figure highlighted, however, the impact on TFP appeared to be increasing over time so Model 2 seems more appropriate. This model's results suggest that the MDP's opening is associated with an approximately 12% increase in TFP five years later. The estimates from both models would be judged to be statistically different from zero by conventional criteria and are unaffected by the changes in the specifications. Furthermore, the entries in the "Pre-trend" row demonstrate that the null hypothesis of equal trends in TFP among incumbents in winning and losing counties cannot be rejected.

The numbers in square brackets in column 4 measure the average size of the spillover from a MDP opening in millions of 2006\$. This figure is calculated by multiplying the estimated impact by the mean value of incumbent plants' total shipments in winning counties in $\tau = -1$. This calculation indicates that the increase in TFP from a MDP was associated with an increase in total output of about \$170 million per year in Model 1. The Model 2 estimate is even larger, suggesting an increase in output of roughly \$429 million in year $\tau = 5$. These numbers are large, with the Model 2 effect at $\tau = 5$ nearly the average level of MDP output.

Column (5) presents the results from a "naïve" estimator that is based on using plant openings without an explicit counterfactual. Specifically, a set of 47 plant openings were randomly chosen from the *ASM* in the same years and industries as the MDP openings. The remainder of the sample includes all manufacturing plants in the *ASM*, not also owned by the randomly chosen firm, that report data for at least 14 consecutive years. With these data, we fit a regression of the natural log of output on the natural log of inputs, year by 2-digit SIC fixed effects, and plant fixed effects. In Model 1, two additional dummy variables are included for whether the plant is in a winning county 7 to 1 years before the randomly chosen opening or 0 to 5 years after. The reported mean shift is the difference in these two coefficients (i.e., the average change in TFP following the opening). In Model 2, the same two dummy variables are included along with pre- and post-trend variables. The shift in level and trend are reported, along with the pre-trend and the total effect evaluated after 5 years.

This naïve "first-difference" style estimator indicates that the opening of a new plant is associated with a -6% to -8% effect on incumbent plants' TFP, depending on the model. If the estimates from the MDP research design are correct, then this naïve approach understates the extent of spillovers by 13% (Model 1) to 18% (Model 2). Interestingly, the parameter on the "pre-trend" indicates that the TFP of the incumbent plants was on a downward trend in advance of the openings in the counties that attracted these new plants. This is similar to what is observed in our MDP sample of winners. Overall, the primary message is that the absence of a credible research design can lead to misleading inferences in this setting.

It is natural to wonder about the degree of heterogeneity in the treatment effects from the 47 separate case studies that underlie the estimates presented thus far. Figure 2 explores this heterogeneity by plotting case-specific estimates of parameter θ_1 in Model 1 and their 95% confidence interval. Specifically, the Figure plots results from a version of Model 1 that interacts the variable (1(Winner) * 1(τ >= 0)) with indicators for each of the cases. There are 45 estimates of θ_1, one for each case. (Results from two cases were omitted to comply with the Census Bureau's confidentiality rules). The figure reveals that there is heterogeneity in the estimated impacts on TFP of incumbent plants. 27 of the 45 estimates are positive. 13 of the positive estimates would be judged to be statistically different from zero at the 5% level, while the comparable figure for the negative estimates is 9.

We assessed whether the estimated spillover effects are related to characteristics of the MDP. Specifically, we regressed the estimates against three measures of the MDP's size, whether the MDP is owned by a foreign company, and whether it is an auto company. When these multiple measures were included jointly, none were significantly related to the estimated effect of the MDP's opening.[33]

Ultimately, TFP is a residual and residual labeling must be done cautiously. As an alternative way to examine the impact of the MDP on incumbents' productivity, we have estimated directly the impacts of a MDP opening on output (unadjusted for inputs) and inputs. The intent is to contrast the changes in outputs and inputs to shed light on whether productivity increased without imposing the structural assumptions of the production function. Put another way, are the incumbents producing more with less after the MDP's opening? Appendix Table 1 reports on estimates of the impact of a MDP opening on incumbents' output and usage of inputs.

---

[33] Separate regressions of the case specific effects on the MDP's total output or the MDP's total labor force generated statistically significant negative coefficients. This result is consistent with the possibility that when the MDP is very large incumbents are left to hire labor and other inputs that are inferior in unobserved ways. On the other hand, we failed to find any significant differences when separately testing whether the productivity effect varied by the ratio of the MDP's output to county-wide manufacturing output, whether the MDP is owned by a foreign company, or whether the MDP is an auto company.

The estimates are from the Model 1 and Model 2 versions of equation (8) with the key difference that these equations do not include the inputs as covariates. Again, we use the Model 2 results to estimate the impact of the opening 5 years afterwards.

Column (1) reports that the MDP opening is associated with an 8-12% increase in output. Columns (2) through (5) report the results for the four inputs. It is striking that the change in all of the inputs is roughly equal to or less than the increase in output. The Model 2 results are especially noteworthy, because the 8% increase in output is accompanied by no increase in either form of capital. Overall, the results in Appendix Table 1 indicate that, after the MDP's opening, incumbent plants produced more with less; that is, they suggest that these plants became more productive, and this is consistent with the TFP increases uncovered in Table 5.[34]

B. Threats to Validity

Estimates in Table 5 appear to be consistent with significant agglomeration spillovers generated by MDP openings. Although the comparisons in Table 3 and the similarity of the pre-existing trends in TFP in winning and losing counties support the validity of the research design, it is, of course, possible that there is a form of unobserved heterogeneity that accounts for the higher levels of TFP in winning counties after the MDP's opening. Consequently, this subsection investigates several possible alternative interpretations of the estimated spillover effects and explores the robustness of the estimates to a variety of different assumptions. Specifically, we investigate (i) the role of functional form assumptions and the possible presence of unobserved industry and regional shocks; (ii) the endogeneity of inputs; (iii) changes in the price of output; (iv) the possible role of public investment; (v) changes in capital utilization; and (vi) attrition.

(i) Functional Form, Industry Shocks and Regional Shocks. Table 6 reports on a series of specification checks. For convenience, column (1) reports the results from the preferred specification in column (4) specification of Table 5. These estimates are intended to serve as a basis of comparison for the estimates in the remainder of the table.

We begin by generalizing our assumption on technology. Estimates in Table 5 assume a Cobb-Douglas technology. In column (2) of Table 6, the inputs are modeled with the translog functional form. Column (3) is based on a Cobb-Douglas technology but allows the effect of

---

[34] Column (6) of Appendix Table 1 presents evidence on changes in the capital/labor ratio. The model suggests that firms should substitute away from labor and toward capital. The estimated change in the capital/labor ratio is poorly determined, making definitive conclusions unwarranted, but the point estimate is not supportive of this prediction.

each production input to differ at the 2-digit SIC level. This model accounts for possible differences in technology across industries, as well as for possible differences in the quality of inputs used by different industries. For example, it is possible that even if technology was similar across different manufacturers, some industries use more skilled labor than others. Column (4) allows the effect of the inputs to differ in winning/losing counties and before/after the MDP opening.

Columns (5) and (6) add census division by year fixed effects and census division by year by 2-digit industry fixed effects. These specifications aim to purge the spillover effects of unobserved region-wide shocks or region by industry shocks to productivity that might be correlated with the probability of winning a MDP.

Taken together, the results in columns (2) through (6) of Table 6 are striking. The estimates appear to be insensitive to the specific functional form of the production function. None of the specifications contradict the findings from the baseline specification in Table 5. Although many of the estimates are smaller than the baseline ones, the magnitude of the decline is modest. For example, they are all within one standard error of the baseline estimate in both Models 1 and 2. Overall, these results fail to undermine the conclusion from Table 5 that the opening of a MDP leads to a substantial increase in TFP among incumbent plants and this is consistent with theories of spillovers.

(ii) Endogeneity of Inputs. An important conceptual concern is that capital and labor inputs should be treated as endogenous, because the same forces that determine output also determine a firm's optimal choice of inputs (Griliches and Mairesse 1995). Unlike the usual estimation of production functions, our aim is the consistent estimation of the spillover parameters, θ_1 and θ_2, so the endogeneity of capital and labor is only relevant to the extent that it results in biased estimates of these parameters. This subsection employs the productivity literature's state of the art techniques to control for the endogeneity of capital and labor to assess this issue's relevance in this paper's setting.

We do this in two ways. First, in columns (7) and (8) of Table 6, we calculate TFP for each plant by fixing the parameters on the inputs at the relevant input's share of total costs (van Biesebroeck 2004; Syverson 2004; Foster, Haltiwanger, and Syverson 2007). This method may mitigate any bias in the estimation of the parameters on the inputs associated with unobserved demand shocks. In these two columns, the cost shares are calculated at the plant level and the 3-

digit SIC industry level over the full sample, respectively. The estimated spillover effects are largely insensitive to this restriction.

Second, Table 7 presents estimates based on the widely-used methodologies proposed by Olley and Pakes (1996) and Levinsohn and Petrin (2003). These methods are based on the result that, under certain conditions, adjustment for investment or intermediate inputs (e.g., materials) will remove the correlation between input levels and unobserved shocks to output. For example, the column (2) specification adds 4th degree polynomial functions of log building investment and log machinery/equipment investment to the baseline specification.[35] The column (3) specification adds 4th degree polynomials in the two types of log capital stocks to the column (2) equation. The column (4) specification is even richer as it adds all the "own" interactions between polynomials in current investment and capital (i.e., the building investment polynomial is interacted with the building capital stock polynomial, but it is not also interacted with the machinery/equipment polynomials for stocks or investment) to column (3). Column (5) adds a 4th degree function of log materials to the baseline specification. In the column (6) specification, a 4th degree polynomial in materials is fully interacted with 4th degree polynomials in building capital and machinery capital. Column (7) includes fourth-degree polynomials in log materials and log investment and log capital stock for both types of capital (not interacted). Finally, column (8) includes the controls from the columns (4) and (6) specifications.

The estimated spillover effects in Table 7 are generally consistent with the findings from the baseline specification. This finding holds in both Models 1 and 2. Overall, this exercise fails to suggest that the possible endogeneity of labor and capital is the source of the estimated productivity spillovers.

(iii) Changes in the Price of Output. Another concern is that the theoretically correct dependent variable is the *quantity* of output. However, due to the data limitations faced by virtually all of the rest of the productivity literature, the dependent variable in our models is the *value* of output or price multiplied by quantity. Consequently, it is possible that the estimated spillover effect reflects higher prices, instead of higher productivity.

We do not expect this to be a major factor in our context. First, the sample is comprised of manufacturing establishments that generally produce nationally traded goods. Therefore, it is

---

[35] The exact measure used is log(1+investment), so zero values are not dropped. The results are very similar when including polynomial functions of the level of investment and a dummy variable for values equal to zero.

likely that in many cases the price of output is set at the national level, and has little to do with what happens in the county where the goods are produced. In the extreme case of a perfectly competitive industry that produces a nationally traded good, there should be no effect on prices.

Second, we tested whether the size of the estimated productivity effect is larger in industries that are more regional or more concentrated. The idea is that if price increases are possible, then they should be larger in industries that are more local and/or in industries that are less competitive. Consider for example, the case of an industry that sells mainly at the local level (e.g., cement). The opening of a large new plant may ultimately increase incumbents' demand by raising local income (even though the initial effect on demand should be negative). If the industry is not very competitive, the increase in demand may ultimately lead to price increases for the incumbents' output.

To implement the test, we estimated a Model 1 version of equation (8) that interacts $1(\text{Winner})_{pj}$, $1(\tau >= 0))_{jt}$, and $(1(\text{Winner}) * 1(\tau >= 0))_{pjt}$ with incumbents' industry-specific measure of average distance traveled by output between production and consumption. We also conducted a similar exercise where these same variables are interacted with a measure of the incumbent's industry concentration.[36] These specifications fail to produce evidence that the estimated spillover effects are larger in more local industries or more concentrated industries; in fact, there is some evidence for larger effects on incumbent plants that ship their products further. Our conclusion is that price increases do not appear to be the source of the estimated spillover effects.

(iv) Public Investment. State and local governments frequently offer substantial subsidies to new manufacturing plants to locate within their jurisdictions. These incentives can include tax breaks, worker training funds, the construction of roads, and other infrastructure investments. It is possible that some of the public investment in infrastructure benefits firms other than the beneficiary of the incentive package. For example, the construction of a new road intended for a MDP plant may also benefit the productivity of some of the incumbent firms. If the productivity gains we have documented are due to public investment, then it is inappropriate to interpret them as evidence of spillovers.

---

[36] The information on distance is from Weiss (1972). Distance varies between 52 and 1337 miles, with a mean of 498. Examples of regional industries are: hydraulic cement, iron and steel products, metal scrap and waste tailings, ice cream and related frozen desserts, and prefabricated wooden buildings. The information on industry concentration is from the Bureau of Census ("Concentration Ratios", 2002).

To investigate this possibility, we estimated the effect of MDP openings on government total capital expenditures and government construction expenditures with data from the *Annual Survey of Governments*. In models similar to equation (8), we find that the opening of a MDP plant is associated with statistically insignificant increases in capital and construction expenditures. In fact, in most specifications the estimated impact of a MDP opening is negative and statistically insignificant. Even in the specifications that produce positive insignificant estimates, there is no plausible rate of return that could generate a meaningful portion of the productivity gains in winning counties. Based on these measures of public investment, it seems reasonable to conclude that public investment cannot explain the paper's results.

(v) Changes in Capital Utilization. Another potential threat to validity is that incumbent plants may respond to the MDPs by increasing the intensity of their capital usage. This could happen if depressed counties where the existing capital stock had been used below capacity win the MDPs and increase production simply by operating their capital stock closer to capacity. As an indirect test of this possibility, we estimated whether the MDP's opening affected the ratio of the dollar value of energy usage (which is increasing in the use of the capital stock) to the capital stock. Column 7 of Appendix Table 1 reports small and insignificant changes in this measure. Thus, we conclude that greater capacity utilization is unlikely to be the source of the findings of productivity spillovers.

(vi) Attrition of Sample Plants. Differential attrition in the sample of incumbent plants in winning and losing counties could contribute to the measured differential in productivity trends among survivors after the MDP's opening. This attrition could result from plants shutting down operations or from plants continuing operations but dropping out of the group of plants that are surveyed with certainty as part of the *ASM*.[37]

The available evidence suggests that differential attrition is unlikely to explain the finding of spillovers in winning counties. First, in the baseline sample (i.e., the one used to produce the results in column (4) of Table 5 and in Tables 6 and 7), 72% of the winning county plants operating in the year of the MDP opening were still in the sample at its end (i.e., $\tau = 5$). The analogous figure in losing counties is 68%. The slightly larger attrition rate in losing counties is consistent with the paper's primary result. Specifically, one seemingly reasonable interpretation

---

[37] Recall, establishments are sampled with certainty if they are part of a company with manufacturing shipments exceeding $500 million or their total employment was at least 250.

of this result is that the MDP's opening allowed some winning county plants to remain open that would have otherwise closed. Thus to the extent that a MDP opening keeps weaker plants operating, the above analysis will underestimate the overall TFP increase. Second, the estimation of equation (8) on the sample of plants that is present for all years from -7 to +5 yields results that are qualitatively similar to those from the full sample. Third, the null hypothesis of equal trends in TFP among attriting plants in winning and losing counties prior to the MDP's opening cannot be rejected; the TFP trend in winning counties minus the TFP trend in losing counties was -0.0052 (0.0080).[38]

C. Estimates of Spillovers by Economic Distance

What can explain the productivity gains uncovered above? Section I A discussed some possible mechanisms that may be responsible for agglomeration spillovers. Tables 8 and 9 attempt to shed some light on the possible mechanisms by investigating how the measured spillover effect varies as a function of economic distance.

By Industry. Table 8 shows separate estimates from the baseline model for samples of incumbent plants in the MDP's 2-digit industry and all other industries. In general, one might expect that the effects of spillovers decline with economic distance (equation 5'). Although looking within-industry does not shed direct light on which channel is the source of the spillovers, it seems reasonable to presume that spillovers would be larger within an industry.

In examining the 2-digit SIC MDP industry results, it is important to recall that just 16 of the 47 cases have plants in the MDP's 2-digit industry. We also note that there can be substantial heterogeneity in technologies and labor forces among the industries within a 2-digit SIC industry. However, this research design and the available data do not permit an examination at finer industry definitions.

Column 1 of Table 8 repeats the all industries estimates from column (4) of Table 5 and is intended to serve as a basis of comparison. Columns (2) and (3) report on estimates from the

---

[38] In addition to the specification checks described in this section, we also tested whether the results are sensitive to the choice of the date of the MDP's opening. The estimated spillovers are virtually unchanged when we use the year that the plant is first observed in SSEL as the MDP's opening date. When the year of the MDP article in *Site Selection* is used as the plant's opening date, the Model 1 results are nearly identical to those in the Table 5 column (4) specification and roughly 5% in the Model 2 specification. When the estimating equation is unweighted, the evidence in favor of spillover effects is weaker indicating that the spillovers are concentrated among the larger plants in the sample. As discussed above, our view is that the economically relevant concept of spillover is the change in productivity for the average dollar of output, rather than the average plant.

baseline specification for incumbent plants in the MDP's 2-digit industry and all other industries, respectively. The entries in these columns are from the same regression. Just as in Table 5, the numbers in square brackets convert the parameter estimates into millions of 2006$.

The impacts are substantially larger in the own 2-digit industry. For example, the estimated increase in TFP for plants in the same 2-digit industry is a statistically significant 17% in Model 1 and a poorly determined 33% at $\tau = 5$ in Model 2. In contrast, the estimates for plants in other industries are a statistically insignificant 3.3% in Model 1 and marginally significant 8.9% in Model 2. These basic findings in the own 2-digit and other industries are robust to the different specifications in tables 6 and 7.[39]

Figures 3 and 4 provide 2-digit MDP industry and other industry analogues to Figure 1. Importantly, there is not evidence of differential trends in the years before the MDP's opening and statistical tests confirm this visual impression. The 2-digit MDP industry estimates are noisy due to the small sample size, which was also evident in the statistical results. Just as in Figure 1, the estimated impact reflects the continuation of a downward trend in TFP in losing counties and a cessation of the downward trend in winning counties.

By Direct Measure of Economic Proximity. Having found that the spillover is larger for incumbent plants in the MDP industry, we investigate the role of economic proximity more directly by using several explicit measures of economic proximity that capture worker flows, technological proximity, and input-output flows. To ease the interpretation, the economic proximity or linkage variables are standardized to have a mean of zero and standard deviation of one. In all cases, a positive value indicates a "closer" relationship between the industries.

Specifically, we estimate the following equation:

$$(8') \quad \ln\left(Y_{pijt}\right) = \beta_1 \ln\left(L_{pijt}\right) + \beta_2 \ln\left(K^{B}_{pijt}\right) + \beta_3 \ln\left(K^{E}_{pijt}\right) + \beta_4 \ln\left(M_{pijt}\right)$$

$$+ \delta\ 1(\text{Winner})_{pj} + \kappa\ 1(\tau >= 0))_{jt} + \theta_1\ (1(\text{Winner}) * 1(\tau >= 0))_{pjt}$$

$$+ \pi_1\ 1(\text{Winner})_{pj} * \text{Proximity}_{ij} + \pi_2\ (1(\tau >= 0)_{jt} * \text{Proximity}_{ij})$$

$$+ \pi_3\ (1(\text{Winner}) * 1(\tau >= 0)_{pjt} * \text{Proximity}_{ij}) + \alpha_p + \mu_{it} + \lambda_j + \varepsilon_{pijt}$$

where Proximity_{ij} is a measure of economic proximity between the incumbent plant industry and the MDP industry. This equation is simply an augmented version of Model 1 that adds

---

[39] Within the same 2-digit SIC, 71% of incumbents in winning counties and 69% of incumbents in losing counties were still in the sample 5 years after the opening. Additionally, attriting plants within the same 2-digit SIC were also on statistically indistinguishable trends prior to the MDP opening. Thus, differential attrition seems unlikely to explain the 2-digit results.

interactions of the industry linkage variables with $1(\text{Winner})_{pj}$, $1(\tau >= 0))_{jt}$, and $(1(\text{Winner}) * 1(\tau >= 0))_{pjt}$. The coefficient of interest is π_3 , which is the coefficient on the triple interaction between the dummy for winner, the dummy for "after," and the measure of proximity. This coefficient assesses whether "closer" industries benefit more from the MDP's opening. A positive coefficient means that the estimated productivity spillover is larger after the MDP opening for incumbents that are geographically and economically close to the new plant, relative to incumbents that are geographically close but economically distant from the new plant (relative to the same comparison among incumbents in loser counties). A zero coefficient means that the estimated productivity spillover is the same for all the incumbents in a county, regardless of their economic proximity to the new plant.

Table 9 reports estimates of π_3. The first 6 columns include the interactions in one at a time. For example, column (1) suggests that a one standard deviation increase in the *CPS* Worker Transitions variable between incumbent plants' industry and the MDP's industry is associated with a 7% increase in the spillover. This finding is consistent with the theory that spillovers occur through the flow of workers across firms. One possibility is that new workers share ideas on how to organize production or information on new technologies that they learned with their previous employer. This measure tends to be especially high within 2-digit industries, so this finding was foreshadowed by the own 2-digit results in Table 8.

The measures of intellectual or technological linkages indicate meaningful increases in the spillover. The precise mechanism by which these ideas are shared is unclear, although both the flow of workers across firms and the mythical exchange of ideas over beers between workers from different firms are possibilities. Notably, there is more variation in these measures within 2-digit industries than in the *CPS* labor transitions measure.

Columns (5) and (6) provide little support for the flow of goods and services in determining the magnitude of spillovers. Thus, the data fail to support the types of stories where an auto manufacturer encourages (or even forces) its suppliers to adopt more efficient production techniques. Recall, all plants owned by the MDP's firm are dropped from the analysis, so this finding does not rule out this channel within firms. The finding on the importance of labor flows is consistent with the results in Ellison, Glaeser and Kerr (2007) and Dumais, Ellison, and Glaeser (2002), while the finding on input and output flows stands in contrast with these papers' findings.

In the column (7) specification, we include all the interactions simultaneously. The labor flow, the citation pattern, and the technology input interactions all remain positive but now would be judged to be statistically insignificant. The interactions with proximity to customers and suppliers are now both negative.

Overall, this analysis provides some support for the notion that spillovers occur between firms that share workers and between firms that use similar technologies. In terms of Section IC, this evidence is consistent with intellectual externalities, to the extent that they are embodied in workers who move from firm to firm, and to the extent that they occur among firms that use technologies that are reasonably similar. Table 9 seems less consistent with the hypothesis that agglomeration occurs because of proximity to customers and suppliers. We caution against definitive conclusions, because the utilized measures are all imperfect proxies for the potential channels. Further, the possibility of better matches between workers and firms could not be directly tested with these data.

D. Entry and Labor Costs as Indirect Tests of Spillovers

The paper has uncovered economically sizable productivity gains for incumbent establishments following the opening of the new MDP. For example, a MDP plant opening is associated with a 12% increase in TFP five years later. This effect is even larger for incumbent plants that are in the same 2-digit industry of the new plant and for plants that tend to share workers and technologies. In the presence of positive spillovers, the model has two empirical predictions which this subsection tests.

First, if the spillovers are of a sufficient magnitude (i.e., they are larger than the increase in costs in the short run), the MDP's county should experience entry by new firms (relative to the losing counties). Table 10 tests this prediction. The entries in Panel 1 come from regressions that use data from the Census of Manufactures, which is conducted every five years. The dependent variables are the log of the number of establishments (column 1) and the log of total manufacturing output (column 2) in the county, respectively. In both columns, all plants owned by the MDP's firm are excluded from the dependent variable. The sample is comprised of observations from winning and losing counties only. The covariates include a full set of county fixed effects, year fixed effects, case fixed effects, and an indicator for whether the observation is from after the MDP's opening. The parameter of interest is associated with the interaction of indicators for an observation from a winning county and the post-opening indicator, so it is a difference in differences estimator of the impact of the MDP's opening.[40]

Column (1) reports that the number of manufacturing plants increased by roughly 12.5% in winning counties after the MDP plant's opening. A limitation of this measure is that it assumes that all plants are of an equal size. The total value of output is economically more meaningful, because it treats an increase in output at an existing plant and a new plant equally. As column (2) highlights, the opening of a MDP plant is associated with a roughly 14.5% increase in total output in the manufacturing sector although this is not estimated precisely.

Overall, these results are consistent with the TFP results of substantial spillovers in that it appears that the MDP attracted new economic activity to the winning counties (relative to losing ones) in the manufacturing sector.[41] Presumably, this new activity located in the winning counties to gain access to the spillovers.

The second prediction is that if the spillovers are positive, the prices of local inputs will increase as firms compete for these factors of production. The most important locally supplied input for manufacturing plants is labor. Column (3) in Panel 2 of Table 10 reports the results from regressions of the log wage using data from the 1970, 1980, 1990, and 2000 Censuses of Population from the winning and losing counties.[42] These data are preferable to the measure of labor costs reported in the Census of Manufacturers (i.e., the aggregate wage bill for production and non-production workers), which does not provide information on the quality of the labor force (e.g., education and experience). Specifically, we estimate models for ln(wage) and control for dummies for interactions of worker age and year, age-squared and year, education and year, sex and race and Hispanic and citizen, and case fixed effects. We also include indicators for whether the observation is from a winning county, occurs after the MDP's opening, and the

---

[40] Because data is available every 5 years, depending on the Census year relative to the MDP opening, the sample years are 1 – 5 years before the MDP opening and 4 – 8 years after the MDP opening. Thus, each MDP opening is associated with one earlier date and one later date. Models are weighted by the number of plants in the county in years -6 to -10 and column 4 is weighted by the county's total manufacturing output in years -6 to -10.

[41] It is possible that the MDP's spillovers extended beyond manufacturing. In this case, it might be reasonable to expect increased entry in other sectors too.

[42] The sample is limited to individuals who worked last year, worked more than 26 weeks, usually work more than 20 hours per week, are not in school, are at work, and who work for wages in the private sector. One important limitation of the Census data is that they lack exact county identifiers for counties with populations below 100,000. Instead, it is possible to identify PUMAs in the Census, which in rural areas can include several counties. This introduces significant measurement error, which is partly responsible for the imprecision of the estimate.

interaction of these two indicators.[43] This interaction is the focus of the regression and is an adjusted difference in differences estimator of the impact of the MDP's opening on wages. This equation is analogous to the Model 1 version of equation (8) that was used to analyze TFP.

The estimate indicates that after adjusting for observable heterogeneity, wages increase by 2.7% in winning counties after the MDP's opening. This effect appears quantitatively sizable and is marginally statistically significant. The multiplication of the estimated 2.7% wage increase by the average labor earnings in winning counties implies that the quality-adjusted annual wage bill for employers in all industries increased by roughly $151 million after the MDP's opening. This finding is consistent with positive spillovers and an upward sloping labor supply curve, perhaps due to imperfect mobility of labor (as in Section I).

It is possible to use the estimated increase in wages to make some back of the envelope calculations of the MDP's impact on incumbent plants' profits. Recall, the Model 1 result in Table 5 indicated an increase in TFP of approximately 4.8% (we focus on Model 1 because it is impossible to estimate a version of Model 2 with the decennial population Census data). If we assume that workers are homogenous or that high and low skill workers are perfectly substitutable in production, then the labor market-wide increase in wages applies throughout the manufacturing sector. In our sample, labor accounts for roughly 23% of total costs, so the estimated 2.7% increase in skill adjusted wages implies that manufacturers' costs increased by approximately 0.62%. The increased production costs due to higher wages are therefore 13% of the gain in TFP.

These calculations demonstrate that the gains in TFP do not translate directly to profits due to the higher costs of local inputs. Since the prices and quality of other inputs are not observable, it is not possible to determine the total increase in production costs. Further, the observations on wages occurs only a few years after the MDP's opening and the impact on wages may increase more as production expands to gain access to the spillover. For these reasons, this back of the envelope calculation should be interpreted as a lower bound of the increase in input costs. In the long run, an equilibrium requires that the total impact on profits is zero.

---

[43] The pre-period is defined as the most recent census before the MDP opening. The post-period is defined as the most recent census 3 or more years after the MDP opening. Thus, the sample years are 1 – 10 years before the MDP opening and 3 – 12 years after the MDP opening.

VI. Summary and Interpretation

This paper makes three principal contributions. This section summarizes them and places them in some context.

The first and most robust result is that the successful attraction of a "Million Dollar Plant" is on average associated with increased TFP for incumbent plants. This finding is robust to a battery of specification tests. In the preferred Model 2, the estimates suggest that incumbents' TFP in winning counties was about 12% higher. This translates into an additional $430 million in annual output five years after the MDP's opening. This is an economically large number and it is natural to wonder whether it is "too large" to be plausible.

There are several related issues worth noting when considering this possibility. First, there may be an unobserved factor correlated with the MDP's opening that can explain these results and would invalidate the identifying assumption. The likelihood of this possibility is diminished by the similarity of the pre-trends in TFP documented in Figure 1 and the balancing of many of the ex-ante observable characteristics of winning and losing counties and their incumbent plants. Nevertheless, this possibility cannot be dismissed as would be the case in a randomized experiment.

Second, it is possible that the estimated impacts on TFP are influenced by changes in the quality of the workforce employed by incumbent plants. The sign of this change is *a priori* unclear. On the one hand, the MDP may attract higher quality workers to the county, allowing the incumbents to upgrade the quality of their workforce. On the other hand, the MDP could hire away the incumbents' best workers. Regardless of which force prevails, this could affect the estimated impact on TFP because labor is measured as total hours by production workers (recall that education or other measures of skill are not included in the *ASM*).

Third, the external validity of the results is unknown. In particular, the 47 MDP plants in the sample differ from the average manufacturing plant in many respects. Perhaps most importantly, these plants generated bidding from local governments, which is a first indication that there may be an ex-ante expectation of significant spillovers. Further, these plants are substantially larger than the typical manufacturing plant. The point is that these results are unlikely to generalize to the opening of more typical plants.

Fourth, the estimated impact on TFP is unlikely to be a structural estimate of the MDP's impact on the TFP of incumbents. As Table 10 indicated, the MDP's opening appears to lead to other plants opening in the same county. Thus, the estimated spillover may reflect spillovers from the MDP and these new plants. Consequently, it is likely more appropriate to interpret the TFP effect as the reduced form effect of the MDP's opening and everything that accompanies it rather than the impact of the MDP alone.

The second contribution is a theoretical one that is supported by the data. In particular, the model underscored that increases in TFP do not translate directly into higher profits available for other manufacturing plants. The model predicts that the increases in TFP will be accompanied by increases in local input prices and that these increases are necessary for an equilibrium. The finding of higher prices for quality-adjusted labor in Table 10 is consistent with this prediction. Further, the increased levels of economic activity also documented in Table 10 reflect the increased demand to locate in the winning county that leads to higher local prices and the new equilibrium.

The third contribution is that the paper has shed some light on the channels that underlie the estimated spillovers. Specifically, our tentative conclusion is that the spillovers are larger between firms that share workers and use similar technologies. This is consistent with intellectual externalities, to the extent that they are embodied in workers who move from firm to firm and occur among firms that use technologies that are reasonably similar. Additionally, it is consistent with higher rates of TFP due to improved efficiencies of worker-firm matches.

Finally, these results may have some surprising policy implications. A standard critique of local governments providing subsidies to new plants to locate within their jurisdictions is that it may be rational for the locality but it is welfare decreasing for the nation. Although the economics of this argument are not always transparent, we believe that it generally refers to cases of disequilibrium where local factors of production (e.g., labor or buildings) are unemployed. In this case, tax competitions may be beneficial locally but suboptimal nationally (at least in the absence of significant worker moving costs).

In contrast, the finding of spillovers is the basis of an economic justification for local subsidies, even from a national perspective. Specifically, the MDP plant cannot capture these spillovers on its own and consequently might choose a location where its costs are low but the spillovers are minimal. However, the socially efficient outcome is for these plants to locate where the sum of their profits and the spillovers are greatest. In this setting, national welfare is maximized when payments are made to plants that produce the spillovers so that they internalize this externality in making their location decision.

In thinking about the policy implications, it is important to bear in mind that the estimated 12% gain in TFP is an average effect. As Figure 2 demonstrated, there is substantial variability. For example, the estimated impact is negative in 40% of the cases. The point is that risk averse local governments may be unwilling to provide tax incentives with this distribution of outcomes.

VII. Conclusion

Overall, this paper has documented that there are substantial spillovers flowing from large new plants to incumbent plants. Further, these spillovers are larger between plants that share labor pools and use similar technologies. Thus, this paper has provided evidence consistent with the idea that firms agglomerate in certain localities, at least in part, because they are more productive for being close to other firms.

There are several implications for future research. First, the paper has demonstrated the value of quasi-experiments that plausibly avoid the confounding of spillovers with differences in the determinants of TFP across locations. Second, the paper highlights that tests for the presence of spillovers can be conducted by directly measuring TFP. These tests can serve as an important complement to measurement of rates of coagglomeration that may reflect spillovers, cost shifters, or natural advantages. In this spirit, it is important to determine whether impacts on TFP are evident outside the manufacturing sector. Third, the heterogeneity in the estimated spillovers documented in Figure 2 and the results on the mechanism in Table 9 underscore that there is still much to learn about the structural source of these spillovers. This is an especially fruitful area for future research.

References

Acemoglu, Daron. 1996. "A Microfoundation for Social Increasing returns in Human Capital Accumulation." *Quarterly Journal of Economics*, 11 (August): 779-804.

Audretsch, David and Maryann Feldman. 1996. "R&D Spillovers and the Geography of Innovation and Production." *American Economic Review* 86 (June): 630-640.

Audretsch, David and Maryann Feldman. 2004. "Knowledge Spillovers and the Geography of Innovation." In *Handbook of Urban and Regional Economics* 4, edited by J. Vernon Henderson and Jacques-François Thisse. Amsterdam: North-Holland.

Becker, Randy, John Haltiwanger, Ron Jarmin, Shawn Klimek, and Daniel Wilson. 2005. "Micro and Macro Data Integration: The Case of Capital." *Bureau of the Census: Center for Economic Studies Working Paper 05-02.*

Bloom, Nick, Mark Schankerman, and John Van Reenen. 2007. "Identifying Technology Spillovers and Product Market Rivalry." National Bureau of Economic Research Working Paper No. 13060.

Branstetter, Lee G. 2001. "Are Knowledge Spillovers International or Intranational in Scope? Microeconometric evidence from US and Japan." *Journal of International Economics* 53 (February): 53-79.

Card, David, Kevin F. Hallock, and Enrico Moretti. 2007. "The Geography of Giving: The Effect of Corporate Headquarters on Local Charities." University of California, Berkeley Mimeograph.

Chiang, Hyowook. 2004. "Learning By Doing, Worker Turnover, and Productivity Dynamics." Econometric Society 2004 Far Eastern Meeting Paper No. 593.

Davis, James and J. Vernon Henderson. 2004. "The Agglomeration of Headquarters." *Bureau of the Census: Center for Economic Studies Working Paper 04-02.*

Davis, Steven J., John C. Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction.* Cambridge, MA: MIT Press.

Davis, Donald and David E. Weinstein. 2002. "Bones, Bombs, and Break Points: The Geography of Economic Activity." *American Economic Review* 92 (December).

Dumais, Guy, Ellison, Glenn, and Glaeser, Edward L. 2002. "Geographic Concentration as a Dynamic Process." *Review of Economics and Statistics* 84 (May): 193-204.

Duranton, Giles and Diego Puga. 2004 "Micro-Foundations of Urban Agglomeration Economies." In *Handbook of Urban and Regional Economics* 4, edited by J. Vernon Henderson and Jacques-François Thisse. Amsterdam: North-Holland.

Ellison, Glenn and Edward L. Glaeser. 1997. "Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach." *Journal of Political Economy* 105 (October): 889-927.

Ellison, Glenn and Edward L. Glaeser. 1999. "The Geographic Concentration of Industry: Does Natural Advantage Explain Agglomeration?" *American Economic Review* 89 (May): 311-316.

Ellison, Glenn, Edward L. Glaeser, and William Kerr. 2007. "What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns." National Bureau of Economic Research Working Paper No. 13068.

Foster, Lucia, John C. Haltiwanger, and Chad Syverson. 2007. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" University of Chicago Mimeograph..

Foster, Lucia, John C. Haltiwanger, and C.J. Krizan. 2000. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." Mimeograph.

Glaeser, Edward L. 1999. "Learning in Cities." *Journal of Urban Economics* 46(September): 254-77.

Glaeser, Edward L. 2001. "The Economics of Location-Based Tax Incentives." *Harvard Institute of Economic Research,* Discussion Paper Number 1932.

Glaeser, Edward L. and Janet E. Kohlhase. 2003. "Cities, Regions and the Decline of Transport Costs." *Papers in Regional Science* 83 (January): 197-228.

Glaeser, Edward L. and David C. Mare. 2001. "Cities and Skill." *Journal of Labor Economics* 19(2): 316-342.

Greenstone, Michael and Enrico Moretti. 2004. "Bidding for Industrial Plants: Does Winning a 'Million Dollar Plant' Increase Welfare?" National Bureau of Economic Research Working Paper No. 9844.

Griliches, Zvi. 1958. "Research Cost and Social Returns: Hybrid Corn and Related Innovations." *Journal of Political Economy* 66 (October): 419-431.

Griliches, Zvi and Jacques Mairesse. 1995. "Production Functions: the Search for Identification." National Bureau of Economic Research Working Paper No. 5067.

Grossman, Gene M. and Elhanan Helpman. 1991. *Innovation and Growth in the Global Economy.* Cambridge, MA: MIT Press.

Henderson, J. Vernon. 2001. "Urban Scale Economies." In *Handbook of Urban Studies*, edited by Ronan Paddison. London: SAGE Publications.

Henderson, J. Vernon. 2003. "Marshall's Scale Economies." *Journal of Urban Economics* 55 (January): 1-28.

Henderson, J. Vernon and Duncan Black. 1999. "A Theory of Urban Growth." *Journal of Political Economy*, 107 (April): 252-284

Jaffe, Adam B. 1986. "Technological Opportunities and Spillovers of R&D: Evidence from Firm's Patents, Profit and Market Value." *American Economic Review* 76 (December): 984-1001.

Jaffe, Adam B. and Manuel Trajtenberg. 2002. *Patents, Citations and Innovations: A Window on the Knowledge Economy.* Cambridge, MA: MIT Press.

Jaffe, Adam B., Manuel Trajtenberg, and Michael S. Fogarty. 2000. "Knowledge Spillovers and Patent Citations: Evidence from a Survey of Inventors." *American Economic Review* 90 (May): 215-218.

Jaffe, Adam B., Manuel Trajtenberg and Rebecca M. Henderson. 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citation." *Quarterly Journal of Economics* 108 (August): 577- 598.

Jovanovic, Boyan and Rafael Rob. 1989. "The Growth and Diffusion of Knowledge." *Review of Economic Studies* 56 (October): 569-582.

Krugman, Paul. 1991a. *Geography and Trade.* Cambridge, MA: MIT Press.

Krugman, Paul. 1991b. "Increasing Returns and Economic Geography," *Journal of Political Economy* 99 (June): 483-499.

Levinsohn, James and Amil Petrin. 2003. "Estimating Production Functions Using Inputs to Control for Unobservables." *Review of Economic Studies* 70 (April): 317-342.

Lucas, Robert E. Jr. 1988. "On the Mechanics of Economic Development." *Journal of Monetary Economics,* 22 (July): 3-42.

Marshall, Alfred. 1890. *Principles of Economics.* New York: Macmillan and Co.

McGahan, Anita W. and Brian S. Silverman. 2001. "How Does Innovative Activity Change as Industries Mature?" *International Journal of Industrial Organization* 19 (July): 1141-1160.

Moretti, Enrico. 2004a. "Estimating the external return to higher education: Evidence from cross-sectional and longitudinal data." *Journal of Econometrics* 120 (July-August): 175-212.

Moretti, Enrico. 2004b. "Workers' Education, Spillovers and Productivity: Evidence from Plant-Level Production Functions." *American Economic Review* 94 (June): 656-690.

Moretti, Enrico. 2004c. "Human Capital Externalities in Cities." In *Handbook of Urban and Regional Economics* 4, edited by J. Vernon Henderson and Jacques-François Thisse. Amsterdam: North-Holland.

Mowery, David C. and Arvids A. Ziedonis. 2001. "The Geographic Reach of Market and Non-market Channels of Technology Transfer: Comparing Citations and Licenses of University Patents." National Bureau of Economic Research Working Paper No. 8568.

Olley, Steven G. and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (November): 1263-98.

Ottaviano, Gianmarco and Jacques-François Thisse. 2004. "Agglomeration and Economic Geography." In *Handbook of Urban and Regional Economics* 4, edited by J. Vernon Henderson and Jacques-François Thisse. Amsterdam: North-Holland.

Petrongolo, Barbara and Christopher A. Pissarides. 2005. "Scale Effects in Markets with Search." *The*

Economic Journal, 116 (January): 21-44.

Rauch, James E. 1993. "Productivity Gains from Geographic Concentration of Human Capital: Evidence from the Cities." *Journal of Urban Economics* 34 (November): 380-400.

Roback, Jennifer. 1982. "Wages, Rents and the Quality of Life." *Journal of Political Economy* 90 (December): 1257-1278.

Roback, Jennifer. 1988. "Wages, rents and amenities: Differences among workers and regions." *Economic Inquiry* 26 (January): 23-41.

Rosenbaum, Paul R. and Donald B. Rubin. 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrika* 70 (April): 41-55.

Rosenthal, Stuart S. and William C. Strange. 2001. "The Determinants of Agglomeration." *Journal of Urban Economics* 50 (September): 191-229.

Rosenthal, Stuart S. and William C. Strange. 2004. "Evidence on the Nature and Sources of Agglomeration." In *Handbook of Urban and Regional Economics* 4, edited by J. Vernon Henderson and Jacques-François Thisse. Amsterdam: North-Holland.

Saxenian, Anna Lee. 1994. *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*. Cambridge, MA: Harvard University Press.

Syverson, Chad. 2004a. "Market Structure and Productivity: A Concrete Example." *Journal of Political Economy* 6 (December): 1181-1222.

Syverson, Chad. 2004b. "Product Substitutability and Productivity Dispersion." *Review of Economics and Statistics* 2 (May): 534-550.

van Biesebroeck, Johannes. 2004. "Robustness of Productivity Estimates." National Bureau of Economic Research Working Paper No. 10303.

Venable, Tim. 1992. "BMW Drives into South Carolina with $300 Million Auto Plant." *Site Selection* August: 630-631.

Weiss, Leonard W. 1972. "The Geographic Size of Markets in Manufacturing." *The Review of Economics and Statistics* 54 (August): 245-257.

Zucker, Lynne G., Michael R. Darby, and Marilynn B. Brewer. 1998. "Intellectual Human Capital and the Birth of U.S. Biotechnology Enterprises." *American Economic Review* 88 (March): 290-306.

Figure 1: The Effect of a "Million Dollar Plant" Opening on TFP of All Manufacturing Plants in Winner and Loser Counties.

All Industries: Winners Vs. Losers

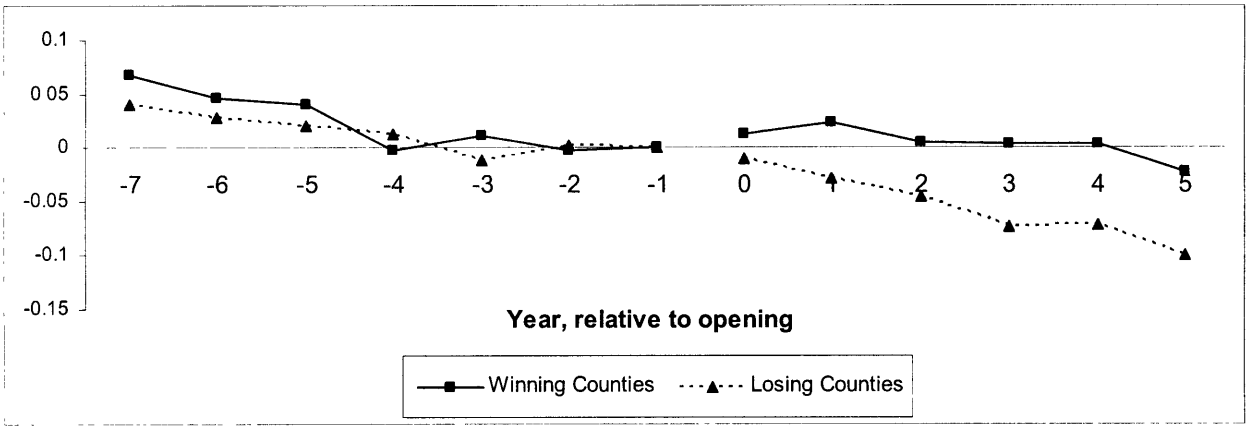

Difference: Winners – Losers

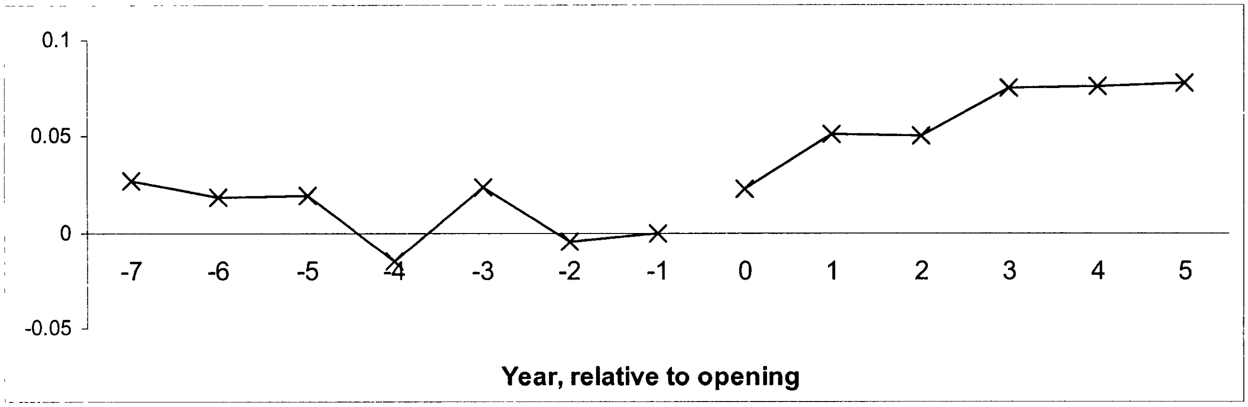

Notes: These figures accompany Table 4.

Figure 2. Distribution of Case-Specific Mean Shift Effects from the Opening of a "Million Dollar Plant"

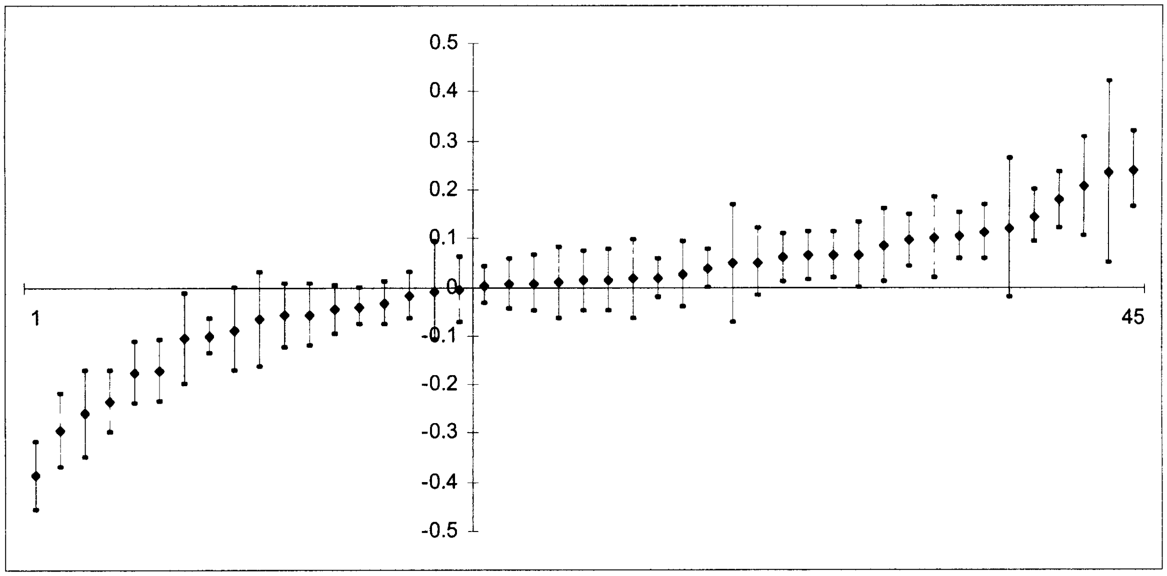

Notes: The figure reports results from a version of Model 1 that estimates the parameter θ_1 for each of the 47 MDP cases. The figure reports only 45 estimates because two cases were dropped for Census confidentiality reasons.

Figure 3: The Effect of a "Million Dollar Plant" Opening on TFP of Manufacturing Plants in the MDP's 2-Digit Industry in Winner and Loser Counties.

2-digit MDP Industry: Winners Vs. Losers

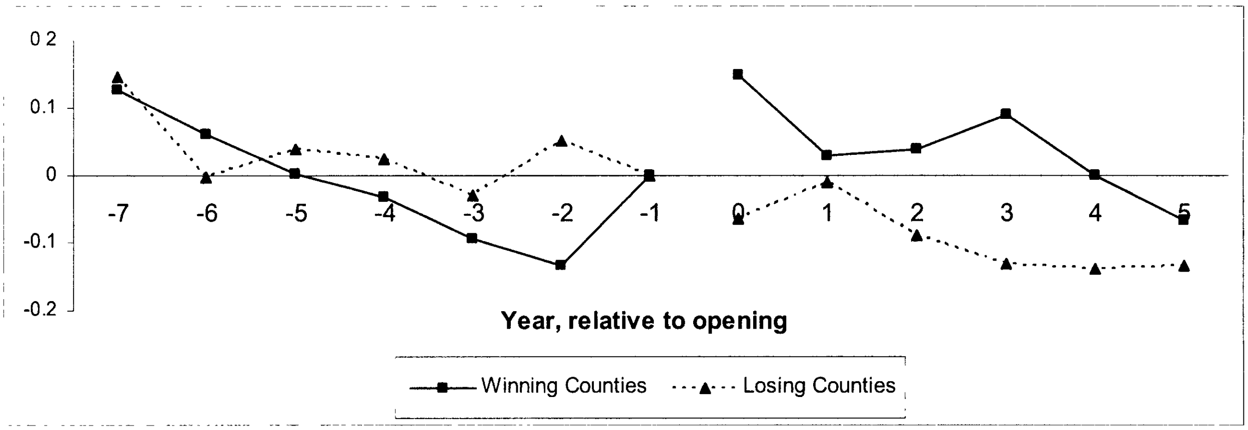

Difference (Winners – Losers)

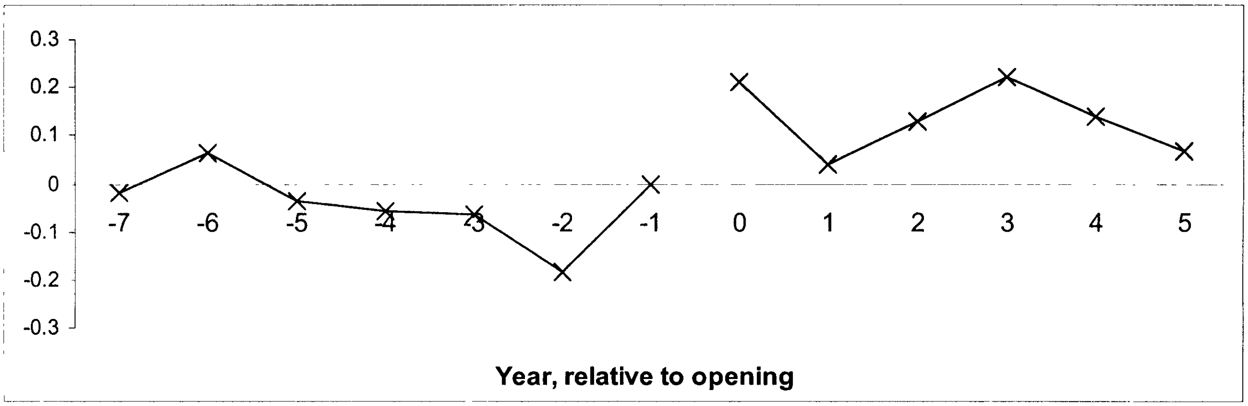

Notes: These figures accompany Table 8, Column 2 (MDP's 2-digit Industry).

Figure 4: The Effect of a "Million Dollar Plant" Opening on TFP of Manufacturing Plants in All Industries, Except the MDP's 2-Digit Industry in Winner and Loser Counties.

Other Industries: Winners Vs. Losers

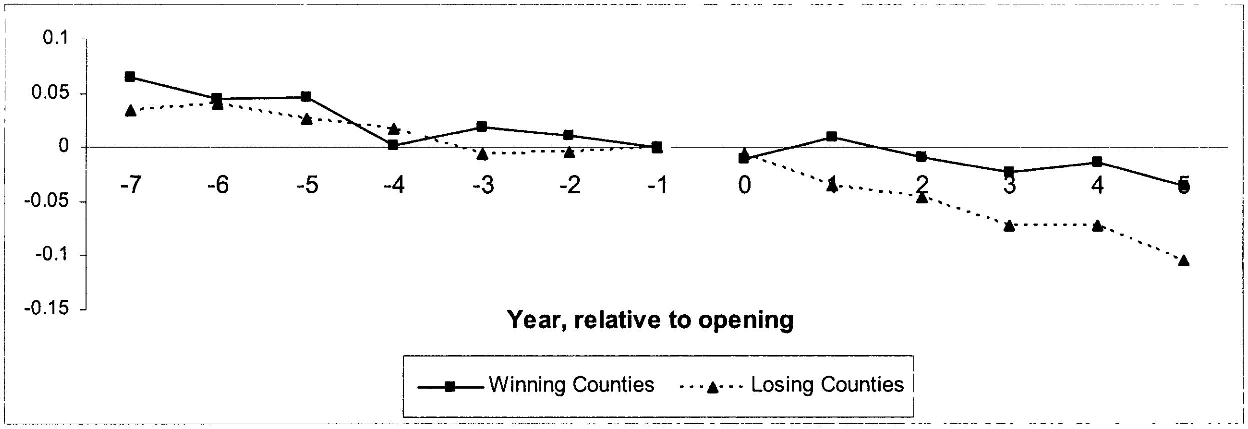

Difference (Winners – Losers)

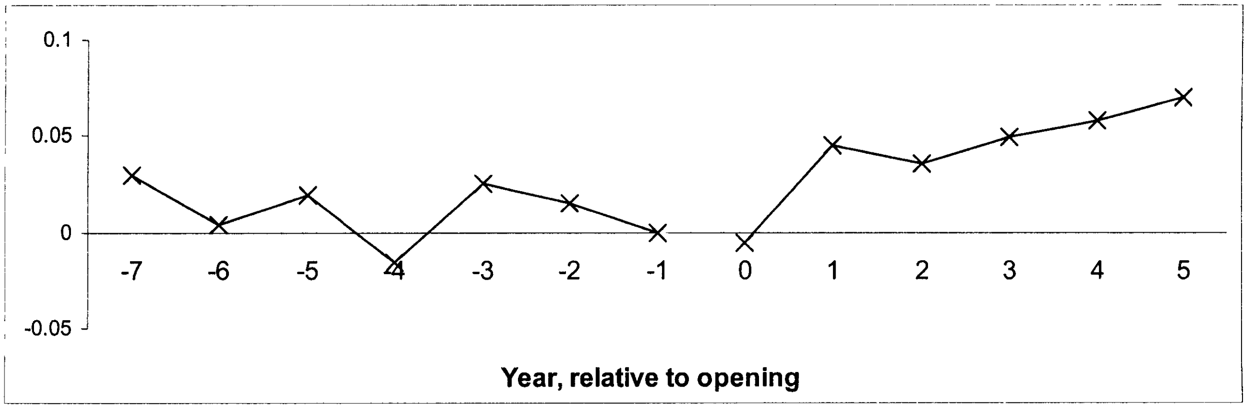

Notes: These figures accompany Table 8, Column 3 (All 2-digit Industries, except the MDP's 2-digit Industry).

Table 1. The "Million Dollar Plant" Sample

	(1)
Sample MDP Openings[1]:	
Across All Industries	47
Within Same 2-digit SIC	16
Across All Industries:	
Number of Loser Counties per Winner County:	
1	31
2+	16
Reported Year – Matched Year: [2]	
-2 to -1	20
0	15
1 to 3	12
Reported Year of MDP Location:	
1981 – 1985	11
1986 – 1989	18
1990 – 1993	18
MDP Characteristics, 5 years after opening:[3]	
Output ($1000)	452801
	(901690)
Output, relative to county output 1 year prior	0.086
	(0.109)
Hours of Labor (1000)	2986
	(6789)

[1] Million Dollar Plant openings that were matched to the Census data and for which there were incumbent plants in both winning and losing counties that are observed in each of the eight years prior to the opening date (the opening date is defined as the earliest of the magazine reported year and the year observed in the SSEL.) This sample is then restricted to include matches for which there were incumbent plants in the Million Dollar Plant's 2-digit SIC in both locations.

[2] Only a few of these differences are 3. Census confidentiality rules prevent being more specific.

[3] Of the original 47 cases, these statistics represent 28 cases. A few very large outlier plants were dropped so that the mean would be more representative of the entire distribution (those dropped had output greater than half of their county's previous output and sometimes much more). Of the remaining cases: most SSEL matches were found in the *ASM* or *CM* but not exactly 5 years after the opening date; a couple of SSEL matches in the 2xxx-3xxx SICs were never found in the *ASM* or *CM*; and a couple of SSEL matches not found were in the 4xxx SICs. The MDP characteristics are similar for cases identifying the effect within same 2-digit SIC. Standard deviations are reported in parentheses. All monetary amounts are in 2006 US dollars.

Table 2. Summary Statistics for Measures of Industry Linkages

Measure of Industry Linkage	Description	Mean All Plants	Mean Only 1st Quartile	Mean Only 4th Quartile	Standard Deviation
Labor Market Pooling:					
CPS Worker Transitions	Proportion of workers leaving a job in this industry that move to the MDP industry (15 months later)	0.119	0.002	0.317	0.249
Intellectual or Technology Spillovers:					
Citation pattern	Percentage of manufactured industry patents that cite patents manufactured in MDP industry	0.022	0.001	0.057	0.033
Technology Input	R&D flows from MDP industry, as a percentage of all private sector technological expenditures	0.022	0.000	0.106	0.084
Technology Output	R&D flows to MDP industry, as a percentage of all original research expenditures	0.011	0.000	0.042	0.035
Proximity to Customers and Suppliers:					
Manufacturing Input	Industry inputs from MDP industry, as a percentage of its manufacturing inputs	0.017	0.000	0.075	0.061
Manufacturing Output	Industry output used by MDP industry, as a percentage of its output to manufacturers	0.042	0.000	0.163	0.139

Notes: *CPS* Worker Transitions was calculated from the frequency of worker industry movements in the rotating *CPS* survey groups. This variation is by Census Industry codes, matched to 2-digit SIC. The last 5 measures of cross-industry relationships were provided by Ellison, Glaeser, and Kerr (NBER Working Paper 13068). These measures are defined in a 3-digit SIC by 3-digit SIC matrix, though much of the variation is at the 2-digit level. In all cases, more positive values indicate a closer relationship between industries. Column 1 reports the mean value of the measure for all incumbent plants matched to their respective MDP. Column 2 reports the mean for the lowest 25% and column 3 reports the mean for the highest 25%. Column 4 reports the standard deviation across all observations. The sample of plants is all incumbent plants, as described for Table 1, for which each industry linkage measure is available for the incumbent plant and its associated MDP. These statistics are calculated when weighting by the incumbent plant's total value of shipments eight years prior to the MDP opening.

Table 3. County & Plant Characteristics by Winner Status, One Year Prior to a "Million Dollar Plant" Opening

	All Plants					Within Same Industry (2-digit SIC)				
	Winning Counties (1)	Losing Counties (2)	All US Counties (3)	(1) – (2) t-stat (4)	(1) – (3) t-stat (5)	Winning Counties (6)	Losing Counties (7)	All US Counties (8)	(6) – (7) t-stat (9)	(6) – (8) t-stat (10)
# of Counties	47	73				16	19			
County Characteristics:										
Total Per-capita Earnings ($)	17,418	20,628	11,259	**-2.05**	**5.79**	20,230	20,528	11,378	-0.11	**4.62**
% Change, over last 6 years	0.074	0.096	0.037	-0.81	1.67	0.076	0.089	0.057	-0.28	0.57
Population	322,745	447,876	82,381	-1.61	**4.33**	357,955	504,342	83,430	-1.17	**3.26**
% Change, over last 6 years	0.102	0.051	0.036	**2.06**	**3.22**	0.070	0.032	0.031	1.18	1.63
Employment-Population ratio	0.535	0.579	0.461	-1.41	**3.49**	0.602	0.569	0.467	0.64	**3.63**
Change, over last 6 years	0.041	0.047	0.023	-0.68	**2.54**	0.045	0.038	0.028	0.39	1.57
Manufacturing Labor Share	0.314	0.251	0.252	**2.35**	**3.12**	0.296	0.227	0.251	1.60	1.17
Change, over last 6 years	-0.014	-0.031	-0.008	1.52	-0.64	-0.030	-0.040	-0.007	0.87	**-3.17**
# of Sample Plants	18.8	25.6	7.98	-1.35	**3.02**	2.75	3.92	2.38	-1.14	0.70
Plant Characteristics:										
Output ($1000)	190,039	181,454	123,187	0.25	**2.14**	217,950	178,958	132,571	0.41	1.25
% Change, over last 6 years	0.082	0.082	0.118	0.01	-0.97	-0.061	0.177	0.182	-1.23	**-3.38**
Hours of labor (1000s)	1,508	1,168	877	1.52	**2.43**	1,738	1,198	1,050	0.92	1.33
% Change, over last 6 years	0.122	0.081	0.115	0.81	0.14	0.160	0.023	0.144	0.85	0.13

Notes: For each case to be weighted equally, counties are weighted by the inverse of their number per-case. Similarly, plants are weighted by the inverse of their number per-county multiplied by the inverse of the number of counties per-case. The sample includes all plants reporting data in the *ASM* for each year between the MDP opening and eight years prior. Excluded are all plants owned by the firm opening a MDP. Also excluded are all plants from two uncommon 2-digit SIC values so that subsequently estimated clustered variance matrices would always be positive definite. The sample of all United States counties excludes winning counties and counties with no manufacturing plant reporting data in the *ASM* for nine consecutive years. These other United States counties are given equal weight within years and are weighted across years to represent the years of MDP openings. Reported t-statistics are calculated from standard errors clustered at the county level. All monetary amounts are in 2006 US dollars.

Table 4. Incumbent Plant Productivity, Relative to the Year of a MDP Opening

Event Year	In Winning Counties (1)	In Losing Counties (2)	Difference (1) – (2) (3)
τ = -7	0.067	0.040	0.027
	(0.058)	(0.053)	(0.032)
τ = -6	0.047	0.028	0.018
	(0.044)	(0.046)	(0.023)
τ = -5	0.041	0.021	0.020
	(0.036)	(0.040)	(0.025)
τ = -4	-0.003	0.012	-0.015
	(0.030)	(0.030)	(0.024)
τ = -3	0.011	-0.013	0.024
	(0.022)	(0.022)	(0.021)
τ = -2	-0.003	0.001	-0.005
	(0.027)	(0.011)	(0.028)
τ = -1	0	0	0
τ = 0	0.013	-0.010	0.023
	(0.018)	(0.011)	(0.019)
τ = 1	0.023	-0.028	0.051*
	(0.026)	(0.024)	(0.023)
τ = 2	0.004	-0.046	0.050
	(0.036)	(0.046)	(0.033)
τ = 3	0.003	-0.073	0.076+
	(0.047)	(0.057)	(0.043)
τ = 4	0.004	-0.072	0.076*
	(0.053)	(0.062)	(0.033)
τ = 5	-0.023	-0.100	0.077*
	(0.069)	(0.067)	(0.035)
R-squared	0.9861		
Observations	28732		

Notes: Standard errors are clustered at the county level. Columns 1 and 2 report coefficients from the same regression: the natural log of output is regressed on the natural log of inputs (all worker hours, building capital, machinery capital, materials), year x 2-digit SIC fixed effects, plant fixed effects, case fixed effects, and the reported dummy variables for whether the plant is in a winning or losing county in each year relative to the MDP opening. When a plant is a winner or loser more than once, it receives a dummy variable for each incident. Plant-year observations are weighted by the plant's total value of shipments eight years prior to the MDP opening. Data on plants in all cases is only available 8 years prior to the MDP opening and 5 years after. Capital stocks were calculated using the permanent inventory method from early book values and subsequent investment. The sample of incumbent plants is the same as in columns 1 – 2 of Table 3. ** denotes significance at 1% level, * denotes significance at 5% level, + denotes significance at 10% level.

Table 5. The Effect of the Opening of a MDP Plant on the Productivity of Incumbent Plants

	All Counties MDP Winners – MDP Losers		MDP Counties MDP Winners – MDP Losers		All Counties Random Winners
	(1)	(2)	(3)	(4)	(5)
Model 1:					
Mean Shift	0.0442+	0.0435+	0.0524*	0.0477*	-0.0824**
	(0.0233)	(0.0235)	(0.0225)	(0.0231)	(0.0177)
				[$170m]	
R-squared	0.9811	0.9812	0.9812	0.9860	0.9828
Observations (plant x year)	418064	418064	50842	28732	426853
Model 2:					
Effect after 5 years	0.1301*	0.1324*	0.1355**	0.1203*	-0.0559+
	(0.0533)	(0.0529)	(0.0477)	(0.0517)	(0.0299)
				[$429m]	
Level Change	0.0277	0.0251	0.0255	0.0290	-0.0197
	(0.0241)	(0.0221)	(0.0186)	(0.0210)	(0.0312)
Trend Break	0.0171+	0.0179*	0.0183*	0.0152+	-0.0060
	(0.0091)	(0.0088)	(0.0078)	(0.0079)	(0.0072)
Pre-trend	-0.0057	-0.0058	-0.0048	-0.0044	-0.0057**
	(0.0046)	(0.0046)	(0.0046)	(0.0044)	(0.0029)
R-squared	0.9811	0.9812	0.9813	0.9861	0.9828
Observations (plant x year)	418064	418064	50842	28732	426853
Plant & Ind-Year FEs	YES	YES	YES	YES	YES
Case FEs	NO	YES	YES	YES	N/A
Years Included	All	All	All	$-7 \leq \tau \leq 5$	All

Notes: The table reports results from the fitting of several versions of equation (8). Specifically, entries are from a regression of the natural log of output on the natural log of inputs, year x 2-digit SIC fixed effects, plant fixed effects, and case fixed effects. In Model 1, two additional dummy variables are included for whether the plant is in a winning county 7 to 1 years before the MDP opening or 0 to 5 years after. The reported mean shift indicates the difference in these two coefficients, i.e., the average change in TFP following the opening. In Model 2, the same two dummy variables are included along with pre- and post-trend variables. The shift in level and trend are reported, along with the pre-trend and the total effect evaluated after 5 years. In columns (1), (2), and (5), the sample is composed of all manufacturing plants in the *ASM* that report data for 14 consecutive years, excluding all plants owned by the MDP firm. In these models, additional control variables are included for the event years outside the range from $\tau = -7$ through $\tau = 5$ (i.e., -20 to -8 and 6 to 17). Column (2) adds the case fixed effects that equal 1 during the period that τ ranges from -7 through 5. In columns (3) and (4), the sample is restricted to include only plants in counties that won or lost a MDP. This forces the industry-year fixed effects to be estimated solely from plants in these counties. Incumbent plants are now required to be in the data only when the MDP opens and all 8 years prior (not also for 14 consecutive years, though this does not change the results). For column (4), the sample is restricted further to include only plant-year observations within the period of interest (where τ ranges from -7 to 5). This forces the industry-year fixed effects to be estimated solely on plant by year observations that identify the parameters of interest. In column (5), a set of 47 plant openings in the entire country were randomly chosen from the *ASM* in the same years and industries as the MDP openings. For all regressions, plant-year observations are weighted by the plant's total value of shipments eight years prior to the opening. Plants not in a winning or losing county are weighted by their total value of shipments in that year. All plants from two uncommon 2-digit SIC values were excluded so that estimated clustered variance-covariance matrices would always be positive definite. In brackets is the value in 2006 US$ from the estimated increase in productivity: the percent increase is multiplied by the total value of output for the affected incumbent plants in the winning counties. Standard error clustered at the county level in parenthesis. ** denotes significance at 1% level, * denotes significance at 5% level, + denotes significance at 10% level.

Table 6 The Effect of the Opening of a MDP Plant on the Productivity of Incumbent Plants. Robustness to Different Specifications

	Baseline specification (1)	Translog functional form (2)	Input - industry interactions (3)	Input-winner, input-post (4)	Region - Year FE (5)	Region - Year - Industry FE (6)	Fixed Input Cost Shares: plant level (7)	Fixed Input Cost Shares: SIC-3 level (8)
Model 1: Mean Shift	0.0477*	0.0471*	0.0406+	0.0571*	0.0442+	0.0369+	0.0364	0.0325
	(0.0231)	(0.0226)	(0.0220)	(0.0245)	(0.0230)	(0.0215)	(0.0228)	(0.0241)
Model 2: After 5 years	0.1203*	0.1053+	0.0977*	0.1177*	0.1176*	0.0879*	0.0971	0.0938+
	(0.0517)	(0.0535)	(0.0487)	(0.0538)	(0.0520)	(0.0442)	(0.0656)	(0.0538)

Notes: The table reports results from the fitting of several versions of equation (8). Column 1 reports estimates from the same specification as in column 4 of Table 5. Column 2 uses a translog functional form. Column 3 allows the effect of each input to differ at the 2-digit SIC level. Column 4 allows for the effect of inputs to differ in winning/losing counties and before/after the MDP opening. Column 5 controls for region (9 census divisions) by year fixed effects. Column 6 includes census division by year by industry fixed effects. Column 7 reports estimates when fixing the coefficient on each input to be its average cost share for each plant over the sample period. Per-period capital costs were calculated from capital rental rates using BLS data. Column 8 fixes the coefficient to be the average across all plants in each 3-digit SIC. Standard error clustered at the county level in parenthesis. ** denotes significance at 1% level, * denotes significance at 5% level, + denotes significance at 10% level.

Table 7. The Effect of the Opening of a MDP Plant on the Productivity of Incumbent Plants. Olley-Pakes (1996) and Levinsohn-Petrin (2003) Specifications.

	Baseline specification (1)	Investment (2)	Investment, Capital (3)	Investment-Capital Interactions (4)	Materials (5)	Materials-Capital Interactions (6)	Investment, Materials, Capital (7)	Materials-Capital and Investment-Capital Interactions (8)
Model 1: Mean Shift	0.0477*	0.0460*	0.0431+	0.0399+	0.0451+	0.0399+	0.0410+	0.0397*
	(0.0231)	(0.0230)	(0.0226)	(0.0213)	(0.0230)	(0.0216)	(0.0222)	(0.0199)
Model 2: After 5 years	0.1203*	0.1149*	0.0971*	0.0989*	0.1153*	0.1004*	0.0919+	0.1004*
	(0.0517)	(0.0529)	(0.0490)	(0.0482)	(0.0526)	(0.0487)	(0.0493)	(0.0487)

Notes: The table reports results from the fitting of several versions of equation (8). Column (1) reports estimates from the same specification as in column 4 of Table 5. To this baseline specification, the column (2) specification adds 4th degree polynomial functions of log building investment and log machinery/equipment investment. The column (3) specification adds 4th degree polynomials in the two different capital stocks to the polynomials in investment in the column (2) equation. The column (4) specification adds all the "own" interactions between polynomials in current investment and capital (i.e., the building investment polynomial is interacted with the building capital stock polynomial, but is not also interacted with the machinery/equipment polynomials for stocks or investment). Column (5) adds a 4th degree function of log materials to the baseline specification. In the column (6) specification, a 4th degree polynomial in materials is fully interacted with 4th degree polynomials in building capital and machinery capital. Column (7) includes fourth-degree polynomials in log materials and in log investment, and log capital stock for both types of capital (not interacted). Finally, column (8) includes the controls from the columns (4) and (6) specifications. Standard error clustered at the county level in parenthesis. ** denotes significance at 1% level, * denotes significance at 5% level, + denotes significance at 10% level.

Table 8 The Effect of the Opening of a MDP Plant on the Productivity of Incumbent Plants, For Incumbent Plants in the MDP's 2-digit Industry and All Other Industries

	All Industries (1)	MDP's 2-digit Industry (2)	All Other 2-Digit Industries (3)
Model 1:			
Mean Shift	0.0477*	0.1700*	0.0326
	(0.0231)	(0.0743)	(0.0253)
	[$170m]	[$102m]	[$104m]
R-squared	0.9860	0.9861	
Observations	28732	28732	
Model 2:			
Effect after 5 years	0.1203*	0.3289	0.0889+
	(0.0517)	(0.2684)	(0.0504)
	[$429m]	[$197m]	[$283m]
Level Change	0.0290	0.2814**	0.0004
	(0.0210)	(0.0895)	(0.0171)
Trend Break	0.0152+	0.0079	0.0147+
	(0.0079)	(0.0344)	(0.0081)
Pre-trend	-0.0044	-0.0174	-0.0026
	(0.0044)	(0.0265)	(0.0036)
R-squared	0.9861	0.9862	
Observations	28732	28732	

Notes: The table reports results from the fitting of several versions of equation (8). For comparison, Column 1 reports the same results as column 4 of Table 5. Columns 2 and 3 report estimates from the same regression, which fully interacts the winner/loser and pre/post variables with indicators for whether the incumbent plant is in the same 2-digit industry as the MDP or a different industry. Standard error clustered at the county level in parenthesis. The numbers in brackets are the value (2006 US$) from the estimated increase in productivity: the percent increase is multiplied by the total value of output for the affected incumbent plants in the winning counties. ** denotes significance at 1% level, * denotes significance at 5% level, + denotes significance at 10% level.

Table 9. How the Effect of the Opening of a MDP Plant on the Productivity of Incumbent Plants Varies with Measures of Economic Distance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
CPS Worker Transitions	0.0701**						0.0374
	(0.0237)						(0.0260)
Citation pattern		0.0545**					0.0256
		(0.0192)					(0.0208)
Technology Input			0.0320+				0.0501
			(0.0173)				(0.0421)
Technology Output				0.0596**			0.0004
				(0.0216)			(0.0434)
Manufacturing Input					0.0060		-0.0473
					(0.0123)		(0.0289)
Manufacturing Output						0.0150	-0.0145
						(0.0196)	(0.0230)
R-squared	0.9852	0.9852	0.9851	0.9852	0.9851	0.9852	0.9853
Observations	23397	23397	23397	23397	23397	23397	23397

Notes: The table reports results from the fitting of several versions of equation (8). Building on the Model 1 specification in column 4 of Table 5, each column adds the reported interaction terms between winner/loser and pre/post status with the indicated measures of how an incumbent plant's industry is linked to its associated MDP's industry. For assigning this linkage measure, the incumbent plant's industry is held fixed at its industry the year prior to the MDP opening. Whenever a plant is a winner or loser more than once, it receives an additive dummy variable and interaction term for each occurrence. These industry linkage measures are defined and described in Table 2 and here the measures are normalized to have a mean of zero and a standard deviation of one. The sample of plants is that in column 4 of Table 5, but it is restricted to plants that have industry linkage data for each measure. Standard errors clustered at the county level are reported in parentheses. ** denotes significance at 1% level, * denotes significance at 5% level, + denotes significance at 10% level.

Table 10. The Effect of the Opening of a MDP Plant on Wages and Number of Plants in the County

	Panel 1 (Census of Manufacturers)		Panel 2 (Census of Population)
	Dep. Var.: Log(Plants) (1)	Dep. Var.: Log(Total Output) (2)	Dep. Var.: log(Wage) (3)
Difference-in-difference	0.1255*	0.1454	0.0268+
	(0.0550)	(0.0900)	(0.0139)
R-squared	0.9984	0.9931	0.3623
Observations	209	209	1057999

Notes: The table reports results from the fitting of three regressions. In Panel 1, the dependent variables are the log of number of establishments and the log of total manufacturing output in the county, respectively. Controls include case effects, county, and year fixed effects. Entries are the county-level diff in diff estimates for winning a MDP, based on data from the Census of Manufacturers. Because data is available every 5 years, depending on the Census year relative to the MDP opening, the sample years are 1 – 5 years before the MDP opening and 4 – 8 years after the MDP opening. Thus, each MDP opening is associated with one earlier date and one later date. The column (1) model is weighted by the number of plants in the county in years -6 to -10 and the column (2) model is weighted by the county's total manufacturing output in years -6 to -10. In Panel 2, the dependent variable is log wage and controls include dummies for age*year, age-squared*year, education*year, sex*race*Hispanic*citizen, and case fixed effects. The entry is the county-level diff in diff estimate for winning a MDP. The pre-period is defined as the most recent census before the MDP opening. The post-period is defined as the most recent census 3 or more years after the MDP opening. Thus, the sample years are 1 – 10 years before the MDP opening and 3 – 12 years after the MDP opening. The sample is limited to individuals who worked last year, worked more than 26 weeks, usually work more than 20 hours per week, are not in school, are at work, and who work for wages in the private sector. The number of observations reported refers to unique individuals – some IPUMS county groups include more than one FIPS, so all individuals in a county group were matched to each potential FIPS. The same individual may then appear in more than one FIPS and observations are weighted to give each unique individual the same weight (i.e., an individual appearing twice receives a weight of ½).

Standard errors clustered at the county level are reported in parentheses. ** denotes significance at 1% level, * denotes significance at 5% level, + denotes significance at 10% level.

Appendix Table 1. The Effect of the Opening of a MDP Plant on Inputs and Output (Unadjusted for Inputs) of Incumbent Plants

	Output (1)	Worker Hours (2)	Machinery Capital (3)	Building Capital (4)	Materials (5)	Capital / Labor (6)	Energy / Capital (7)
Model 1: Mean Shift	0.1200**	0.0789*	0.0401	0.1327+	0.0911**	-0.0251	0.0008
	(0.0354)	(0.0357)	(0.0348)	(0.0691)	(0.0302)	(0.0418)	(0.0372)
Model 2: After 5 years	0.0826+	0.0562	-0.0089	-0.0077	0.0509	-0.1310	0.0585
	(0.0478)	(0.0469)	(0.0300)	(0.0375)	(0.0541)	(0.1228)	(0.0493)

Notes: Standard errors clustered at the county level are reported in parentheses. All outcome variables are run in logs. In columns (6) and (7), "Capital" is the combined stock of building and machinery capital. In column (7), "energy" is the combined cost of electricity and fuels: 163 plant-year observations are excluded that have missing or zero values for energy. See the text for more details. ** denotes significance at 1% level, * denotes significance at 5% level, + denotes significance at 10% level.