

Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices

Ali Hortaçsu and Chad Syverson

University of Chicago and National Bureau of Economic Research

This paper empirically investigates the possible market power effects of vertical integration proposed in the theoretical literature on vertical foreclosure. It uses a rich data set of cement and ready-mixed concrete plants that spans several decades to perform a detailed case study. There is little evidence that foreclosure is quantitatively important in these industries. Instead, prices fall, quantities rise, and entry rates remain unchanged when markets become more integrated. These patterns are consistent, however, with an alternative efficiency-based mechanism. Namely, higher-productivity producers are more likely to vertically integrate and are also larger, more likely to survive, and more likely to charge lower prices. We find evidence that integrated producers' productivity advantage is tied to improved logistics coordination afforded by large local concrete operations. Interestingly, this benefit is not due to firms' vertical structures per se: nonvertical firms with large local concrete operations have similarly high productivity levels.

We thank Daron Acemoglu, Tim Bresnahan, Jeremy Fox, John Haltiwanger, Justine Hastings, Tom Holmes, Tom Hubbard, Eddie Lazear, Canice Prendergast, Lynn Riggs, Mike Riordan, Mark Roberts, Nancy Rose, Kathryn Shaw, Mike Whinston, numerous seminar participants, and anonymous referees for helpful comments. Jeremy Shapiro provided valuable research assistance. Hortaçsu thanks the Center for Industrial Organization at Northwestern University for its hospitality, and Syverson is grateful for financial support from the National Science Foundation (award SES-0519062), the Olin Foundation, and the Stigler Center. The research in this paper was conducted while the authors were Special Sworn Status researchers of the U.S. Census Bureau at the Chicago Census Research Data Center. 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 ensure that no confidential data are revealed. Support for this research at the Chicago Research Data Center from NSF (awards SES-0004335 and ITR-0427889) is also gratefully acknowledged.

[*Journal of Political Economy*, 2007, vol. 115, no. 2]

© 2007 by The University of Chicago. All rights reserved. 0022-3808/2007/11502-0003\$10.00

I. Introduction

Is vertical integration a device for firms to create and harness market power, or does it enhance efficiency and improve social welfare? This paper looks at this issue in two vertically linked industries, using a rich data set on their producers, by investigating patterns of prices, productivity, entry, exit, and scale across integrated and unintegrated firms.

The reasons for, and results of, vertical integration have been a topic of considerable attention since Coase's (1937) landmark paper. Since then, economic theories of vertical mergers have evolved both in response to and as drivers of antitrust policy. Along with theoretical developments in efficiency-based theories of vertical integration, recent years have also seen the introduction of new theoretical arguments for the possible harm that can arise through vertical integration's foreclosure, or market power, effect.¹

This study utilizes integration episodes in the cement (Standard Industrial Classification [SIC] 3241) and ready-mixed concrete (SIC 3273) industries between 1963 and 1997 as an empirical laboratory to investigate the causes and consequences of vertical mergers.² A first look at simple patterns in the data is instructive. Table 1 shows how the prices and quantities sold in downstream industry (ready-mixed concrete) markets are related to the extent of vertical integration in those markets. The table reports coefficients obtained by regressing local average (logged) concrete prices or total (logged) quantities, depending on the specification, on either the market share or the number of vertically integrated firms in the downstream market. Specifications are reported both excluding and including market fixed effects.³ (Year fixed effects are included in all specifications to account for aggregate movements. We will further detail our sample, construction of the variables, and the market definition below.)

As the table makes apparent, the greater the presence of vertically integrated firms in a market, whether measured by their market share or number, the lower the prices and the higher the quantities in the

¹ Efficiency reasons for integration are espoused in, e.g., Williamson (1971) and Grossman and Hart (1986). Work on foreclosure includes Salinger (1988), Hart and Tirole (1990), Ordover, Saloner, and Salop (1990), Bolton and Whinston (1993), Riordan (1998), and Nocke and White (forthcoming). Snyder (1995), Bernheim and Whinston (2004), and Rey and Tirole (2007) offer comprehensive surveys.

² While often interchanged colloquially, cement is *not* concrete. Cement—made by baking limestone and clay or shale together in a kiln and grinding the result into a powder—is a single but important ingredient in the production of concrete. Ready-mixed concrete is produced by mixing cement with sand, gravel, water, and chemical admixtures and is what is contained in the familiar trucks with the spinning barrels on their backs. Thus cement is the upstream industry and ready-mixed the downstream industry.

³ Coefficients in the latter case thus reflect how concrete prices and quantities change *within a market* when that market becomes more vertically integrated.

TABLE 1
MARKET-LEVEL RELATIONSHIPS BETWEEN AVERAGE PRICES, TOTAL QUANTITIES, AND THE
EXTENT OF INTEGRATION

	QUANTITY-WEIGHTED AVERAGE PRICE				TOTAL QUANTITY			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
R^2	.065	.070	.430	.429	.072	.185	.890	.891
Market share of vertically integrated firms	-.143* (.029)		-.083* (.041)		1.532* (.290)		.331* (.154)	
Number of vertically integrated firms		-.037* (.006)		-.014 (.011)		.637* (.043)		.153* (.034)
Market fixed effects?	No	No	Yes	Yes	No	No	Yes	Yes

NOTE.—The table shows the coefficients obtained by regressing quantity-weighted average concrete prices and total concrete quantities sold in a market on either the market share or the number of vertically integrated firms operating in the market. The sample consists of 1,873 market-year observations. Standard errors are clustered by market.

* Significant at the 5 percent level.

local ready-mixed concrete market. These specifications are very simple and subject to deeper investigation below, but they do suggest at a first glance that factors other than foreclosure may be driving patterns in the data. We will see that these results are in fact robust, with foreclosure having little role in explaining what happens when firms in these industries vertically integrate.

While foreclosure effects do not appear to explain the data, an efficiency-based explanation may. Indeed, as we will show, once we explicitly account for efficiency differences, the impact of vertical integration on market-level outcomes ceases to be statistically and economically significant in most of our specifications. Plant-level evidence also reveals a strong relationship between vertical integration and productivity, especially downstream. Integrated ready-mixed concrete plants are more productive than unintegrated plants, even those in the same market.

What, then, is the efficiency benefit of vertical integration? We argue in Section V that in these particular industries, the benefit might not arise because of *vertical* integration per se. Instead, efficiency gains appear to be sourced in firms' abilities to operate multiple ready-mixed plants in the same market. This allows firms to harness scale economies tied to logistical coordination if they have the necessary managerial talent: deliveries of the industry's highly perishable product, typically ordered on very short notice by consumers at multiple locations, can be made more efficiently by having a central dispatch office that substitutes production and delivery among the firm's several local plants. We show how this efficiency gain is consistent with several features of our data as well as with the stated intentions and operational practices of firms with multiple concrete plants in the same market. While these

are in many cases vertically integrated firms, the evidence suggests that they need not have a vertical structure to benefit from such gains. This is reflected most tellingly in the evidence we document that vertically integrated concrete plants are no more efficient than those owned by non-vertically integrated multiunit firms with total local concrete sales of equal size.⁴

The cement and ready-mixed concrete industries have several practical features that make them favorable for a case study, and we will discuss these further below. Perhaps the most salient motivation for investigating these industries is that vertical integration among their producers has been the focus of substantial policy attention. U.S. antitrust authorities challenged the legality of several mergers in the late 1960s and early 1970s (no other industry saw as many vertical merger cases brought), and more recent cases have been considered by competition authorities in Europe, New Zealand, and Brazil. Thus our findings may be interesting from a policy evaluation perspective.⁵

This paper builds on an empirical literature that, compared to the resurgent theoretical literature surrounding the vertical foreclosure debate, has been relatively thin. Grimm, Winston, and Evans (1992) study railroad destination pairs served by a single firm and find that increased interline competition—competing railroads that connect intermediate points on the otherwise monopolized route—reduces welfare distortion. Waterman and Weiss (1996) and Chipty (2001) show that U.S. cable television systems integrated with content providers are more likely to include their own suppliers' paid content. Chipty, however, argues that higher-quality programming is offered in integrated markets, resulting in higher consumer surplus. Hastings and Gilbert (2005) utilize both within- and across-market variation in integration status among producers of wholesale gasoline. They examine the effect of vertical mergers on the wholesale price paid by competing independent gasoline retailers, and they find that the closer the competitor, the higher the wholesale price it has to pay.⁶ Asker (2004) tests for foreclosure due to exclusive dealing relationships in beer distribution and finds no significant evidence that exclusive dealing increases market power. Rosengren and

⁴ As explained further in Secs. II and III, our data allow us to link plants at the firm level. A firm that owns multiple production plants in any industry (or industries) is a *multiunit* firm. Vertically integrated firms, which own concrete *and* cement plants specifically, are a subset of multiunit firms.

⁵ It should be made clear that we are not trying to evaluate the appropriateness of the antitrust authorities' actions in any particular case. We are investigating *all* mergers between U.S. producers in the two industries, not just those that were the subject of legal action.

⁶ Hastings (2004) analyzes the impact of a multimarket vertical merger on retail gasoline prices in California. She finds that reducing the market share of unintegrated retailers led to higher retail prices. However, she does not draw a connection between her work and the foreclosure literature, instead ascribing the finding to product differentiation between branded and unbranded gasoline.

Meehan (1994) also find no support for foreclosure theory in event studies of the effects of vertical merger announcements and antitrust challenges on the stock prices of merging and rival firms. Snyder (1995), however, uses similar methods on a different set of industries and finds support for foreclosure.⁷

Section II documents cement's and ready-mixed concrete's patterns of integration over the sample period. It also discusses the evolution of economic theories of vertical integration over the same time frame, because this academic debate was mirrored by adjustments in policies aimed at vertical integration in the two industries. Section III describes the data used in the analysis. In Section IV we test the implications of foreclosure in the industry and compare foreclosure's predictions to those of an alternative efficiency-based explanation. We follow with a section examining this alternative in detail (Sec. V). A short discussion section (Sec. VI) presents conclusions.

II. Vertical Integration in Cement and Ready-Mixed Concrete: History, Policy, and Theory

Our data span 34 years, from 1963 to 1997. Over this time the cement and ready-mixed industries experienced two distinct periods of integration, separated by an over decade-long period of initial disintegration and then stability. Table 2 shows this evolution. Its first four rows report the fraction of cement and concrete industry plants and sales accounted for by vertically integrated producers (i.e., plants in firms owning both cement and ready-mixed plants). Notice that these increases in vertical integration coincide with growth in concrete plants in all *multiunit* firms (those owning two or more plants, regardless of industry). While all vertically integrated plants by our definition are in multiunit firms, other organizational structures meet this condition as well: firms that own several ready-mixed plants or firms owning a concrete plant(s) plus

⁷ The aforementioned policy actions against cement and concrete producers in the 1960s and 1970s also spurred at the time a small literature on the impact of forward integration by cement producers. These papers preceded the recent theoretical literature on foreclosure and as such do not directly test for such effects. They did, however, center on the issue of whether and how integration might enhance market power. Allen (1971) reviews evidence brought forth in the Federal Trade Commission's 1966 report on these industries (FTC 1966). While institutionally instructive, this study does not conduct formal statistical tests of economic hypotheses. McBride (1983) attempts to do so, finding a negative correlation between average cement prices in 17 markets and the cumulative number of market ready-mixed plants acquired by cement firms. He explains these results as resulting from the fact that forward integration into concrete makes it easier for cement plants to "adjust" their prices to accommodate demand swings, presumably as a result of (undocumented) collusive agreements in cement pricing that are easier to monitor than concrete prices. However, Johnson and Parkman (1987) argue that accounting for pre-existing price trends makes this result statistically insignificant.

TABLE 2
EVOLUTION OF VERTICAL INTEGRATION IN THE CEMENT AND READY-MIXED CONCRETE INDUSTRIES

	1963	1967	1972	1977	1982	1987	1992	1997
Cement plants that are vertically integrated (%)	21.9	47.4	41.9	34.8	32.5	35.2	49.5	30.5
Cement sales from vertically integrated producers (%)	25.2	51.2	48.4	41.0	49.5	51.3	75.1	55.4
Ready-mixed plants that are vertically integrated (%)	1.8	3.2	3.8	3.1	3.0	5.5	11.1	10.6
Ready-mixed sales from vertically integrated producers (%)	6.1	8.9	10.0	8.7	8.5	11.3	14.4	14.2
Ready-mixed plants in multiunit firms (%)	24.8	26.4	32.2	34.3	35.4	41.7	49.6	55.6
Ready-mixed sales from plants in multiunit firms (%)	40.1	46.3	52.4	54.0	50.9	57.5	61.3	65.0

NOTE.—The table shows the fraction of plants (or sales) accounted for by firms of various organizational types in the cement and ready-mixed concrete industries.

plants in another industry or industries. Thus the overall trend in vertical integration in concrete was matched by similar horizontal consolidation. This is a point we shall return to below when we discuss the productivity advantage of integrated producers.

The first vertical merger wave, driven by forward integration by cement producers, occurred in the early and mid-1960s. It received substantial attention from antitrust authorities. The FTC brought 15 antitrust cases during the 1960s against cement companies that had purchased concrete firms. Each case ended in divestiture of ready-mixed plants. In its report on these industries (FTC 1966), the commission cited several likely anticompetitive effects of vertical mergers, including limitation of unintegrated cement firms' downstream market access, increased entry costs for unintegrated concrete firms, and reduced bargaining power for these unintegrated concrete producers with their suppliers.

The vigorous enforcement action accompanied a chilling of merger activity in the sector throughout the 1970s. At the same time, however, the economic foundations of the so-called naive foreclosure theory exemplified by the FTC report were under attack by "Chicago School" critiques. Allen (1971), Posner (1976), and Bork (1978) pointed out that in cases of fixed-proportions technology (as cement is for ready-mixed concrete), a monopolist upstream producer cannot raise its profits by monopolizing its downstream market.⁸ Thus vertical mergers would occur only if there are efficiency gains.

⁸ Vernon and Graham (1971) show that if the downstream firm can substitute away from the monopolist supplier's input, a vertical acquisition may increase the monopolist's profits, though with ambiguous welfare effects.

These and similar arguments influenced antitrust authorities in the Reagan and first Bush administrations and softened official views toward vertical mergers considerably. In 1985, the FTC explicitly eased its enforcement policy regarding cement and ready-mixed vertical mergers.⁹ Industry firms responded by reintegrating. The fraction of cement plants in vertically integrated firms rose from 32.5 percent to 49.5 percent between 1982 and 1992, and the fraction of sales grew from 49.5 to 75.1 percent. For ready-mixed the corresponding growth for plants went from 3.0 percent to 11.1 percent and the share of industry sales from 8.5 to 14.4 percent.¹⁰

The theoretical debate pendulum swung back in the other direction in the late 1980s, as several authors formulated game-theoretic models to formalize certain conditions, robust to Chicago School criticisms, under which vertical mergers would have anticompetitive effects. An influential example of these newer foreclosure models is the “ex post monopolization” model of Hart and Tirole (1990).¹¹ Their framework has upstream homogeneous-good producers competing in prices. Importantly, for reasons that will become clear in a moment, they are assumed to have different marginal costs. In the cement-concrete context, these cost differences might arise because of productivity differences or because in a particular market a local cement producer has a transport cost advantage over more distant producers. Without loss of generality, one can think of the local upstream industry comprising a monopolist U that is disciplined by the potential entry of higher-cost competitors.

Suppose that this monopolist supplies two local downstream producers, D_1 and D_2 . The structure of the game is as follows: U offers each D_i a (possibly nonlinear) tariff $T_i(\cdot)$. Producers D_i then order quantities

⁹ The announcement cites “developments in economic thinking” as being in part behind the decision. It also states that cement mergers after 1977 were no longer subject to special consideration (FTC 1985, 21507).

¹⁰ The decline in integration seen from 1992 to 1997 was due to a demand-driven bout of unprecedented entry into the cement industry. The year 1997 saw a net gain of 61 plants over 1992’s 218 plants, which was roughly what had been the industry’s long-run average. These new plants were primarily small, unintegrated plants specializing in grinding clinker (an intermediate material in the cement-making process) that was kilned elsewhere, sometimes overseas.

¹¹ We thank Michael Whinston for an insightful discussion that led us to this model. Incidentally, in a section on applications, Hart and Tirole explicitly cite cement and ready-mixed mergers as involving the anticompetitive effects they highlight. The model is also described in detail in Snyder (1995), Bernheim and Whinston (2004), and Rey and Tirole (2007) and is used in the experimental study of Martin, Normann, and Snyder (2001). O’Brien and Shaffer (1992) show that the Hart-Tirole model’s assumption of Cournot competition downstream is not necessary to obtain foreclosure (we thank a referee for pointing us to this paper). They derive a result similar to the one we sketch here with a downstream industry comprising differentiated-product Bertrand competitors. Given the spatial nature of ready-mixed concrete competition, this may also be a reasonable modeling framework for our purposes.

q_i and pay $T_i(q_i)$. The D_i then produce q_i , observe each other's outputs, and set prices. Downstream competition is modeled as Bertrand with capacity constraints, which yields the Cournot outcome under conditions described in Kreps and Scheinkman (1983).

Let (Q^m, p^m) be the monopoly price and quantity if U also had monopoly power in the downstream market. Monopolist U would like to offer the quantity and tariff schedule $(q_i, T_i) = (Q^m/2, p^m Q^m/2)$ to the downstream producers so as to achieve monopoly profits. However, Hart and Tirole point out that offering this contract may not be credible if contracts are secret or can be secretly renegotiated ex post. If one downstream firm did agree to the above contract terms, the monopolist would then have an incentive to sell more than $Q^m/2$ to the other downstream firm.¹² But this action would lower the profits of the first downstream firm, making it unwilling to sign such a contract in the first place. Hart and Tirole show that the equilibrium with secretly renegotiable contracts has U selling Cournot quantities to the downstream firms and earning less than the monopoly profit.

However, U can commit to limiting total supply to the monopoly quantity by vertically integrating. An integrated U could sell Q^m through its subsidiary and sell nothing to the other downstream firm. This extreme outcome is unlikely in reality because the outside downstream firm could try to buy input from another upstream supplier. But since U has a cost advantage upstream—and this is where the asymmetric cost assumption is important—the integrated firm can (and will) also sell the input to the outside downstream firm. However, it will be at a price that just undercuts the higher-cost upstream competitor. This leads to an asymmetric Cournot outcome downstream, with the unintegrated downstream firm facing higher costs than its integrated competitor. The greater the upstream cost asymmetry, the closer equilibrium gets to the monopoly outcome.

The end results of integration, therefore, are higher average prices and lower quantities in the downstream output market. Moreover, if the downstream technology has a fixed operating cost (or if no other upstream supplier is available), exit of unintegrated downstream firms is more likely. While the above version of the model is silent about what happens to (potential) competitors in the upstream market, Hart and Tirole also provide a variant—the “scarce needs” version, in which two upstream competitors have the same marginal costs but different fixed entry costs—in which vertical integration leads to the nonentry of one of the upstream firms.

¹² The monopolist takes $Q^m/2$ as given for the contracted firm and reoptimizes quantity for the other firm using the “right-shifted” version of the residual demand curve. This yields an optimum quantity that is greater than $Q^m/2$.

Thus if foreclosure effects are dominant, we should expect that vertical integration will be associated with higher prices, lower quantities of the downstream good, higher exit rates among unintegrated producers, and lower entry rates downstream and upstream.

Note that if vertical integration increases productive efficiency, we would instead expect lower prices and higher quantities of concrete. However, one may still expect higher exit rates among unintegrated producers if they do not benefit from the efficiency enhancements experienced by integrated rivals. Entry rates may also be lower if the potential entrants are not efficient enough to survive in the market, though the opposite may be true if acquisition of existing downstream plants provides an avenue of entry and expansion by low-cost producers in other markets (including cement producers that want to enter into the concrete business).

It is also reasonable to expect that foreclosure and efficiency effects could simultaneously operate in many market settings. As detailed in the next two sections, our empirical analysis will address this by including explicit controls for plant-level productivity to look at the impact of integration on prices, quantities, entry, and exit, net of efficiency improvements. We suspend until Section V the investigation into the logistics coordination mechanism that ties vertical integration to productivity.

III. Data

A. *Plant-Level Ownership, Productivity, and Prices*

The core of our analysis uses plant-level microdata from the 1963, 1967, 1972, 1977, 1982, 1987, 1992, and 1997 *Census of Manufactures*. The census is comprehensive; we observe every U.S. cement and ready-mixed concrete plant operating in the respective census years. A typical census has 220 cement and 5,200 ready-mixed plants, though these relatively stable levels hide substantial gross plant entry and exit in both industries. The census microdata contain a wealth of information on plants' production activities that we exploit below. Crucially here, they also contain firm identification numbers for each plant. (Plants, or "establishments" in Census Bureau terminology, are unique physical locations in which products are made. A firm can own one or more plants.) Thus we are able to observe when a single firm owns plants in both industries, in which case we consider the firm and its component plants vertically integrated.

The comprehensiveness of the *Census of Manufactures* is extremely useful. It allows us to observe each plant's integration status in each census year, and because the data contain permanent plant identifiers

that are invariant to ownership changes, we can track changes in this status over time. Entry and exit of plants between census years are completely observable as well, allowing us to look at vertical integration's long-run impact on markets.

Besides ownership information, the census contains data on plant revenues, several labor input measures, book values of equipment and structures capital stocks, expenditures on inputs, and state and county codes. We use these data directly and also to calculate other plant-level technology measures such as labor and total factor productivity (TFP), measures of overhead labor, and capital-labor ratios. Details of their construction are in the Data Appendix. Further, for some ready-mixed concrete plants, we have access to auxiliary files containing the total value and the physical quantity of product shipments at a highly detailed (seven-digit SIC) level. We can therefore measure plant output in either dollars or physical units (cubic yards, in this case), as well as calculate the plant's average unit price (measured on a free-on-board basis and deflated to 1987 dollars using the industry-level price index from the NBER Productivity Database).¹³

B. Local Markets in the Cement and Ready-Mixed Concrete Industries

A useful attribute of the cement and ready-mixed industries as a forum for testing foreclosure theory is that, because of the high transport costs inherent to the industries, they contain many local markets. This naturally raises the empirical issue of defining these markets. We use different market definitions for cement and concrete because while transport costs are high in both industries, they are higher for concrete, and shipment patterns reflect this.

For ready-mixed concrete, we use the component economic areas (CEAs) created by the Bureau of Economic Analysis (BEA). CEAs are usually collections of counties, but not always, centered on metropolitan statistical areas (MSAs), and they are mutually exclusive and exhaustive of the land area of the United States. The roughly 3,200 U.S. counties are grouped into 348 such markets. The procedures used to create CEAs ensure that counties in a given CEA are economically intertwined; see Johnson (1995) for detailed information on their construction. We be-

¹³ Conveniently, most plants in the ready-mixed concrete industry are highly specialized; virtually all their revenue—roughly 95 percent on average—comes from sales of ready-mixed, which is itself a seven-digit product. Other concrete products such as block, prefabricated structural members, and pipe are typically made by producers in concrete industries other than SIC 3273. Likewise, the amount of ready-mixed produced at these other industries' plants is minuscule. Unlike the ready-mixed industry, cement plants produce a number of seven-digit products (different cement types based on their chemical composition). This makes cement plants' average unit output prices more difficult to construct and less comparable across plants.

lieve that CEAs are well-founded, albeit imperfect, market area measures for an industry in which maximum ideal delivery distances are 30–45-minute drives from the plant. See Syverson (2004) for an extended discussion of the trade-offs involved in defining ready-mixed markets. Since most, though not all, of the empirical exercises below deal with downstream markets, CEAs are the primary geographic demarcation used in the paper.

For cement, with its lower transport costs, we use larger market areas. Our primary market definition is that used by the U.S. Geological Survey in its *Minerals Yearbook* annual surveys of the industry (see, e.g., USGS 2004). These 26 markets are essentially states or collections of states, with some states with particularly large numbers of cement plants (California, Pennsylvania, and Texas) split into separate markets. This reflects an apparent USGS objective of roughly equating the number of plants across markets. Because of our concern that this objective may sometimes get in the way of defining the true geographic markets present in the industry, in some specifications we use economic areas (EAs), which are aggregations of CEAs.¹⁴ There are considerably more EAs (172) than USGS cement markets, and they are not beholden to the latter's state boundary or producer equalization conditions.

Table 3 shows some summary measures of market structure for our geographic markets in both industries. The average number of ready-mixed plants in a CEA market-year observation in our sample is 14.6, and the median is 10. The largest market has 125 ready-mixed plants. The respective values for firm counts are 11.8, 8, and 110. As suggested by the above discussion about the fraction of integrated ready-mixed plants, most markets are not integrated: the average number of integrated firms in a market is 0.3, and the median is zero (though certainly in the sample's later years these values are higher). If we condition on at least one integrated firm operating in the concrete market, however, the average market share of integrated firms is substantial: 31.5 percent. The median Herfindahl-Hirschman index (HHI) for concrete markets in our sample is 2,464.¹⁵

For the USGS cement markets, the average number of plants is 8.3,

¹⁴ While there are on average two CEAs per EA, larger and denser EAs typically have more CEAs than those in less populated areas. For example, the Kansas City EA comprises the Kansas City (KS-MO), Lawrence (KS), and St. Joseph (MO) CEAs, whereas the Bangor (ME) CEA is the only one in the Bangor EA.

¹⁵ We estimated some descriptive multivariate regressions to gauge the relationship between market structure and the extent of integration. The market share of integrated firms is higher in more concentrated markets (as measured by HHI) as well as in those with high demand (i.e., larger markets; see below how we construct this demand measure from local construction activity data). However, there is no clear relationship between integrated firms' share of the local concrete market and concentration in the *cement* market in which the CEA is contained.

TABLE 3
SUMMARY STATISTICS OF MARKET STRUCTURE IN THE CEMENT AND READY-MIXED
INDUSTRIES

	Average	First Quartile	Second Quartile	Third Quartile	Maximum
A. Ready-Mixed Concrete (2,767 CEA-Year Markets)					
Number of plants	14.6	5	10	17	125
Number of firms	11.8	5	8	14	110
Number of vertically integrated firms	.3	0	0	0	6
Number of vertically integrated firms (if >0)	1.4	1	1	2	6
Market share of verti- cally integrated firms (if >0)	.314	.128	.270	.430	1.000
HHI	3,001	1,585	2,464	3,832	10,000
B. Cement (208 USGS-Year Markets)					
Number of plants	8.3	5	7	11	24
Number of firms	7.1	5	6.5	9	23
Number of vertically integrated firms	2.5	1	2	4	7
HHI	2,795	1,990	2,362	3,348	10,000

NOTE.—The table shows moments of the distribution of market structure measures for the cement and ready-mixed concrete industries. See the text for details and market definition.

the median 7, and the maximum 24. A larger fraction of firms are vertically integrated: an average of 2.5 per market (of 7.1 firms total). The median HHI is 2,362. The counts are of course lower (and the concentration numbers higher) for EAs, since these are notably smaller than USGS markets.

In most of the empirical tests below, we control for activity in the local construction sector (SICs 15–17), which buys the vast majority of the cement and concrete industries' outputs. These demand measures are built from County Business Patterns employment data for the sector that we aggregate to the market level. Total market demand (size) is measured simply as logged total construction employment, and market density is the (log of the) number of construction-sector workers per square mile in the market.¹⁶

¹⁶ As discussed in Syverson (2004, forthcoming), construction activity is likely exogenous to the specifics of local concrete competition because while the construction sector accounts for virtually all ready-mixed sales, concrete's cost share among the sector's intermediate inputs is small. Hence any changes in concrete prices should not by themselves create notable changes in the level of construction activity. We expect that similar effects operate on a broader geographic scale in cement.

IV. The Extent of Integration and Market Power

Our empirical approach is straightforward. We compare patterns in the data to the predictions of foreclosure theory as well as the alternative explanation that increases in integration reflect the expansion of larger and more efficient producers. Our ability to distinguish foreclosure from the alternative rests on two key sets of implications. First, foreclosure and productivity/efficiency theories have opposing predictions regarding the response of downstream prices and quantities to integration. Second, while the two stories' predictions for firm turnover are similar, foreclosure implies an effect of vertical integration per se. Productivity/efficiency explanations, on the other hand, imply that any observed effect of integration on turnover acts only inasmuch as productivity and integration are connected; hence integration's effects should disappear once productivity is controlled for.¹⁷

A. *Does Vertical Integration Raise Prices?*

Foreclosure implies a positive relationship between downstream prices and the extent of vertical integration in a market, whereas an efficiency-based explanation for integration predicts a negative relationship. As seen in table 1 above, the basic patterns in the data are not consistent with foreclosure. We now investigate whether these market-level results hold up to deeper scrutiny. We first take advantage of our detailed plant-level data to see if there is a relation between vertical integration and prices among producers in the same market. Next, we examine the efficiency and vertical integration link directly in plant-level productivity data. Finally, we revisit the market-level evidence displayed in table 1, this time using a richer set of controls and paying more attention to potential endogeneity problems.

¹⁷ A third distinguishing implication of foreclosure models is that unintegrated downstream producers face higher input costs. We can in principle test this in our data, because a subset of plants in the *Census of Manufactures* answer questions regarding their total expenditures on and purchased quantities of detailed intermediate inputs, allowing us to compute unit cement prices paid by concrete producers. When we do so, we find that, rather than unintegrated concrete producers reporting higher cement input prices, there are no significant differences across integrated and unintegrated plants. While on its face this is not what is suggested by foreclosure, it is far from clear what the reported cement prices from integrated concrete makers reflect. The Census Bureau instructs establishments to report the value of internal materials transfers "at their full economic value (the value assigned by the shipping plant, plus the cost of freight and other handling charges)." It is possible that these results simply reflect integrated plants "marking to the market" their internal transfer prices.

TABLE 4
VERTICAL INTEGRATION AND READY-MIXED CONCRETE PRICES: PLANT-LEVEL RESULTS

	WITHIN-MARKET DIFFERENCE				CHANGE FOR CONTINUERS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Observations	12,553	12,553	8,555	8,555	4,025	4,025	2,439	2,439
R^2	.424	.425	.501	.501	.456	.460	.575	.576
Vertical integration indicator	-.022* (.006)	-.017* (.006)	-.006 (.007)	-.006 (.007)	.011 (.035)	.011 (.035)	.003 (.034)	.005 (.034)
Multiunit indicator		-.012* (.004)		-.001 (.005)		-.037 (.020)		-.033 (.028)
TFP			-.214* (.015)	-.215* (.015)			-.237* (.028)	-.237* (.028)
	INTEGRATED VS. UNINTEGRATED ENTRANTS				INTEGRATED ENTRANTS VS. UNINTEGRATED INCUMBENTS			
	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Observations	2,771	2,771	2,025	2,025	7,490	7,490	6,104	6,104
R^2	.563	.566	.655	.655	.430	.430	.523	.523
Vertical integration indicator	-.037* (.018)	-.025 (.019)	-.020 (.020)	-.012 (.020)	-.023 (.012)	-.023 (.012)	-.012 (.012)	-.012 (.012)
Multiunit indicator		-.032* (.012)		-.025 (.014)		NA		NA
TFP			-.218* (.031)	-.215* (.031)			-.221* (.019)	-.221* (.019)

NOTE.—The table shows the results from regressing plant-level concrete prices on a number of variables. Vertical integration (multiunit) indicator is equal to one if the plant is in a vertically integrated (multiunit/multiplant) firm and zero otherwise; TFP is the plant's quantity-based total factor productivity. See the text for details. In the vertically integrated entrants vs. non-vertically integrated incumbents comparison, there are no observations of new multiunit entrants that are not also vertically integrated entrants. All regressions include market-year fixed effects.

* Significant at the 5 percent level.

1. Plant-Level Evidence

We start our look at the plant-level relationship between downstream prices and vertical integration by regressing ready-mixed concrete plants' (logged) prices on an indicator for the plant's vertical integration status and a full set of market-year fixed effects. The vertical integration dummy coefficient therefore captures the mean price difference across integrated and unintegrated producers in the same market and time period. This specification is useful in that it compares producers that are facing the same market-level demand and supply conditions, which the market-level regressions in table 1 do not fully control for.

The results of this exercise are shown in column 1 of table 4. On average, integrated producers charge prices 2.2 percent lower than their unintegrated competitors in the same market. Thus the negative relationship between downstream prices and integration seen across markets also holds even among producers within them.

We further investigate these plant-level patterns with several additional

comparisons, again all within the same market-year. Column 5 of the table compares price *growth rates* among continuing plants (i.e., those that operated in both the previous and current censuses) that were unintegrated in the previous census. The vertical integration indicator in this case takes a value of one if a plant became integrated between the previous and current censuses. The estimated coefficient implies that these newly integrated plants have somewhat—1.1 percent—more positive price changes than their cohorts that remain unintegrated, but the difference is not statistically significant. Columns 9 and 13 contrast the prices of integrated entrants (entrants are defined as plants appearing for the first time in the current census) to those of unintegrated entrants and unintegrated incumbents, respectively. As with the comparison in column 1, integrated producers charge lower prices than those in the unintegrated comparison groups. Integrated firms' new plants price 3.7 percent below new plants of unintegrated firms and 2.3 percent below integrated incumbents (though this last difference is only marginally significant, with a *p*-value of 0.07). Integrated producers' lower prices thus arise not because formerly unintegrated existing plants see large price drops upon becoming integrated, but instead because newly built integrated plants tend to charge lower prices, and this difference persists over time.

To get at the source of integrated plants' lower prices, we introduce two additional controls into each of the four plant-level price comparisons above. These are a dummy indicating if a plant is owned by a multiunit firm and the plant's TFP level. The multiunit dummy captures, in a rough measure, the size of the firm that owns the plant. Recall from the discussion in Section II that all vertically integrated plants are in multiplant firms, but nonvertical organizational structures meet this condition as well. So when we include the multiunit indicator in a specification, the vertical integration coefficient reflects differences between integrated plants and other plants in multiunit firms without cement operations. The TFP measure reflects a producer's efficiency, in other words, the inverse of its cost level. We use a physical TFP measure—that is, the output measure in the TFP index is the physical output of the plant—rather than the literature standard revenue-based output measure. As Foster, Haltiwanger, and Syverson (2005) discuss, the latter attributes price differences across industry plants to output differences. This is something one would like to avoid, and we can do so since we can measure the actual cubic yards of concrete that plants ship.

The remaining columns of table 4 show each of the four plant-level price comparisons with the multiunit status dummy and physical TFP measure included in the specifications, first separately and then to-

TABLE 5
PLANT-LEVEL TFP COMPARISONS

	Within-Market Difference (1)	Change for Continuers (2)	Integrated vs. Unintegrated Entrants (3)	Integrated Entrants vs. Unintegrated Incumbents (4)
Observations	8,555	2,439	2,025	6,104
R^2	.308	.419	.573	.352
Vertical integration indicator	.043* (.014)	.102 (.055)	.054 (.045)	.046* (.028)

NOTE.—The table shows the results from regressing ready-mixed concrete plants' TFP levels on an indicator for the plant being in a vertically integrated firm. All regressions include market-year fixed effects.

* Significant at the 5 percent level.

gether.¹⁸ Controlling for a plant's multiunit status explains some of the price difference between plants within a market, to the point of making the vertical integration indicator insignificant in the comparison of entering plants' prices; but prices of vertically integrated plants remain significantly lower than those of other multiunit plants in the broadest comparison. More starkly, TFP always enters negatively into the regressions and reduces the magnitude and eliminates the statistical significance of the vertical integration indicator. These results suggest that vertically integrated producers charge lower prices because to some degree all plants in multiunit firms do, but especially because they have higher productivity levels.¹⁹

2. Are Integrated Plants More Productive?

To look at the productivity-integration relationship more closely, we repeat the four sets of comparisons above, except we replace prices with plants' TFP levels.²⁰ (Again all specifications include CEA-year fixed effects, so results reflect variation among plants in the same market-year.) The results are shown in table 5. Column 1 indicates that vertically

¹⁸ The number of observations in the specifications including plant TFP is smaller because capital stocks, which we need to construct a TFP measure, were not available in the 1967 *Census of Manufactures*, whereas price data were.

¹⁹ The fact that we construct plant-level prices as the ratio of reported revenues to physical quantities means that if there is measurement error in quantities, this will create a negative correlation between measured prices and quantities (and, by implication, TFP as well). This could potentially lead to these specifications overstating the actual correlation between prices and productivity levels. However, work in Foster et al. (2005) and Syverson (forthcoming) indicates that the negative relationship is not a spurious result driven by measurement error. Further, as we show below, the negative price-productivity correlation also holds for market-level averages, where the averaging process should lessen any influence of measurement error.

²⁰ We will explore the connection between vertical integration and firm size in detail below.

integrated plants have TFP levels that are 4.3 percent higher on average than their unintegrated competitors. In other words, they produce 104.3 cubic yards of ready-mixed concrete from the same inputs that unintegrated producers in the same market use to produce 100 cubic yards. Comparing TFP changes among formerly unintegrated plants in column 2 shows that plants that become integrated see roughly 10 percent faster productivity growth, though this difference is only marginally statistically significant. TFP gaps similar to those in the entire sample are observed when we restrict the comparison to between integrated entrants and either unintegrated entrants or incumbents. The significance of these differences is weaker, however, likely because of in part the reduced sample sizes.

These patterns are consistent with the alternative explanation to foreclosure forwarded above: namely, that increases in integration reflect expansion of more efficient producers. Their lower costs drive down average market prices directly, because they have lower costs and can charge less, but also indirectly, by reducing their competitors' optimal prices in equilibrium.

3. Market-Level Evidence Revisited

With these plant-level findings in hand, we now revisit market-level analyses of the types shown in table 1. We check the robustness of these earlier results by adding more controls. We include the HHI concentration measure of the CEA-year observation to control for the horizontal structure of the local concrete market. As discussed above, the extent of vertical integration in the downstream market is related to downstream (though not upstream) concentration. Controlling for this correlation lets us see the vertical integration–price relationship that is orthogonal to any changes in horizontal concentration. We also control for demand density (described above) in the local market; Syverson (forthcoming) shows how it affects average prices through competition-driven selection and markups. We cluster standard errors by CEA market to account for any temporal correlation in unobservables.

The results are shown in table 6. Columns 1, 2, 6, and 7 repeat the market-level price specifications from table 1 (for both extent of integration measures) while including the additional controls. Just as in table 1, more integrated concrete markets have lower average prices.²¹

²¹ For space and clarity reasons, we do not report the HHI or demand density coefficients. HHI typically enters positively, indicating that, all else equal, higher concentration coincides with higher prices. However, it is usually not statistically significant, and its economic magnitude is modest: a 2,250-point increase in the HHI—the interquartile range—is associated with a roughly 1 percent increase in prices. Density too is typically insignificant, except in the specification excluding market fixed effects, in

TABLE 6
VERTICAL INTEGRATION AND READY-MIXED CONCRETE PRICES: MARKET-LEVEL RESULTS
BENCHMARK SPECIFICATION

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Observations	1,870	1,870	1,870	1,550	1,550	1,870	1,870	1,870	1,550	1,550
R^2	.087	.433	.434	.573	.573	.087	.432	.432	.573	.573
Market share of vertically integrated firms	-.125*	-.090*	-.086*	-.043	-.043					
Market share of multiunit firms	(.028)	(.041)	(.041)	(.039)	(.039)					
			-.015	.001	.001					
			(.022)	(.024)	(.024)					
Number of vertically integrated firms						-.028*	-.015	-.013	-.009	-.007
Number of multiunit firms						(.007)	(.011)	(.011)	(.009)	(.009)
								-.003	-.004	-.004
								(.004)		(.004)
Quantity-weighted average TFP				-.293*	-.293*				-.294*	-.294*
				(.054)	(.054)				(.054)	(.054)
Market fixed effects?	No	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes

NOTE.—The table shows the coefficients obtained by regressing quantity-weighted average concrete prices in a market on either the market share or the number of vertically integrated firms operating in the market. The market share or number of multiunit firms and the quantity-weighted average TFP in the market are also included in some specifications. All regressions control for the HHI and density of demand in the market as well as year effects (coefficients not reported). Standard errors are clustered by market.

* Significant at the 5 percent level.

The magnitudes of the coefficients track the earlier values, though the specifications that include market fixed effects have smaller estimates that in one case become insignificant.²² The implied average price differences are substantial. For example, column 1 implies that going from a market with no integrated producers to one in which integrated firms hold a market share of 0.315 (the average share conditional on at least one vertically integrated producer being in a market) corresponds to a 4 percent decline in the average price of ready-mixed. Likewise, the coefficient in column 6 indicates that an additional integrated firm operating in a market corresponds to a roughly 3 percent drop in average prices. These estimated magnitudes are roughly 25–35 percent of the between-market standard deviation in average prices of 10.2 percent.

As suggested by the plant-level analyses above, we investigate how controlling for multiunit status and TFP affects the price-integration relationship. For market-level controls analogous to those above, we use the market share (or number) of multiplant firms and the quantity-weighted average TFP in the market. Again we add each control separately to the regressions and then control for them simultaneously. The results are in columns 3–5 and 8–10 (again the sample size falls because we cannot measure TFP in the 1967 *Census of Manufactures*).

The results are similar to those in the plant-level analyses above. The market share or number of multiunit firms enters negatively and slightly reduces the estimated impact of vertical integration, but the greatest impact is found when average TFP in the market is included in the specification. The average TFP coefficients are negative; intuitively, growth in the average productivity of a market's producers is associated with declining average prices. Our TFP control therefore does seem to capture efficiency effects on market outcomes. Moreover, the coefficients on the extent of integration measures remain negative but shrink in size and become insignificant. Hence if changes in average TFP sufficiently capture the efficiency gains that accompany vertical integration, the results suggest that the remaining price effects of integration are essentially zero. These results suggest two inferences. First, any impact of the extent of vertical integration in a market appears to be driven by the increase in average productivity levels among market producers rather than something tied to integration per se. Second, it does not

which case it enters negatively. It is positive and insignificant otherwise. The sign switches because within-market changes in density actually reflect only changes in the level of demand, since market areas are fixed. *Ceteris paribus*, higher demand will raise prices in a market.

²² The smaller coefficients with market fixed effects could result either because there are market unobservables correlated positively with vertical integration and negatively with prices (more on this below) or because attenuation bias from measurement error (say if our market definitions imperfectly capture the true geographic markets) is exacerbated when we identify vertical integration effects only from within-market changes.

appear that efficiency gains from vertical mergers or other unobservables are simply hiding foreclosure effects in the raw data. Once we control for the change in average efficiency, we still do not see a positive impact of integration on prices.

These results are not on their face consistent with the positive relationship between prices and integration implied by foreclosure. Of course, the decision to vertically integrate is a choice made by firms, and the extent of integration is not random across our sample. Perhaps integration does facilitate foreclosure, but there are market-specific unobservables correlated with both higher integration intensity and lower prices. We make a few observations on this point. First, any unobservables would have to have particularly strong correlations with higher integration intensity and lower prices to swamp the positive price effects of foreclosure, since the observed net integration-price correlations are substantially negative. Second, it is unclear what the candidates for such unobservables might be. Suppose that the price drops are driven by unobservable inward shifts in market demand (and recall that we are already controlling for demand density, which is positively correlated with demand across markets and whose growth within markets equals demand growth since areas are fixed). If so, these unobservable demand contractions would have to induce large, expanding firms—which, as we will discuss later, aptly describes integrated firms in our sample—to prefer integrating into such markets. But one typically expects entry and expansion to follow positive demand shocks, not negative ones. Nonetheless, possible endogeneity concerns should be minimized to the greatest extent possible given the data. Including market fixed effects goes some distance in this direction, since it removes any market-level influences on prices and the extent of integration that are fixed over time. However, we take additional approaches toward this end. These are discussed below.²³

Our first approach uses variation in antitrust authorities' enforcement regimes (some of which we discussed briefly above in Sec. II) and spatial demand patterns within markets to instrument for market structure. For the sake of brevity, we detail the construction of the instrument set in the Data Appendix, but their logic can be summarized succinctly. The antitrust enforcement instruments use the stated policy rules of the U.S. antitrust agencies (both the FTC and the Department of Justice) to identify instances in which the mechanical nature of the rules creates plausibly exogenous variation in the extent of antitrust scrutiny across our sample markets, either in cross section or over time. These policies

²³ Vertical integration is not the only market structure attribute that is endogenous: the market share of multiunit firms and concentration reflected in the HHI are as well. Our methods also attempt to address this endogeneity.

include rules applied both economywide and specifically to the cement and ready-mixed industries. The spatial demand instruments work off the facts that construction activity should be exogenous to outcomes of the ready-mixed market because of concrete's small share in construction costs (as discussed above), and that demand variation across areas within a given market plausibly affects the returns to a firm owning multiple concrete plants. Both the antitrust policy and spatial demand instruments should have explanatory power over firms' preferred organizational structures within and across markets—that is, whether they are multiunit or vertically integrated, and how large a market share they can have within a given market—and be orthogonal to other local demand or supply shocks. The first-stage results, not reported here, indicate that interactions of firm and market structures, and the extent of vertical integration in a market in particular, are related to the instruments in the expected ways. For example, the extent of integration in concrete markets with a larger share of firms specifically marked by the FTC for merger scrutiny is lower during the years in which this policy was in force.

The results of the instrumented specification are shown in columns 1–4 of table 7. (All specifications control for demand, downstream HHI, and market fixed effects.) The implied impact of vertical integration on prices is negative, as with the benchmark results above. Moreover, the magnitude of the implied relationship between integration and prices is larger than in table 6 (in the fully saturated specification, moving a market from integrated firms having zero market share to those having a 31.5 percent share implies a 25 percent drop in average price). We note, however, that the standard errors are much larger here than in the benchmark specification. This lack of precision, which is even more pronounced in our market-level quantity regressions discussed in subsection *B* below, may arise because our instruments have difficulty predicting changes in vertical integration that are orthogonal to changes in horizontal structure (as captured by the share of multiunit firms and market HHI).²⁴

Another approach to address endogeneity concerns estimates price effects using the subsample of CEA-years involved in multimarket mergers. The idea, as with Hastings and Gilbert (2005), is to identify inte-

²⁴ As pointed out by Shea (1997), in the presence of multiple endogenous variables in a specification, instruments need not just be relevant to each endogenous variable; they must have *linearly independent* relevance to each. We constructed the R^2 relevance statistics Shea suggests to measure this and found, indeed, that after we accounted for the market share of multiunit firms and market HHI (two other endogenous variables), the instruments did not have a high degree of explanatory power over the market share of vertically integrated producers. This suggests caution in interpreting the results from this specification, particularly with regard to the validity of the standard errors, even though the instruments pass the commonly used *F*-test in the first stage.

TABLE 7
VERTICAL INTEGRATION AND READY-MIXED CONCRETE PRICES: MARKET-LEVEL RESULTS
ROBUSTNESS CHECKS

	ANTITRUST AND SPATIAL DEMAND INSTRUMENTAL VARIABLES				MULTIPLE-MARKET MERGERS				ADDING STATE-YEAR EFFECTS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Observations	1,870	1,870	1,550	1,550	142	142	133	133	1,870	1,870	1,550	1,550
R^2	NA	NA	NA	NA	.453	.472	.584	.602	.586	.586	.681	.681
Market share of vertically integrated firms	-.468* (.217)	-.722* (.308)	-.707* (.220)	-.914* (.331)	-.164* (.049)	-.132* (.054)	-.098 (.051)	-.068 (.051)	-.071 (.044)	-.065 (.045)	-.022 (.044)	-.019 (.044)
Market share of multiunit firms		.285 (.220)		.343 (.348)		-.141 (.105)		-.136 (.089)		-.022 (.025)		-.008 (.027)
Quantity-weighted average TFP			-.287* (.019)	-.296* (.043)			-.363* (.067)	-.364* (.066)			-.269* (.056)	-.269* (.056)

NOTE.—The table shows the coefficients obtained by regressing quantity-weighted average concrete prices in a market on the market share of vertically integrated firms operating in the market. The market share of multiunit firms and the quantity-weighted average TFP in the market are also included in some specifications. All regressions control for the HHI and density of demand in the market as well as year effects (coefficients not reported). The various identification approaches, listed in the header row, are detailed in the text. Standard errors are clustered by market.

* Significant at the 5 percent level.

gration's impacts using arguably random variation in the extent of integration across the several markets that experience a single-firm merger. For example, consider the simple case of a cement firm integrating by purchasing a concrete firm that owns two plants, each operating in a different market. One plant has a market share of 20 percent in its local market, the other a 30 percent share. If the cement firm's decision to integrate with this particular downstream firm was made on the basis of considerations of the target firm as a whole rather than in its specific markets, differences in the extent of integration due to the merger—integrated firms' market share will grow 10 percentage points more in the second market than in the first—will be exogenous to market-specific outcomes. Hastings and Gilbert use this type of variation from a particular merger in the wholesale gasoline industry that affected 14 markets. While most of the multimarket mergers in our sample involve fewer markets, we observe not one but many such mergers during our observation period.

We define a multiple-market merger as having occurred if a single firm newly integrates plants in two or more markets during the same intercensus period. This happened 29 times in our sample, each merger involving an average of just under five markets, giving us a sample of 142 market-years with which to estimate vertical integration's effects. We run the price regressions as before on this subsample but now include merger fixed effects in the regressions. Thus all the variation in vertical integration intensity used to identify integration's impact comes from across-market variation *within* particular merger episodes. As in the example above, the specification uses the fact that the market share vertical integration measure was 10 percentage points higher in the second market than in the first and compares this to the difference in average prices between the two markets.

The results using the sample of multimarket merger episodes are shown in columns 5–8 of table 7. Again, all specifications control for the market demand and HHI, though these coefficients are not reported. A higher market share of vertically integrated firms corresponds to lower prices. (Similar results, not reported here, are obtained when the number of vertically integrated firms is used to measure the extent of integration in the market.) The magnitude of the vertical integration coefficient is notably similar to that in the benchmark specification.²⁵ When we add controls for the share of multiunit firms and the average

²⁵ Recall that the exogeneity argument here is that firms make merger decisions on a whole-firm basis, not on the expected market structure change in any particular market they are merging into. Obviously we cannot test this proposition directly. However, arguments that particular market outcomes are the primary concern would have to explain, as is implied by the results, why the merging firms choose to most intensively enter markets with idiosyncratically low prices.

TFP in the market, we also find results similar to those above. The multiunit share is negatively (albeit weakly) related to average prices and explains part of the correlation between vertical integration and prices. Average TFP is also associated with lower prices (we lose nine CEA-year markets and four merger episodes in the smaller sample), and including it makes the extent of integration measure insignificant. When we control for both multiunit share and average TFP, integration is negatively but insignificantly connected to average prices.

Our third approach to address possible endogeneity is to add state-year fixed effects to our benchmark specification. These control for any state-specific unobserved demand or supply shocks that might be correlated with the extent of integration and average prices. One example of such shocks would be local government construction spending that for one reason or another is not reflected in our construction employment (demand) measure. Note that we still include CEA effects as well; in effect, the specification identifies vertical integration's impact from intertemporal price and integration variations across different CEAs in the same state.²⁶

The results are in columns 9–12 of table 7. Not surprisingly, the reduced variation we have to estimate the link between integration and prices leads to less statistical significance; in no case does the vertical integration coefficient have a *p*-value below 0.11. However, the qualitative features of the results echo those seen in the benchmark. Vertical integration is negatively related to average concrete prices, as are the share of multiunit plants and the average TFP level in the market, though only the latter is significantly so. This is again consistent with efficiency effects driving the negative relationship between prices and the extent of integration.

On the whole, the many results presented in this section do not point to foreclosure-driven price increases. Instead, vertical integration is associated with lower prices at both the plant and market levels. We do find evidence that this negative relationship is not driven by vertical integration per se, but rather because vertical integration is tied to larger, more productive firms. This is consistent with the alternative efficiency-based explanation of the observed patterns in the data. In

²⁶ A measurement issue arises here because states do not completely nest every CEA (recall that CEA boundaries need not coincide with state boundaries). To implement this empirical specification, we have to assign a “primary” state to each CEA that is not wholly contained within one state. We do so by counting the number of counties of each state in every multistate CEA and assigning the primary state designation to the state with the most counties. In the few cases in which the maximum counts were equal, the state with the primary influence was apparent because it had a large metropolitan area in the CEA. The state-year dummies were constructed by interacting year effects with either this “modal-state” designation for multistate CEAs or the unique state containing the other CEAs.

TABLE 8
VERTICAL INTEGRATION AND READY-MIXED CONCRETE QUANTITIES: PLANT-LEVEL RESULTS

	WITHIN-MARKET DIFFERENCE				CHANGE FOR CONTINUERS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Observations	12,553	12,553	8,555	8,555	4,025	4,025	2,439	2,439
R^2	.359	.362	.376	.377	.457	.457	.539	.540
Vertical integration indicator	.486* (.046)	.426* (.047)	.399* (.046)	.370* (.047)	-.191* (.096)	-.192* (.096)	-.208* (.098)	-.204* (.097)
Multiunit indicator		.160* (.025)		.088* (.028)		-.002 (.064)		-.087 (.082)
TFP			.548* (.052)	.538* (.052)			.487* (.076)	.488* (.077)
	INTEGRATED VS. UNINTEGRATED ENTRANTS				INTEGRATED ENTRANTS VS. UNINTEGRATED INCUMBENTS			
	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Observations	2,771	2,771	2,025	2,025	7,490	7,490	6,104	6,104
R^2	.543	.544	.601	.602	.367	.367	.392	.392
Vertical integration indicator	.324* (.119)	.276* (.122)	.240 (.134)	.212 (.136)	.042 (.086)	.042 (.086)	-.070 (.012)	-.070 (.012)
Multiunit indicator		.127 (.077)		.082 (.089)		NA		NA
TFP			.592* (.132)	.583* (.132)			.533* (.068)	.533* (.068)

NOTE.—The table shows the results from regressing plant-level physical output on a number of variables. Vertical integration (multiunit) indicator is equal to one if the plant is in a vertically integrated (multiunit/multiplant) firm and zero otherwise; TFP is the plant's quantity-based total factor productivity. See the text for details. In the vertically integrated entrants vs. non-vertically integrated incumbents comparison, there are no observations of new multiunit entrants that are not also vertically integrated entrants. All regressions include market-year fixed effects.

* Significant at the 5 percent level.

Section V we will discuss why vertical integration, productivity, and firm size might be linked in these industries.

B. Does Vertical Integration Reduce Quantities?

We use specifications similar to those above to investigate how vertical integration affects market quantities.

We start again by documenting the plant-level patterns. Table 8 repeats the comparisons of table 4, except using plants' physical outputs (logged cubic yards of ready-mixed concrete). Column 1 shows that integrated ready-mixed producers make considerably more output than their un-integrated competitors in the same market. However, as seen in column 5, the output of ready-mixed plants that become integrated actually *shrinks* (some 18 percent) relative to that of unintegrated continuers. This is counterbalanced, though, by the building of new plants by the

integrated firm; as can be seen in the estimates in column 9, new integrated plants are larger than plants of unintegrated producers.²⁷

When we add controls for multiunit status and the plant's TFP, both are associated with higher plant output, with multiunit status sometimes statistically significant and TFP always so. Further, in both the overall and entrant-only comparisons, their inclusion reduces the implied quantity impact of vertical integration. This is qualitatively similar to their effect in the price comparisons above. In the broadest comparison (col. 4), however, there is still a statistically distinct quantity difference between vertically integrated plants and other multiunit plants, even when we control for TFP. Including the controls has little impact on how much newly integrated firms' outputs fall relative to plants that remain unintegrated, even though productivity gains are strongly correlated with output growth (col. 8).

For the market-level analysis, we replace the weighted-average price as the dependent variable with the total (logged) physical quantity of ready-mixed concrete sold in the market.²⁸ Just as with the market-level price regressions, we control for year effects, the market's HHI, and local demand as measured by logged construction employment (levels rather than density here). Again we do not report these coefficients for space reasons; HHI is significantly negative and demand, not surprisingly, significantly positive in every specification. As with the price regressions, we estimate several specifications that include our benchmark estimations and those that address possible endogeneity of firm and market structure.

The results are presented in tables 9 and 10. In the benchmark specifications in table 9, greater vertical integration measured by either the market share or the number of vertically integrated firms is associated with higher output levels in the local concrete market. This again reflects

²⁷ Note that there is a less direct connection between plant- and market-level outcomes in quantity terms than with prices. Firms' total quantities produced in a market can change either intensively, through the amount produced by individual plants, or extensively, with the number of plants the firm owns. The expected output change due to a new integrated firm entering a market depends on the size of the output reductions at its newly acquired plants relative to the output added by its newly built plants. In our sample, the net revenue changes when a new vertically integrated firm comes into a market are on average positive. This is consistent with the market-level evidence below. We discuss further in Sec. V why an integrated producer in the same market over the same period might purchase and shrink existing plants while also building new ones.

²⁸ Here we must be particularly careful with regard to measurement issues. Because the product supplement of the *Census of Manufactures* does not cover all plants, we cannot directly calculate total market quantity simply by aggregating plants' reported quantities. We therefore use plant revenue instead. Yet revenue data alone are problematic because they incorporate price variation across plants: because our price deflator is not market specific, unadjusted revenue overstates (understates) output in markets with higher (lower) than average prices. Our solution is to measure total quantity in a market as market revenue divided by the market's quantity-weighted average price that we use above.

TABLE 9
VERTICAL INTEGRATION AND READY-MIXED CONCRETE QUANTITIES: MARKET-LEVEL RESULTS
BENCHMARK SPECIFICATIONS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Observations	1,870	1,870	1,870	1,550	1,550	1,870	1,870	1,870	1,550	1,550
R^2	.781	.904	.907	.923	.926	.785	.905	.908	.924	.927
Market share of vertically integrated firms	.688*	.335*	.238	.068	-.011					
Market share of multiunit firms	(.133)	(.143)	(.137)	(.150)	(.143)					
			.332*		.324*					
			(.080)		(.081)					
Number of vertically integrated firms						.190*	.121*	.085*	.060	.024
						(.025)	(.033)	(.031)	(.032)	(.029)
Number of multiunit firms								.057*		.062*
								(.010)		(.011)
Quantity-weighted average TFP				.380*	.367*				.381*	.383*
				(.088)	(.084)				(.087)	(.084)
Market fixed effects?	No	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes

NOTE.—The table shows the coefficients obtained by regressing the total (logged) quantity of ready-mixed concrete sold in a market on either the market share or the number of vertically integrated firms operating in the market. The market share or number of multiunit firms and the quantity-weighted average TFP in the market are also included in some specifications. All regressions control for the HHI and density of demand in the market as well as year effects (coefficients not reported). Standard errors are clustered by market.

* Significant at the 5 percent level.

TABLE 10
VERTICAL INTEGRATION AND READY-MIXED CONCRETE QUANTITIES: MARKET-LEVEL RESULTS
ROBUSTNESS CHECKS

	ANTITRUST AND SPATIAL DEMAND INSTRUMENTAL VARIABLES				MULTIPLE-MARKET MERGERS				ADDING STATE-YEAR EFFECTS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Observations	1,870	1,870	1,550	1,550	142	142	133	133	1,870	1,870	1,550	1,550
R^2	NA	NA	NA	NA	.829	.836	.858	.863	.923	.929	.942	.944
Market share of vertically inte- grated firms	.710 (.812)	-.643 (1.063)	.228 (.744)	.416 (.952)	.490 (.425)	.348 (.406)	.407 (.400)	.292 (.375)	.403* (.179)	.273 (.172)	.168 (.210)	.056 (.203)
Market share of multiunit firms		1.527* (.759)		-.311 (.914)		.634 (.379)		.507 (.353)		.366* (.089)		.362* (.090)
Quantity-weighted average TFP			.331* (.065)	.339* (.072)			.564* (.232)	.572* (.230)			.350* (.102)	.339* (.097)

NOTE.—The table shows the coefficients obtained by regressing the total (logged) quantity of ready-mixed concrete sold in a market on either the market share or the number of vertically integrated firms operating in the market. The market share of multiunit firms and the quantity-weighted average TFP in the market are also included in some specifications. All regressions control for the HHI and density of demand in the market as well as year effects (coefficients not reported). The various identification approaches, listed in the header row, are detailed in the text. Standard errors are clustered by market.

* Significant at the 5 percent level.

the patterns seen in table 1. This is true in the cross section, as seen in columns 1 and 6, and within markets over time, as in columns 2 and 7. The estimates imply that going from a completely unintegrated market to one in which integrated firms have a 31.5 percent share implies a rise in total quantities sold of roughly 10 percent. Alternatively, going from zero to one vertically integrated firm in the concrete market corresponds to a quantity increase of about 12 percent.²⁹ Similar patterns are seen in response to changes in integration within markets. Columns 3 and 4 of the table show the results when market fixed effects are included. Output increases significantly when a market becomes more integrated.

When we control for the market share (or number) of multiunit firms and the quantity-weighted average TFP in the market, again we see, as with the market-level price tests above, that these measures explain the link between market outcome and the extent of vertical integration. The multiunit measure always enters as significantly positive and substantially reduces the magnitude of the coefficient on the extent of integration measure. Average TFP follows the same pattern. When included together, these controls drive the coefficient on the vertical integration measure essentially to zero. Unlike the average price results, where average TFP explained most of the relationship between integration and prices, both multiunit firms and average TFP have independent, significant correlations with total market quantities.

In the robustness checks in table 10, we obtain a positive point estimate of the effect of vertical integration on total market quantities in all but one of the specifications. As in table 9, we also find that adding the TFP and multiunit controls reduces the point estimate on vertical integration. Statistical significance, however, is obtained less readily. The antitrust enforcement and spatial demand instrumentation strategy yields imprecise point estimates (see the discussion above). In the multimarket merger specification, while the coefficients on the extent of integration measure are all positive and with a size similar to that of the benchmark results, they are not significant. (The multimarket merger specifications that used the number of vertically integrated firms

²⁹ We note that the above argument that market-level construction activity is likely exogenous to concrete prices is still consistent with a nonzero quantity response to changes in concrete prices. To see why, note that the exogeneity argument implies that the derivative of construction output with respect to concrete prices is roughly zero. This derivative can be restated as the product of two derivatives: construction output with respect to concrete output and concrete output with respect to concrete prices. Thus even when the latter is nonzero (i.e., concrete quantities are responsive to concrete prices), the total derivative is roughly zero if the former derivative is small enough. This former derivative is determined by the construction sector's production function. The very small cost share of concrete in total construction costs (around 2 percent in the data tables of the Benchmark Input-Output Accounts published by the BEA) suggests that this derivative is in fact quite small.

to measure the extent of integration, not reported here for brevity, did yield significantly positive estimated impacts in the regression without multiunit and TFP controls. The estimates also become smaller and insignificant once we added these controls.) Finally, the specifications that control for state-year fixed effects follow the benchmark results quite closely, qualitatively and quantitatively.

In sum, these quantity tests, as with the price tests above, give no indication that vertical integration in the cement and ready-mixed industries facilitates the exercise of market power. Vertically integrated plants sell more and charge lower prices on average than their unintegrated competitors in the same markets. The net quantity impact of the entry of a vertically integrated firm into a market is positive because the decline in output from acquired plants is compensated by additional output from newly built plants. Furthermore, at the market level, increases in the extent of integration in a market are associated with greater total quantities sold. As with the price results, any such vertical integration effect does not seem attributable to vertical integration itself, but rather to concomitant increases in the presence of larger, more efficient firms in the market. (As we will detail below, this appears to arise because such firms are able to harness logistical efficiencies.) There is also no evidence that foreclosure impacts exist but are simply being empirically swamped by efficiency gains.

C. Do Integration's Effects Differ across Markets with Different Ex Ante Foreclosure Potential?

The tests above assume that vertical integration has common price and quantity implications for all markets. However, the model exposited earlier suggests that foreclosure potential may be linked to upstream and downstream market structure. The upstream firm's ability to leverage market power gains by integrating depends in part on the extent of its preintegration market power. Further, the heightened vertical antitrust enforcement regime during part of our sample may also have limited upstream producers' foreclosure abilities during that period. We test in this subsection whether there are noticeable differences in price and quantity effects across four types of markets that theory suggests offer greater ex ante foreclosure potential.

The first specification takes the model's polar case quite literally. It tests whether price and quantity impacts are systematically different in ready-mixed markets that both (a) are served by a single, vertically integrated cement firm and (b) have a small number of ready-mixed firms—specifically, fewer than five. Market power gains are presumably largest in these markets since unintegrated ready-mixed producers are more vulnerable to being foreclosed. The second specification is a looser

version of this cement monopoly test. We allow the integration coefficients to differ depending on whether the local cement industry HHI is above or below its mean value in our sample.³⁰ The third variation allows differential impacts in markets in which cement imports, the availability of which could limit foreclosure power, are less likely to be a viable alternative for cement buyers. We consider imports to be most easily available in EAs containing a port that was one of the 25 largest reception points for imported cement (using 2002 data from the U.S. Customs Service) *and* if the year was 1977 or later, since before that time cement imports were a trivial share of national consumption.³¹ Our fourth and final market breakdown takes advantage of the antitrust enforcement policy variation discussed above and allows different price and quantity effects during census years other than 1967, 1972, and 1977, the period of the strictest stance toward vertical mergers by U.S. antitrust authorities. If enforcement or potential enforcement limits foreclosure possibilities, we should be more likely to see foreclosure-type price and quantity responses during the lax enforcement period.

Table 11 shows the results of this exercise. Each set of estimates was obtained using the benchmark specification and sample of the previous two subsections, except here the price and quantity effects are allowed to vary depending on whether or not one of the conditions above is met. The “interacted” markets in each specification are chosen so that foreclosure effects should be more likely in those markets than in the remainder of the sample.³²

When integration coefficients are allowed to differ in markets with

³⁰ For these two specifications we consider cement plants’ service areas to be BEA economic areas (described above) rather than the much larger USGS cement markets. Cement monopolies are very rare in the USGS markets because of the way they are constructed, and furthermore the model’s assumed transport cost–driven advantage of local cement monopolists is likely to be much stronger when considered within the smaller market definition. We also realize that the upstream market structure measures used here are themselves endogenous outcomes that might have been influenced by the vertical structure itself rather than simply capturing pre-existing conditions. However, the variation in foreclosure potential across markets in our two other specifications below is essentially exogenous.

³¹ While imported cement may not always be consumed in the market in which it is brought to port, the fraction moved to other markets is likely to be small. Land transport of cement is expensive relative to its value, and “ports” in the Customs data include interior cities on navigable waterways, such as St. Louis, Minneapolis, and Cleveland. One possible factor that would counteract the ability of cement consumers to substitute imports is the inclusion in imports of clinker (cement in an intermediate, postkiln, pellet-like stage) that is purchased and then ground into cement and sold by domestic cement firms. Unfortunately, the data do not reveal importers’ identities.

³² Here we use the specification without market fixed effects to highlight differences across markets with differing foreclosure potential. All specifications control for the HHI in the ready-mixed concrete market, demand density (price regressions) or demand (quantity regressions), and year effects. Only specifications using the market share measure of the extent of vertical integration are shown for space reasons, though similar results are obtained using the number of integrated firms.

TABLE 11
MARKET EX ANTE FORECLOSURE POTENTIAL, VERTICAL INTEGRATION, AND CONCRETE PRICES AND QUANTITIES

	VERTICALLY INTE- GRATED CEMENT MONOPOLIST		CEMENT HHI ABOVE MEAN		LOW IMPORT PRESENCE		LIGHT ANTITRUST EN- FORCEMENT PERIOD	
	Quantity- Weighted Average Price (1)	Total Quantity (2)	Quantity- Weighted Average Price (3)	Total Quantity (4)	Quantity- Weighted Average Price (5)	Total Quantity (6)	Quantity- Weighted Average Price (7)	Total Quantity (8)
R^2	.087	.783	.087	.784	.096	.783	.088	.782
Vertical integration intensity	-.132* (.035)	.769* (.161)	-.125* (.029)	.758* (.134)	-.085 (.058)	.551* (.251)	-.107* (.032)	.884* (.205)
(Vertical integration intensity) × indicator	.022 (.050)	-.158 (.207)	.008 (.126)	-.583 (.351)	-.057 (.063)	.201 (.266)	-.030 (.043)	-.322 (.206)
Implied effect in interacted markets	-.110* (.041)	.611* (.172)	-.117 (.124)	.175 (.329)	-.143* (.029)	.752* (.145)	-.137* (.037)	.562* (.131)

NOTE. — This table shows market-level price and quantity regressions in which the coefficient on the extent of integration measure is allowed to vary for differing types of markets. All specifications are estimated from a sample of 1,870 CEA market-years. All regressions control for the HHI and density of demand in the market as well as year effects (coefficients not reported). See the text for details. Standard errors are clustered by market.

* Significant at the 5 percent level.

vertically integrated monopolist cement firms and few ready-mixed firms (cols. 1 and 2), the implied price effects are not significantly different in these markets than in the rest of the sample. The negative price-integration and positive quantity-integration relationships remain even in those markets that most closely match the polar case of the model above. (Both of their magnitudes are about 20 percent smaller than in the rest of the sample, but the differences are insignificant.) When this comparison is broadened to all markets with above-average cement HHIs, again there are no statistically significant differences in the price- and quantity-integration correlations. However, a combination of the magnitude of the interaction terms and their imprecision leads to the implied price and quantity effects of vertical integration being insignificant—and in the case of quantities, considerably smaller in size—in markets with high cement HHIs. In the third set of interactions, markets in which importing cement is unlikely to be an option still see negative (positive) and significant associations between average prices (total quantities) and the share of vertically integrated firms. The magnitudes of the coefficients are slightly larger in no-import markets, but again the differences are insignificant. Finally, when we allow differential price and quantity effects during periods of lax antitrust policy attitudes toward vertical mergers, we do not see foreclosure's implications surfacing in these periods. The relationships between market-level outcomes and the extent of integration remain the same in these years as during the tougher enforcement regime. The implied price effect is slightly larger, the quantity effect slightly smaller, but neither significantly so.

Therefore, in final-goods markets that theory indicates are a priori more susceptible to foreclosure, we still see a negative correlation between integration and prices and a positive correlation between integration and quantities. In none of the specifications is the magnitude of the negative price-integration relationship statistically different in "high-foreclosure potential" markets from the remainder of the sample. In three of four specifications, integration does appear to be associated with smaller (though still positive) quantity gains in such markets, but this difference is statistically weak.

D. Does Vertical Integration Force Exit of Unintegrated Producers?

Foreclosure models predict that firms can use leverage afforded by vertical integration to force unintegrated rivals out of business in both the upstream and downstream industries. We test this implication in this subsection.

The first test directly compares exit probabilities across integrated and unintegrated producers. We regress an indicator for plant death by the next census (i.e., it does not exist in the following census) on

TABLE 12
INTEGRATION AND EXIT PROBABILITIES: LIKELIHOOD OF EXIT ACROSS INTEGRATED AND UNINTEGRATED PLANTS

	CONCRETE ($N = 24,042$)				CEMENT ($N = 1,010$)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
R^2	.091	.097	.092	.098	.446	.472	.447	.473
Vertical integration indicator	-.053* (.014)	-.004 (.014)	-.050* (.014)	-.003 (.014)	-.083* (.036)	.005 (.036)	-.082* (.036)	.005 (.036)
Multiunit indicator		-.087* (.007)		-.084* (.007)		-.233* (.054)		-.232* (.054)
TFP (revenue-based)			-.075* (.013)	-.058* (.013)			-.033 (.051)	-.023 (.050)

NOTE.—The table shows results from linear probability models in which the dependent variable is an indicator equal to one if the plant exits by the next *Census of Manufactures*. Market-year fixed effects are included in all regressions.

* Significant at the 5 percent level.

an integration status indicator as well as a full set of market-year fixed effects.³³

Table 12 contains the results for both concrete and cement plants. Integrated plants in both industries are less likely to exit than their unintegrated competitors in the same market-year. The implied differences are nontrivial: integrated concrete plants have a 5.3-percentage-point lower probability of exit over a five-year intercensus period and cement plants an 8.3-percentage-point lower probability. As points of comparison, the unconditional intercensus exit rates across all plants are 30.3 percent in ready-mixed concrete and 16.4 percent in cement. Unlike the price and quantity results, these exit patterns are consistent with foreclosure.

We add the multiunit indicator and plant TFP controls as shown in columns 2–4 and 6–8.³⁴ Both are negatively correlated with exit in both industries. Interestingly, here multiunit status has greater influence on the vertical integration coefficient than TFP does. It is significant in every specification, and controlling for it alone is enough to make vertical integration's correlation with exit fall to zero in both industries. When we control for both factors simultaneously, the vertical integration indicators for both concrete and cement plants remain statistically insignificant and have small magnitudes.

So again we see that controlling for multiunit status and TFP explains

³³ We do not count as exits plants that leave the cement or concrete industries but continue to operate in a different industry. The exits with which we concern ourselves—complete disappearances of plants from the *Census of Manufactures* universe—correspond to the cessation of economic activity at the former plant's location. We have also run Probit exit models and found results similar to those reported here.

³⁴ Note that in these regressions, we use revenue-based TFP rather than the physical output-based measure used above, since we do not have physical output for cement plants. Moreover, revenue-based TFP may be the better control variable here since it captures both technology and demand fundamentals, both of which affect survival.

TABLE 13
 INTEGRATION AND EXIT PROBABILITIES: MARKET-LEVEL EXIT
 RATES OF NON-VERTICALLY INTEGRATED PLANTS

	Concrete	Cement
Observations	1,725	126
R^2	.265	.537
Market share of vertically integrated firms	.033 (.061)	.084 (.163)

NOTE.—The dependent variable is constructed as the number of non-vertically integrated exiting plants divided by the average number of non-vertically integrated plants in the market over the last two censuses. Market fixed effects are included in both specifications. Standard errors are clustered by market.

away any impact of vertical integration not connected to size or productivity. Vertically integrated plants are more likely to survive not because of any particular facet of being integrated, but instead because they are in large firms and have higher productivity levels. This is inconsistent with foreclosure, which implies that integration itself has a survival benefit. A productivity-based story of vertical merger effects, on the other hand, does not predict a differential impact of integration except through its correlation with productivity and size.³⁵

We next look at how market-level vertical integration affects exit of unintegrated producers. We regress the exit rate of unintegrated producers in a market on the extent of integration, controlling for market and year fixed effects, the industry's HHI in the local market, and demand growth between the present and future censuses.³⁶ The results are in table 13.

The extent of integration is positively related to unintegrated producers' exit rates in both industries, reflecting in part the differences in likelihood of across-plant exits seen above. However, the link here is statistically insignificant and economically small: in concrete (cement), raising the market share of vertically integrated producers from

³⁵ These results are consistent with earlier empirical evidence. Many studies have found that more productive plants are less likely to exit. See, e.g., Baily, Hulten, and Campbell (1992), Olley and Pakes (1996), Aw, Chen, and Roberts (2001), and several of the papers reviewed in Bartelsman and Doms (2000). Further, a separate benefit of size for survival was observed in Dunne, Roberts, and Samuelson (1989), perhaps because larger firms have a greater ability to raise financial capital internally or because firm size acts as a sort of "permanent TFP" proxy for the firm, capturing the long-run component of productivity that is conflated with noise in the contemporaneous measure.

³⁶ For the sake of symmetry with our entry rate analysis below, we construct the market-level exit rate by dividing the number of exiting unintegrated plants by the average number of unintegrated plants in the market in the current and future censuses rather than just the count in the current period. For small percentage changes, both measures are closely comparable. We also weight observations by market demand so as not to give undue influence to markets with very small numbers of plants, where even a single plant exit can create a large exit rate. We do not report coefficients on the controls for space reasons; concrete HHI is negatively related to exit rates, but the results are otherwise unremarkable.

zero to one—that is, across the extremes of the data—implies only a 3.3- (8.4-) percentage-point increase in the exit rate of unintegrated producers. This is a fraction of the average plant-level exit rates above, especially for concrete. Controlling for multiunit status and average TFP in the market makes little difference in the estimate of the vertical integration coefficient, and the controls themselves have insignificant impacts on exit rates; we suppress these results for space reasons.

We see that, as foreclosure predicts, being vertically integrated is correlated with a lower exit probability for a plant in either industry. However, this relationship is driven by the fact that integrated plants are in larger firms and have higher productivity levels. Once this is controlled for, there is no longer per se a survival benefit of vertical integration. Further, increases in the extent of integration within a market do not lead to significantly higher exit rates among unintegrated producers in that market.

E. Does Vertical Integration Reduce Entry?

Many foreclosure models, as well as the FTC report on cement-concrete integration, imply that vertical mergers can lead to higher entry barriers. This prediction has strong implications for long-run welfare, since giving firms the ability to preserve their market power by restricting the entry of new competitors is socially harmful. We explore the link between entry and vertical integration here.

We compute two types of entry rates for each market. One is plant-based and is calculated by dividing the number of new plants in the market by the average number of plants over the current and previous censuses.³⁷ The second entry rate is firm-based. Here definition issues are more complex. Unlike plants, which are geographically unique, firms can operate in several markets. We must therefore take a stand on how to treat entry of existing firms into new geographic markets. For example, if a company that has operated in Texas for 15 years builds a plant in Ohio, is that entry? We consider it as such because of the highly geographically segmented nature of these industries' product markets. Note that this entry could occur not only by the firm building a new plant but also by its purchase of an existing one. That is, a market could experience firm entry but not plant entry. Firm entry rates for a

³⁷ Again we use the average number of establishments in the denominator rather than just the prior value. This keeps small markets that would otherwise have infinite or extremely high computed entry rates from distorting the results. This transformation preserves relative orderings. We also weight by market demand in our regressions to further adjust for the small-market effect.

market are the number of entry episodes divided by the average number of firms operating in the market across the current and past censuses.³⁸

The empirical specification is similar to the market-level exit tests above: we regress entry rates in a market-year, which are based on entry between the past and current censuses, on the market share of vertically integrated establishments in the previous census. As in the exit regressions, we control for market and year fixed effects as well as market demand growth over the period to capture short-run entry impacts of demand fluctuations. We estimate specifications using both overall entry rates, which include entry by integrated and unintegrated producers, and entry rates of unintegrated producers alone.

The results are presented in table 14, with panel A showing results for overall entry rates and panel B for unintegrated producers. There are no significant connections between entry rates and the market share of vertically integrated firms. The economic magnitude of the vertical integration coefficients is also small. As with the market-level exit rate results, even a move from one extreme to the other—that is, a market shifting from a situation in which integrated producers are completely absent to one with 100 percent market share—implies small entry rate changes compared to typical values in the data. For example, even the largest implied effect, that on the plant-level entry rates of unintegrated cement producers, moves entry rates by less than half of their mean and a third of their standard deviation. Moreover, the coefficients are not even negative for concrete. In unreported specifications for the concrete industry in which we add multiunit firms' market shares, these shares are tied to lower entry rates (and are typically significant). However, this also made the implied vertical integration effect *more* positive.

F. Discussion

The evidence suggests that foreclosure effects are not quantitatively important when cement and ready-mixed producers vertically integrate. More extensive integration in a market is not tied to higher prices and

³⁸ There is an additional data difficulty regarding computation of firm-level entry rates. While firm identifiers in the census remain constant if the plant-owning entity does not change, there is an exception to this. Single-unit and multiunit firms are numbered according to different systems, so when a plant in a single-unit firm becomes part of a multiunit firm, its firm identifier changes from one format to the other. This presents no measurement problems when a formerly single-unit plant is bought by a previously multiunit firm; the plant simply assumes the identifier of the purchasing firm. However, if the plant is purchased by a firm that was previously only a single-unit firm or if the plant's owner itself becomes multiunit by purchasing or opening another plant, the plant's firm identifier could change even though no new firm actually entered the market. We therefore count as firm entry episodes only those instances in which a firm identifier changes *and* the plant either was part of a multiunit firm in the previous census or was previously a single-unit firm and its new firm identifier existed in the previous census.

TABLE 14
INTEGRATION AND ENTRY RATES

	CONCRETE		CEMENT	
	Plant-Level Entry Rate	Firm-Level Entry Rate	Plant-Level Entry Rate	Firm-Level Entry Rate
A. Overall Entry Rate (Number of Entrants Divided by Average Plants/Firms in the Market over the Last Two Censuses)				
Observations	1,725	1,725	127	127
R^2	.321	.354	.282	.423
Market share of integrated firms in base year	.046 (.068)	.025 (.090)	-.024 (.127)	-.024 (.152)
B. Unintegrated Producer Entry Rate (Number of Unin- tegrated Entrants Divided by Average Plants/Firms in Market over the Last Two Censuses)				
Observations	1,725	1,725	127	127
R^2	.330	.366	.234	.332
Market share of integrated firms in base year	.067 (.064)	.051 (.071)	-.073 (.106)	-.043 (.115)

NOTE.—This table shows the impact of integration on entry rates. Panel A shows the results of regressing overall plant- and firm-level entry rates in a market (see the text for definitions) on the market share of integrated firms in the base year. Panel B shows similar results for entry rates of unintegrated producers alone. Observations are weighted by market demand. (Similar patterns were observed using lagged number of vertically integrated firms as the explanatory variable, not reported here.) Standard errors are clustered by market.

lower quantities, but rather the opposite. These patterns are robust to numerous empirical specifications and sources of identification. Even when we focus on markets in which one would expect foreclosure opportunities to be greater, the negative correlation between integration and prices remains. Exit rates are lower for integrated plants in both the upstream and downstream industries, as foreclosure would predict, but we find that this can be explained by the fact that integrated plants are more productive and are in larger firms, not because of a particular effect of vertical integration itself. Finally, entry rates do not fall when markets become more integrated.

It seems instead that the data are consistent with a story in which any effects of vertical integration on outcomes arise because integrated firms are larger and, especially, more productive. As these more efficient producers expand their presence in a market, they pass on part of their cost advantages to their customers in lower prices. This reduces average prices in the market directly and also induces their higher-cost unintegrated producers to lower their prices.³⁹ Lower prices in turn increase quantities sold. Those unintegrated producers that cannot remain profitable when facing greater low-cost competition exit.

³⁹ A negative correlation between price and productivity levels has also been documented in previous empirical work (e.g., Roberts and Supina 1996; Eslava et al. 2004; Foster et al. 2005; Syverson, forthcoming).

TABLE 15
DIFFERENCES BETWEEN INTEGRATED AND UNINTEGRATED PRODUCERS

	Labor Productivity	Real Revenue	Total Employment	Total Hours	Nonproduction Worker Ratio	Capital- Labor Ratio
A. Ready-Mixed Concrete						
Observations	31,061	31,770	31,485	31,075	31,480	18,703
R^2	.217	.213	.168	.166	.181	.331
Vertical inte- gration indicator	.292* (.019)	.280* (.034)	.015 (.033)	.047 (.034)	-.018* (.007)	.099* (.030)
B. Cement						
Observations	1,476	1,484	1,480	1,476	1,480	902
R^2	.552	.56	.584	.574	.492	.546
Vertical inte- gration indicator	.116* (.046)	.820* (.106)	.652* (.097)	.648* (.098)	-.034* (.012)	.284* (.098)

NOTE.—This table reports differences in key dependent variables (listed at the head of each column) across integrated and unintegrated producers in the ready-mixed concrete and cement industries. The reported coefficients are those for an indicator variable denoting that a plant is in a vertically integrated firm. Market-year fixed effects are included in all specifications.

* Significant at the 5 percent level.

V. Why Do Integrated Producers Have Higher Productivity Levels?

The fact that the productivity advantage of vertically integrated producers, rather than integration itself, can explain the patterns in the data raises its own issue. What efficiency gains are harnessed through integration, or is the connection merely a correlation? We look more closely at how integrated producers differ from their unintegrated counterparts to try to answer this question.

A. How Are Integrated Producers Different?

We first compare a number of plant-level technology measures across integrated and unintegrated producers. As with the plant-level price, TFP, and quantity comparisons above, we regress the variable of interest on an indicator for integration status and a full set of market-year fixed effects. This again focuses on comparisons among integrated and unintegrated producers in the same market.

These comparisons for each industry's plants are shown in table 15. Differences among ready-mixed concrete plants are in panel A and among cement producers in panel B.⁴⁰ Notice first the labor productivity differences. Integrated ready-mixed producers' labor productivity levels

⁴⁰ The sample size differences across dependent variables reflect variation in the availability of the underlying production data necessary to compute the measures. See the discussion in the Data Appendix.

are 33 percent (29 log points) higher than those of their unintegrated counterparts in the same market. This is considerably larger than the TFP differences seen in table 5 and reflects that fact that, as can also be seen in the table, integrated plants are more capital intensive. In the cement industry, integrated producers also have higher labor productivity levels, by just under 10 percent, than unintegrated cement plants in their market.

Besides being more productive, integrated ready-mixed plants do a larger volume of business than unintegrated plants in their markets; their real revenue is on average about 30 percent higher. Despite their higher output, integrated ready-mixed plants do not employ significantly more labor (workers or hours), as manifested in the higher labor productivity result. There are differences in labor *composition*, however: non-production workers constitute a smaller share of the labor force in integrated plants.⁴¹ This suggests that integrated producers may be able to harness gains coming from the consolidation of administrative duties, a point to which we shall return below. In a comparison of cement plants, integrated producers are also considerably larger in terms of total sales, have a smaller share of nonproduction workers, and have higher capital-labor ratios than their unintegrated counterparts. Unlike concrete plants, they also hire more labor.

Integrated plants are indeed different. However, these patterns are less useful for determining whether the differences are caused by vertical integration. As above, we look at additional comparisons that are less likely to confound the causal impacts of vertical integration with selection on pre-existing heterogeneity. Panel A of table 16 compares plants that were unintegrated in the previous census but became integrated to those that remained unintegrated (we control for market-year fixed effects in this and all these comparisons). Those becoming integrated in the interim had 10 percent faster labor productivity growth than their market counterparts that remained unintegrated. As in the physical quantity comparisons above, sales revenue at these newly integrated ready-mixed plants shrink (at least in relative terms; the reported coefficients give relative growth rates).⁴² Also as previously mentioned,

⁴¹ The U.S. Census Bureau defines production workers (all other employees are considered nonproduction workers) as “workers (up through the line-supervisor level) engaged in fabricating, processing, assembling, inspecting, receiving, packing, warehousing, shipping (but not delivering), maintenance, repair, janitorial, guard services, product development, auxiliary production for plant’s own use (e.g., power plant), record keeping, and other closely associated services (including truck drivers delivering ready-mixed concrete). Exclude proprietors and partners” (form MC-31000).

⁴² While in the absence of randomized integration at the plant level we cannot be sure that these estimates reflect exclusively causal effects, our productivity-based explanation for the results above requires only that vertical integration is correlated with productivity. We also find it difficult to form a reasonable explanation for the observed patterns that

TABLE 16
BECOMING INTEGRATED: READY-MIXED CONCRETE CONTINUERS AND ENTRANTS

	GROWTH OF					
	Labor Productivity	Real Revenue	Total Employment	Total Hours	Nonproduction Worker Ratio	Capital-Labor Ratio
A. Changes among Continuers (Conditioning on Being Unintegrated in Previous Census)						
Observations	15,919	16,358	16,274	15,933	9,166	16,271
R^2	.194	.274	.204	.217	.189	.180
Vertical integration indicator	.105* (.049)	-.399* (.060)	-.396* (.061)	-.439* (.064)	-.030 (.022)	.018 (.076)
B. Integrated Entrants Compared to Unintegrated Entrants						
Observations	7,681	8,005	7,871	7,687	7,870	5,405
R^2	.330	.339	.325	.323	.34	.433
Vertical integration indicator	.336* (.047)	.125 (.079)	-.157* (.073)	-.162* (.078)	-.044* (.014)	.263* (.073)
C. Integrated Entrants Compared to Unintegrated Incumbents						
Observations	18,038	18,310	18,220	18,045	18,217	12,760
R^2	.240	.228	.193	.195	.222	.378
Vertical integration indicator	.358* (.038)	-.281* (.063)	-.510* (.060)	-.514* (.064)	-.061* (.012)	.259* (.058)

NOTE.—This table reports differences in key dependent variables (listed at the head of each column) across integrated and unintegrated producers. Panel A compares growth rates across integrated and unintegrated continuers (plants that survive for two consecutive censuses). Panel B compares integrated and unintegrated entrants (plants appearing in their first census). Panel C compares integrated entrants to unintegrated incumbents. Market-year fixed effects are included in all specifications.

* Significant at the 5 percent level.

though, the decline in sales at acquired concrete plants is on average more than compensated by new entrants. Panels B and C of table 16 compare integrated entrants to, respectively, unintegrated entrants and unintegrated incumbents. Integrated entrants have higher sales than unintegrated entrants, though they hire less labor. New ready-mixed plants in integrated firms are also noticeably more productive than both unintegrated entrants and unintegrated incumbents. As with the comparisons among all producers, integrated entrants have lower nonproduction worker ratios and are more capital intensive than both comparison groups.

Similar comparisons for cement plants—not reported here—found no significant differences among either continuing plants or entrants. This suggests, interestingly, that the differences seen between integrated and unintegrated cement producers are driven by selection into inte-

implies reverse causation, i.e., why integrating firms would explicitly target plants that are expected to increase in productivity *and shrink* in the near future.

gration based on pre-existing differences, whereas some of the differences among ready-mixed concrete plants are more likely to be a result of becoming integrated (or being newly built by an integrated firm). Moreover, vertically integrated firms may have the largest impact on concrete markets in terms of productivity through their newly built plants. We return to these points below.

B. A Specific Mechanism

The findings above, along with a consideration of practical features of the ready-mixed concrete industry, suggest an explanation for the observed integration-productivity link: improved logistical coordination. Operating in the ready-mixed industry requires delivering a perishable product to time-sensitive buyers in multiple locations. Having several plants in a local area and coordinating deliveries through a central office could benefit a firm by consolidating overhead (one dispatcher might handle deliveries from several plants that would each have separate dispatchers in single-unit firms) and allowing more efficient use of available resources through cross-plant substitution.⁴³

This type of productivity gain is consistent with the above findings. Integrated producers' lower nonproduction worker ratios suggest that they have lower overhead labor requirements. The pursuit of coordination efficiencies explains why integrated producers might reshuffle production locations in markets by reducing output at purchased plants and replacing it with production from their newly built plants, perhaps located nearer spatially shifting demand within the market. Such efficiencies are also used to motivate some stated practices of integrated producers. For instance, the large integrated producer Lafarge (a firm that was just starting to develop a sizable presence in the United States at the end of our sample period) describes its preferred placement of ready-mixed operations in its 2004 20-F filing: "We aim to place our ready mix concrete plants in clusters in each micro market in which we operate in order to optimize our delivery flexibility, capacity and backup capability. . . . We evaluate each micro market in which we operate periodically and dismantle and move plants to locations where they can

⁴³ In an earlier version of this paper, we explored in detail two other potential sources of integrated producers' productivity advantage: elimination of double marginalization and improved inventory management and production smoothing. We did not find evidence that either was the source of the observed productivity gap. See Hortaçsu and Syverson (2006) for details.

be used more profitably” (Lafarge 2005, 39). Plant clustering allows firms to better take advantage of any coordination efficiencies.⁴⁴

This raises the question, however, of why *cement* firms need to be the coordinating body. A possible answer is that they need not be. The ready-mixed industry has seen considerable consolidation over our sample—in 1963, 3,999 firms owned 4,621 ready-mixed plants, but by 1997, only 2,898 firms owned 5,252 plants—and much of this consolidation came through horizontal rather than vertical mergers. As seen in table 2 above, the fraction of ready-mixed plants in multiunit firms more than doubled between 1963 and 1997, from 24.8 to 55.6 percent. The corresponding change in the fraction of vertically integrated plants was from 1.8 to 10.6 percent. Clearly, then, the majority of merged ready-mixed concrete plants were folded into firms without cement operations. Coordination and its possible efficiency gains, if they are tied primarily to firm size, have therefore not been exclusive to vertically integrated firms.⁴⁵

We explore this notion further by repeating our earlier comparisons across integrated and unintegrated plants, but this time adding to the specification the total ready-mixed sales of the plant’s firm in the CEA. If logistics coordination benefits are facilitated by having a greater extent of local concrete operations, controlling for the size of the firm’s local ready-mixed sales will let us see whether the relationships observed above between vertical integration and concrete plants’ prices, productivity levels, scale, and technologies are tied to this.

The results of this exercise, shown in table 17, are instructive, especially in comparison to the earlier plant-level price, TFP, and quantity contrasts as well as panel A of table 15. When we control for local concrete sales, the magnitude of vertically integrated plants’ price and

⁴⁴ We note that nothing about these efficiency gains requires spatial proximity between an integrated firm’s cement and concrete operations, or even that the upstream division supply the downstream division at all. Indeed, Lafarge owns about two dozen ready-mixed plants in Colorado and Wyoming, even though its nearest cement distribution terminal is 490 miles away in Oklahoma City. We contacted some of these plants, and they confirmed that they receive all their cement inputs from sources outside of Lafarge. A similar situation exists for Florida Rock Industries, an integrated producer that owns ready-mixed plants in Florida, southern Georgia, Virginia, and Maryland. The firm’s cement facilities, however, are all located in Florida, and the company’s financial filings state that the company’s concrete operations purchased cement from 10 outside suppliers (Florida Rock Industries 2005). Phone calls to management verified complete external cement sourcing among the company’s 30 or so ready-mixed plants in Virginia and Maryland. We also note that these practices are hard to reconcile with foreclosure motives. Foreclosure necessarily implies that integrated downstream producers are at the very least being partially supplied by upstream producers in their firm.

⁴⁵ That is not to say that a large firm in *any* industry could harness coordinating economies, of course. But cement shares a final demand sector with ready-mixed, and the two industries do have other key elements in common, such as the prominence of logistical concerns.

TABLE 17
DIFFERENCES BETWEEN INTEGRATED AND UNINTEGRATED READY-MIXED CONCRETE
PLANTS, CONTROLLING FOR FIRM'S TOTAL SALES IN LOCAL MARKET

	Output Price	TFP	Physical Output	Labor Productivity	Nonproduction Worker Ratio	Capital- Labor Ratio
Observations	12,553	8,555	12,553	31,061	31,468	18,702
R^2	.426	.329	.624	.292	.187	.35
Vertical integra- tion indicator	-.012 (.007)	-.015 (.014)	-.237* (.042)	.097* (.019)	-.039* (.007)	-.037 (.031)
Firm's local ready- mixed sales	-.009* (.002)	.055* (.004)	.633* (.008)	.161* (.004)	.017* (.001)	.113* (.006)

NOTE.—This table reports differences in key dependent variables (listed at the head of each column) across integrated and unintegrated producers in the ready-mixed concrete industry. The reported coefficients are those for a vertically integrated plant dummy and the coefficient on the (log of the) total concrete sales of the plant's firm in the local market. Market-year fixed effects are included in all specifications.

* Significant at the 5 percent level.

productivity differences falls and their significance disappears. In other words, while plants in vertically integrated firms charge lower prices and are more productive on average than unintegrated plants, they do not have significantly different prices or TFP levels than plants in unintegrated firms having equally sized local concrete sales. Recall that most of the plant- and market-level relationships between vertical integration and outcomes above become insignificant once we control for multiunit status and TFP. These results suggest that firm size and TFP are integrally related, specifically, that the firms' *horizontal* size downstream is linked to how efficient its individual concrete plants are. Their productivity advantage in turn allows them to charge lower prices and makes them more likely to survive.⁴⁶ Some differences do remain, however, between integrated and unintegrated plants whose firms have similar downstream sales: integrated plants are smaller (i.e., the same sales are spread across more plants), they still have a labor productivity advantage, and they use less overhead labor.

On the whole, this evidence, both anecdotal and from our data, appears to suggest that *vertical* integration per se might not be the primary source of integrated producers' productivity advantage. Rather, integrated producers—as with any producer with sufficient size and perhaps the necessary complementary managerial talent, even one with a purely horizontal firm organization—seem to be able to take advantage of the

⁴⁶ The strictest interpretation of the logistics coordination story would imply that the large-firm and TFP impacts would be perfectly collinear, since one completely explains the other. We do not observe this in the data, probably because, among other reasons, multiunit status is a noisy measure of a firm in the downstream market, and there are almost surely other firm-specific determinants of productivity that are not related to firm size.

logistical coordination gains from having geographically clustered plants whose logistics are centrally managed.

VI. Concluding Remarks

We have used unusually detailed data to investigate vertical integration in the cement and ready-mixed concrete industries, with a particular focus on the possible foreclosure effects of integration. These industries hold several features that make them an excellent candidate for an empirical case study on the issue. Their downstream markets are highly geographically segmented (especially for ready-mixed, where the vast majority of output is shipped less than 100 miles), providing us with considerable variation to empirically identify effects of interest. Second, cement and ready-mixed are relatively homogeneous in physical attributes and have little brand differentiation. Therefore, the competitive effects we find—on prices, for instance—more likely arise from changes in market structure than from alterations in product mix. Third, we have access to detailed plant-level production information for these industries, affording more variation in vertical market structures (thousands of producers operating in hundreds of local markets over a 34-year time span) than was typically available to previous researchers and allowing us to study interacting effects that previous studies by necessity examined in isolation. We also explore for the first time elements of theoretical models that have not yet (to our knowledge) been studied empirically, including vertical integration's long-run competitive impacts on entry and exit. Finally, vertical integration among producers in these two industries has been the focus of substantial policy attention both in the United States and elsewhere.

We find little evidence that foreclosure takes place when firms integrate in these industries. Instead, prices fall, quantities rise, and entry rates do not change when markets become more integrated (measured by either the market share or the number of vertically integrated firms in the market).

These patterns are consistent with an alternative, efficiency-based mechanism: namely, that greater vertical integration in a market corresponds to the expansion of larger, more productive integrated firms as they take market share from higher-priced, less efficient producers. This alternative explanation is reflected in the fact that, once we control for firm size and productivity impacts, there is very little role for vertical integration itself in explaining plant- and market-level outcomes. We went on to show evidence suggesting that linking integration, firm size, and productivity are gains in logistical coordination efficiencies tied to having multiple concrete plants in the same market. However, these

efficiencies appear not to be tied to vertical integration itself, but rather firm size in the downstream market more generally.

Our results, if correct, point to an interesting trade-off for the social welfare implications of vertical integration. It may be that in the cement and ready-mixed industries, vertical integration is just one way for large, efficient firms to enter a market. In fact there has been extensive consolidation in the ready-mixed industry over the past four decades, and vertical integration has been one way in which industry plants have been folded into larger firms. As we show above, the churning process underlying such consolidations may have had significant impacts on productivity in the industry. This presents again the familiar difficulty of balancing efficiency and market power considerations in merger analysis.

We conclude by emphasizing that one should be careful when trying to generalize the results in this paper to other industries or markets. Vertical integration decisions of firms are driven by many different considerations, of which vertical foreclosure may or may not be important in any particular case. Our study is an attempt to study the short- and long-term effects of the decision to integrate in these particular industries. The fact that we do not find evidence of foreclosure here does not mean that market power gains made possible through vertical integration are not quantitatively important elsewhere.

Data Appendix

We describe here details on the construction of our production variables. We first note that while the *Census of Manufactures* microdata contain enormous amounts of production information, they do not offer full coverage for every variable discussed below. Very small plants (typically with fewer than five employees)—called administrative record establishments—have imputed data for most production variables. Administrative record plants amount to roughly one-sixth of cement establishments and one-third of ready-mixed plants, but because of their small size they constitute a much smaller share of employment (0.6 percent in cement—these are almost surely grinding-only plants without kilns—and 5.1 percent in ready-mixed) and output (0.8 and 4.3 percent in cement and ready-mixed, respectively). Because of the imputations, we exclude administrative record plants from analyses that compare production variables such as productivity, output, or prices. However, we are of course able to use these plants when computing entry and exit rates or integration status. Additionally, not every variable was collected in each census. For example, equipment capital stocks were not collected in 1963 and 1997, making it impossible to compute TFP values and capital-labor ratios during these years. We have removed likely imputes (they are not explicitly flagged) from our sample using methods described in Roberts and Supina (1996) and Syverson (forthcoming).

Labor hours.—Production worker hours are reported directly in the *Census of Manufactures* microdata. This value is then scaled up to total hours by multiplying by the ratio of total employees to production workers. This assumes, in essence,

that average nonproduction worker hours equal average production worker hours within plants.

Capital stocks.—Equipment and building capital stocks are plants' reported book values of each capital type deflated by the book to real value ratio for the corresponding three-digit industry. (These industry-level equipment and structures stocks are taken from published BEA data.) Any reported machinery or building rentals by the plant are inflated to stocks by dividing by a type-specific rental rate.⁴⁷ The total productive capital stock k_{it} is the sum of the equipment and structures stocks.

Real materials and energy use.—Materials and energy inputs are plants' reported expenditures on each divided by their respective industry-level deflators from the NBER Productivity Database.

Labor productivity.—We measure labor productivity as plant output per worker-hour, where output is the real value of shipments and hours are constructed as above.

Total factor productivity (revenue- and physical quantity-based).—We measure productivity using a standard TFP index. Plant TFP is computed as its logged output minus a weighted sum of its logged labor, capital, materials, and energy inputs. That is,

$$\text{TFP}_{it} = y_{it} - \alpha_l l_{it} - \alpha_k k_{it} - \alpha_m m_{it} - \alpha_e e_{it}$$

where the weights α_j are the input elasticities of input $j \in \{l, k, m, e\}$. While inputs are plant-specific, we use industry-level input cost shares to measure the input elasticities.⁴⁸ These cost shares are computed using reported industry-level labor, materials, and energy expenditures from the NBER Productivity Database (which is itself constructed from the *Census of Manufactures*). Capital expenditures are constructed as the reported industry equipment and building stocks multiplied by their respective capital rental rates in cement and ready-mixed concrete's corresponding two-digit industry.

For both industries, we construct TFP measures using plants' reported revenues (deflated to 1987 dollars using price indexes from the NBER Productivity Database) as an output measure. This is the standard measure used in the literature and is the one we use in the investigation of plant exit patterns. (This revenue-based measure is used both so that the results from both industries can be compared and because as explained in Foster et al. [2005], revenue-based productivity measures embody both technology- and demand-side profitability

⁴⁷ Capital rental rates are taken from unpublished data constructed by the Bureau of Labor Statistics for use in computing its multifactor productivity series. Formulas, related methodology, and data sources are described in U.S. Bureau of Labor Statistics (1983) and Harper, Berndt, and Wood (1989).

⁴⁸ This index makes a few implicit assumptions that are worth discussing further. One is that returns to scale are constant. If the scale elasticity were instead different from one, each of the input elasticities α_j should be multiplied by the scale elasticity. In earlier work, Syverson (2004) finds that returns to scale in ready-mixed concrete are essentially constant. Also, using industry-level input elasticities assumes that all plants in the industry use the same technology. While a common assumption in the literature, the fact that we observe different capital intensities at vertically integrated plants could raise questions about whether this is appropriate. However, different input intensities can also result from producers facing different factor prices. In this case, it would imply that integrated producers face lower capital prices than unintegrated producers. Given that integrated plants are owned by firms that are on average substantially larger than those owning unintegrated plants, and as such may have access to easier credit, we find such a factor price difference plausible.

fundamentals that determine exit probabilities.) However, for ready-mixed plants when data are available, we also construct a TFP measure based on the plant's physical output. This removes the influence of within-industry price variation on the output measure. This is the TFP concept used in all other exercises that use only ready-mixed concrete plants.

We must make one adjustment to the output data when computing physical-output TFP. Since ready-mixed plants can produce multiple products (though most do not, as discussed above), but inputs are reported on an establishment-wide rather than product-specific basis, we must impute the share of inputs allocated to ready-mixed production in multiproduct plants. We do so by dividing reported ready-mixed output by its share of total establishment sales. This adjustment in effect assumes that inputs are used proportionately to each product's revenue share. (For example, a plant producing 1,000 cubic yards of ready-mixed concrete accounting for 80 percent of its revenues will have the same TFP as a completely specialized plant producing 1,250 cubic yards with the same measured inputs.)

Output and factor prices.—We use product-level revenue and physical production and consumption data from the *Census of Manufactures* Product and Materials Supplements to compute ready-mixed plants' unit concrete output prices and cement factor input prices. We then adjust these to a common 1987 basis using the corresponding four-digit industry-level shipments deflators from the NBER Productivity Database.

There are two important notes regarding these calculated unit prices. First, the value of shipments (sales revenue) data are collected on a free-on-board basis, that is, exclusive of any shipping costs. Prices should reflect not the delivered cost of the ready-mixed but rather what one could buy it for at the plant gate. Second, the unit prices are annual averages. This can be shown to be equivalent to a quantity-weighted average of all transaction prices charged by the plant during the year. We do not observe product-specific data for administrative record plants, so they are dropped from the analysis as in the core sample. We also remove a small number of gross outliers having prices greater than five times or less than one-fifth the median in a given year and limit the sample to those plants with ready-mixed sales accounting for over one-half of yearly revenues. (This sample criterion is not very restrictive in practice; most ready-mixed producers are specialists.) Finally, we attempt to exclude any non-administrative record plants that have (mostly because of incomplete reporting) physical quantities imputed by the Census Bureau. Unfortunately, these imputes are not flagged. To distinguish and remove imputed product-level data from the sample, we use the techniques described in detail in Roberts and Supina (1996) and Foster et al. (2005).

Imputing construction-sector employment when County Business Patterns data are missing.—County Business Patterns data occasionally have missing observations because of data disclosure regulations. The construction sector's ubiquity and abundance of small firms allow full disclosure of total employment in nearly all counties, however (employment data are withheld in roughly 1.5 percent of the county-year observations). We impute employment when missing by multiplying the number of establishments in each of nine employment ranges (which are always reported) by the midpoint of their respective employment ranges and summing the result. The impact of using imputes is likely to be even less than their proportion indicates, since the typically small nondisclosure counties are less likely to contain ready-mixed plants.

Construction of antitrust policy and spatial demand instruments.—In our instru-

mental variable specifications, we use three sets of instruments. Two use variation in antitrust enforcement across our markets and the other works off spatial patterns of demand (i.e., construction activity) within concrete and cement markets.

The first antitrust policy instrument is an indicator equal to one if the smallest firm in a market has a share of at least 4 percent. This is based on the fact that U.S. antitrust authorities' horizontal merger guidelines have changed over time. During a particularly strict period, any merger involving a firm with more than 4 percent market share was open to challenge. Thus in markets in which the smallest firm was already above that threshold, *any* merger in the market would be suspect. Since the guidelines applied across industries, these guidelines were obviously exogenous to structure in any given ready-mixed concrete market and should be exogenous shifters of horizontal market structure. We interact the indicator with another equal to one during the 1967, 1972, and 1977 censuses since this is the period in which the policy was in effect.

The second antitrust policy instrument comes directly from the FTC's antitrust actions in the cement and ready-mixed industries. In January 1967, in response to industry complaints about the unpredictability of merger challenges, the FTC announced a set of guidelines that promised challenges to any acquisitions by cement producers of any "substantial" ready-mixed concrete company. "Substantial" was defined as a ready-mixed firm that was one of the four largest in its market or one that used over 50,000 barrels of cement per year (*Wall Street Journal* 1967). This rule was dropped by 1977. Since the combined market share of substantial firms plausibly varies across markets for reasons unrelated to equilibrium outcomes (at the very least once exogenous demand is conditioned on), some markets are likely to receive more FTC scrutiny than others simply because of the specifics of the rule. As with the changes in the horizontal merger guidelines, the fact that the rule was in effect only during the 1967–77 censuses also creates intertemporal variation in enforcement likelihoods. We therefore construct for each market-year in our sample the total market share of substantial concrete firms. The first criterion is easily evaluated using plant-level revenue aggregated to the firm level within the market. For the cement use criterion, we employed the census materials supplement. These data allow us to observe the physical quantity of cement purchases made by a subsample of ready-mixed concrete plants. We regressed cement purchases on plant revenues for these plants (which tended to be the largest in the industry) to obtain a total firm revenue value such that firms with greater real sales than this would be expected to buy more than 50,000 barrels of cement. This allowed us to determine which firms met the FTC's second definition of "substantial" even if we did not observe physical cement purchases for the entire firm. Finally, we interact this share of substantial firms with an indicator variable denoting the 1967, 1972, and 1977 censuses as above. This variation in likely FTC scrutiny is used to instrument for the extent of integration.

The third instrument set reflects measures of the dispersion of demand within our sample markets. They are constructed from the county-level construction employment data aggregated to the market level above. We use three measures: the standard deviation in both the (log) level and growth rates of demand across the market's counties, and the HHI of counties' construction employment shares in the market (here, employment is in levels rather than logs). The logic of these instruments is as follows. As with the market-level aggregates, they should be exogenous to ready-mixed market outcomes because of concrete's small share in construction costs. Further, demand (or demand growth) variation across

areas in a market plausibly affects the returns to a firm owning multiple plants. If demand is concentrated in a small space, a firm may be able to serve it sufficiently with a single plant (albeit possibly a large one, depending on the level of demand). But if demand is instead spread spatially, the industry's high transport costs may make it worthwhile to own several plants throughout the market. Both the standard deviation of logged employment and the HHI of employment levels measure the market's spatial demand variation, for a given demand level (which we control for directly in the regressions). The standard deviation in demand growth rates measures the market's demand volatility, which may give firms an incentive to own multiple plants in different locations in a market to diversify away idiosyncratic plant-level risk.

References

- Allen, Bruce T. 1971. "Vertical Integration and Market Foreclosure: The Case of Cement and Concrete." *J. Law and Econ.* 14 (April): 251–74.
- Asker, John. 2004. "Diagnosing Foreclosure Due to Exclusive Dealing." Manuscript, New York Univ., Stern School Bus.
- Aw, Bee Yan, Xiaomin Chen, and Mark J. Roberts. 2001. "Firm-Level Evidence on Productivity Differentials and Turnover in Taiwanese Manufacturing." *J. Development Econ.* 66 (October): 51–86.
- Baily, Martin N., Charles Hulten, and David Campbell. 1992. "Productivity Dynamics in Manufacturing Plants." *Brookings Papers Econ. Activity: Microeconomics*, pp. 187–249.
- Bartelsman, Eric J., and Mark Doms. 2000. "Understanding Productivity: Lessons from Longitudinal Microdata." *J. Econ. Literature* 38 (September): 569–94.
- Bernheim, B. Douglas, and Michael D. Whinston. 2004. "Anticompetitive Exclusion and Foreclosure through Vertical Agreements." Manuscript, Northwestern Univ.
- Bolton, Patrick, and Michael D. Whinston. 1993. "Incomplete Contracts, Vertical Integration, and Supply Assurance." *Rev. Econ. Studies* 60 (January): 121–48.
- Bork, Robert H. 1978. *The Antitrust Paradox: A Policy at War with Itself*. New York: Basic Books.
- Chipty, Tasneem. 2001. "Vertical Integration, Market Foreclosure, and Consumer Welfare in the Cable Television Industry." *A.E.R.* 91 (June): 428–53.
- Coase, Ronald H. 1937. "The Nature of the Firm." *Economica*, n.s., 4 (November): 386–405.
- Dunne, Timothy, Mark J. Roberts, and Larry Samuelson. 1989. "The Growth and Failure of U.S. Manufacturing Plants." *Q.J.E.* 104 (November): 671–98.
- Eslava, Marcela, John Haltiwanger, Adriana Kugler, and Maurice Kugler. 2004. "The Effects of Structural Reforms on Productivity and Profitability Enhancing Reallocation: Evidence from Colombia." *J. Development Econ.* 75 (special issue; December): 333–71.
- Florida Rock Industries. 2005. "10-K Report for 2004." <http://www.flarock.com/flarock/index.asp?bws=72819243&bww=44>.
- Foster, Lucia, John Haltiwanger, and Chad Syverson. 2005. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" Working Paper no. 11555 (August), NBER, Cambridge, MA.
- FTC (U.S. Federal Trade Commission). 1966. *Economic Report on Mergers and Vertical Integration in the Cement Industry*. Washington, DC: U.S. Government Printing Office.

- . 1985. "Cement Industry: Rescission of Enforcement Policy with Respect to Vertical Mergers in the Cement Industry." *Fed. Register* 50 (May): 21507–8.
- Grimm, Curtis M., Clifford Winston, and Carol A. Evans. 1992. "Foreclosure of Railroad Markets: A Test of Chicago Leverage Theory." *J. Law and Econ.* 35 (October): 295–310.
- Grossman, Sanford J., and Oliver D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *J.P.E.* 94 (August): 691–719.
- Harper, Michael J., Ernst R. Berndt, and David O Wood. 1989. "Rates of Return and Capital Aggregation Using Alternative Rental Prices." In *Technology and Capital Formation*, edited by Dale W. Jorgenson and Ralph Landau. Cambridge, MA: MIT Press.
- Hart, Oliver, and Jean Tirole. 1990. "Vertical Integration and Market Foreclosure." *Brookings Papers Econ. Activity* (special issue), pp. 205–76.
- Hastings, Justine S. 2004. "Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California." *A.E.R.* 94 (March): 317–28.
- Hastings, Justine S., and Richard J. Gilbert. 2005. "Market Power, Vertical Integration and the Wholesale Price of Gasoline." *J. Indus. Econ.* 53 (December): 469–92.
- Hortaçsu, Ali, and Chad Syverson. 2006. "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices." Working Paper no. WP-06-21, Center Econ. Studies, Washington, DC.
- Johnson, Kenneth P. 1995. "Redefinition of the BEA Economic Areas." *Survey Current Bus.* 75 (February): 75–81.
- Johnson, Ronald N., and Allen M. Parkman. 1987. "Spatial Competition and Vertical Integration; Cement and Concrete Revisited: Comment." *A.E.R.* 77 (September): 750–53.
- Kreps, David M., and José A. Scheinkman. 1983. "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes." *Bell J. Econ.* 14 (Autumn): 326–37.
- Lafarge. 2005. "Report on Form 20-F." http://www.lafarge.com/lafarge/PUBLICATION/20050407/04072005-publication_finance-20f_2004-uk.pdf.
- Martin, Stephen, Hans-Theo Normann, and Christopher M. Snyder. 2001. "Vertical Foreclosure in Experimental Markets." *Rand J. Econ.* 32 (Autumn): 466–96.
- McBride, Mark E. 1983. "Spatial Competition and Vertical Integration: Cement and Concrete Revisited." *A.E.R.* 73 (December): 1011–22.
- Nocke, Volker, and Lucy White. Forthcoming. "Do Vertical Mergers Facilitate Upstream Collusion?" *A.E.R.*
- O'Brien, Daniel P., and Greg Shaffer. 1992. "Vertical Control with Bilateral Contracts." *Rand J. Econ.* 23 (Autumn): 299–308.
- Olley, G. Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (November): 1263–97.
- Ordover, Janusz A., Garth Saloner, and Steven C. Salop. 1990. "Equilibrium Vertical Foreclosure." *A.E.R.* 80 (March): 127–42.
- Posner, Richard A. 1976. *Antitrust Law: An Economic Perspective*. Chicago: Univ. Chicago Press.
- Rey, Patrick, and Jean Tirole. 2007. "A Primer on Foreclosure." In *Handbook of Industrial Organization*, vol. 3, edited by Mark Armstrong and Robert H. Porter. Amsterdam: North-Holland, forthcoming.

- Riordan, Michael H. 1998. "Anticompetitive Vertical Integration by a Dominant Firm." *A.E.R.* 88 (December): 1232–48.
- Roberts, Mark J., and Dylan Supina. 1996. "Output Price, Markups, and Producer Size." *European Econ. Rev.* 40 (April): 909–21.
- Rosengren, Eric S., and James W. Meehan Jr. 1994. "Empirical Evidence on Vertical Foreclosure." *Econ. Inquiry* 32 (April): 303–17.
- Salinger, Michael A. 1988. "Vertical Mergers and Market Foreclosure." *Q.J.E.* 103 (May): 345–56.
- Shea, John. 1997. "Instrument Relevance in Multivariate Linear Models: A Simple Measure." *Rev. Econ. and Statis.* 79 (May): 348–52.
- Snyder, Christopher M. 1995. "Empirical Studies of Vertical Foreclosure." *Indus. Econ. Conf. Papers and Proc.* (Univ. Melbourne and Bur. Indus. Econ.) 95/23: 98–127.
- Syverson, Chad. 2004. "Market Structure and Productivity: A Concrete Example." *J.P.E.* 112 (December): 1181–1222.
- . Forthcoming. "Prices, Spatial Competition, and Heterogeneous Producers: An Empirical Test." *J. Indus. Econ.*
- U.S. Bureau of Labor Statistics. 1983. *Trends in Multifactor Productivity, 1948–81*. Bull. no. 2178. Washington, DC: Government Printing Office.
- USGS (U.S. Geological Survey). 2004. *Minerals Yearbook*. Vol. 1. *Metals and Minerals*. Washington, DC: U.S. Government Printing Office.
- Vernon, John M., and Daniel A. Graham. 1971. "Profitability of Monopolization by Vertical Integration." *J.P.E.* 79 (July/August): 924–25.
- Wall Street Journal*. 1967. "FTC Announces Formal Merger Guidelines for Cement, Food-Distribution Industries." January 18.
- Waterman, David, and Andrew A. Weiss. 1996. "The Effects of Vertical Integration between Cable Television Systems and Pay Cable Networks." *J. Econometrics* 72 (May): 357–95.
- Williamson, Oliver E. 1971. "The Vertical Integration of Production: Market Failure Considerations." *A.E.R. Papers and Proc.* 61 (May): 112–23.