

Why is Productivity Correlated with Competition?*

Matt Backus[†]

August 29, 2014

Abstract

The correlation between productivity and competition is an oft-observed but ill-understood result. Some suggest that there is a *treatment* effect of competition on measured productivity, e.g. through a reduction of “managerial slack.” Others argue that greater competition makes unproductive establishments exit by reallocating demand to their productive rivals, raising observed average productivity via *selection*. I study the ready-mix concrete industry and offer three perspectives on this ambivalence. First, I model the establishment exit decision to construct a semi-parametric selection correction and quantify the empirical significance of treatment and selection. Second, I use a grouped IV quantile regression to test the distributional predictions of the selection hypothesis. Finally, I look for evidence of a correlation between competition and reallocation using a standard decomposition approach at the market level. I find no evidence greater selection or reallocation in more competitive markets; instead, my results suggest that measured productivity responds directly to competition.

*This is a revised version of the first chapter of my dissertation; the original version is dated November 11, 2011. I am grateful to my advisors Daniel Akerberg, Jeremy Fox, Francine Lafontaine, Natalia Lazzati, and Scott Page for their encouragement and guidance. Thanks also to Jim Adams, Allan Collard-Wexler, Ying Fan, John Haltiwanger, Panle Jia Barwick, Kai-Uwe Kühn, Shawn Klimek, Jagadeesh Sivadasan, Chad Syverson, and seminar participants at Bocconi, Cornell, Duke Fuqua, EARIE 2014, Econometric Society North American and European Summer Meetings 2014, eBay Research Labs, FTC, HBS, IIOC 2014, Iowa, Michigan, SED 2014, Stanford GSB, Temple Fox, Toulouse, University of Buffalo, and Warwick for helpful comments and suggestions. All remaining errors are my own. This paper was written while I was a Special Sworn Status researcher of the US Census Bureau at the University of Michigan Research Data Center. Any opinions and conclusions expressed herein are those of the author and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

[†]Cornell University and eBay Research Labs. E-mail: backus@cornell.edu

1 Introduction

There is a perennial paper in the productivity literature which presents the following result, updated for contemporary innovations in attitudes towards data and econometrics: *firms that are in more competitive markets are more efficient*. This correlation has been identified cross-sectionally across industries (Caves and Barton, 1990; Green and Mayes, 1991), and in panels as well (Nickell, 1996; Hay and Liu, 1997); in the US (Dertouzos et al., 1989) and abroad (Porter, 1990); papers in the trade literature have identified this result using policy changes (Pavcnik, 2002; Sivadasan, 2009), and the correlation remains stark in industry-level studies (Graham et al., 1983; Olley and Pakes, 1996; Fabrizio et al., 2007).

The existence of a positive correlation between competition and productivity is of first-order significance for several reasons. The most salient is the possibility of productive efficiencies of competition, which are, as Williamson (1968) observed in the setting of merger evaluation, infra-marginal and therefore *prima fascia* larger than allocative efficiencies. The potential for such gains could motivate competition policy. Second, the correlation is relevant to recent work on international trade following Melitz (2003), which highlighted productive efficiencies as an important source of gains from trade liberalization. Third and finally, from a business economics standpoint, the correlation offers some leverage on the productivity dispersion puzzle: that establishments in the same industry with the same inputs often produce vastly yield outputs.

Though the existence of a positive correlation between competition and productivity bears on fundamental questions, the mechanism generating it remains controversial. I focus on two leading hypotheses: first, that competition has a direct causal effect on productivity and second, that the correlation is driven by selective attrition of low-productivity establishments in more competitive markets; respectively, the *treatment effect* and the *selection effect*. The treatment effect hypothesis says that competition behaves *as if* it were an input of the production function. Therefore, if one could, *ceteris paribus*, transplant a firm from a less- to a more competitive market, the treatment effect hypothesis implies that it would exhibit an increase in measured productivity. The language “treatment effect” here stands in for real economic phenomena within the firm; in fact, there exists several models consistent with such an effect of competition on productivity— more competitive markets may give firms better incentives to monitor managers or invest in productivity enhancements, or they may create positive informational externalities. Complementary to this, a number of historical studies have documented examples where competition — or the threat of competition —

spurred reorganization, renegotiation of contracts, and higher productivity.

The second hypothesis, the selection effect, operates through selective attrition of low-productivity establishments in highly competitive markets. In particular, it conjectures that market selection on productivity is more aggressive in more competitive markets. Because observation is contingent on survival, this implies that the econometrician will observe a correlation between productivity and competition among *surviving* plants, even when there is no treatment effect of competition on productivity. This hypothesis is corollary to an idea that has become central in models of industry dynamics and trade liberalization: that more competitive markets reallocate demand from low productivity establishments to high productivity ones.

The main contribution of this paper is to show that these two hypotheses— the treatment effect and the selection effect of competition— have economically distinct and econometrically separable implications. I develop two approaches. First, I use an explicit model of the establishment’s exit decision problem to derive a semi-parametric selection correction that quantifies the relative contributions of both hypotheses. The estimator does not require parametric assumptions on demand; instead, the key identifying assumptions are the timing of play and the exogeneity of innovations in establishment-level productivity types. In this sense identification in this approach is similar to that in papers on semi-parametric production function estimation.¹

The second empirical strategy I use is a grouped instrumental variables quantile regression approach at the market level; this allows me to identify the marginal effect of competition on productivity on the entire market-level productivity distribution. I use this to test the hypothesis, implied by selective attrition, that the effect is driven by changes in the left tail of the productivity distribution. In contrast with the first identification strategy, this approach requires no assumptions on the establishment’s decision problem beyond the existence of a threshold exit strategy; however, its conclusions are correspondingly less weaker: besides its intuitive graphical appeal, it can only offer evidence that selection is not the exclusive mechanism.

The natural setting for applying these methods is the ready-mix concrete industry. An important challenge in studying the relationship between competition and productivity is finding sufficient cross-sectional variation in competitive structure at the market level. Though rich data exists for most manufacturing industries, low transportation costs make those indus-

¹In particular, my approach is similar to the timing-based identification in Olley and Pakes (1996); Levinsohn and Petrin (2003), and Akerberg et al. (2006).

tries global in market definition. This unfortunate fact leaves the econometrician with only cross-industry or time series variation in market structure, both of which are suspect. In contrast, idiosyncratically high transportation costs make markets for ready-mix concrete fundamentally *local*, which I exploit to construct geographically defined markets that yield within-industry, cross-market variation. In addition, the availability of homogenous output measures in physical, rather than revenue, terms allows me to estimate physical productivity without knowing prices. This averts the issue raised in Klette and Giliches (1996), that revenue-based output measures are by construction correlated with markups and therefore competition, which would be particularly problematic for my application.

My results show that, in ready-mix concrete, the correlation between competition and productivity is driven by the treatment effect hypothesis, i.e. within-firm changes in productivity in response to competitive conditions. I find no evidence that the selection effect hypothesis drives productivity changes. I explore this further using a standard decomposition of output-weighted productivity at the market level that offers a direct way of measuring reallocation. I show that market-level reallocation, as proxied by the productivity–share covariance term, is unresponsive to exogenous changes in competition, which may explain the absence of a selection effect. Note that this does not imply an absence of market selection altogether. Prior work on ready-mix concrete has already documented a negative relationship between establishment productivity and exit.² My results show that the degree of market selection on physical productivity is not driven by differences in competition— more precisely, not enough to explain the observed correlation between competition and productivity. The most important implication of these results has to do with directions for future research. The treatment effect of competition on productivity, while econometrically identified, remains a black box in this paper. Its significance suggests that studying within-firm adjustment mechanisms is first-order important for understanding productive efficiencies from competition, gains from trade, and establishment-level productivity dispersion.

Section 2 briefly summarizes related literature to motivate the question. Section 3 describes the ready-mix concrete industry, the data used, and measurement issues associated with studying productivity, spatially defined markets, and competition indexes. Section 4 contains the main methodological contributions of the paper as well as results and robustness considerations. Section 5 further explores the selection effect non-result using a decomposition approach to measure reallocation. Section 6 concludes.

²See Foster et al. (2008) and Collard-Wexler (2011), as well as a prior draft of this paper.

2 Related Literature

This paper draws on several literatures which I summarize in three parts. The first set of papers elaborates the treatment effect hypothesis: this includes empirical documentation and a theoretical literature offering a rigorous foundation. The second set of papers is related to the empirical and theoretical basis of the selection effect. Finally, I briefly describe a large and celebrated literature studying productivity using establishment-level data, with special attention to ready-mix concrete.

2.1 Treatment Effect

The idea that productivity is directly related to competition has a long and controversial history in economics. Concluding a discussion of contemporary developments in the theory of monopoly, Hicks (1935) casually suggests that perhaps “the best of all monopoly rents is a quiet life.” Again in Leibenstein (1966) the idea emerges under the title “X-efficiency:” intuitively and empirically — but inexplicably — it seems as if firms in more competitive markets are more motivated to reduce costs. This particular incarnation was assailed in Stigler (1976), which objected to the substitution of “motivation” for sound economic reasoning.³ Indeed, at that time it was difficult to reconcile the empirical phenomenon with optimal choice theory: why would a monopolist have any less incentive to reduce costs?

Subsequent theoretical developments in games with asymmetric information and industry dynamics offered myriad answers to that question. If imperfect monitoring creates a wedge between optimal production and the the second-best (the optimum conditional on incomplete contracts) the relevant question becomes how that wedge depends on competitive conditions in the market. A small literature on principal-agent problems with multiple players following Holmström (1982) develops this idea— if the noise with which the principal observes agents’ behavior is correlated across agents, adding an additional agent has informational externalities that, in the limit, allow the principal to extract optimal effort. In the context of productive efficiencies of competition, owners may be able to more effectively monitor managers when they can compare their performance against that of other firms in the market. Alternatively, Schmidt (1997) proposes risk of liquidation as a way to re-align incentives of the agent; if competitive firms are more likely to go bankrupt, there is less room for shirk-

³See Perelman (2011) for a summary of the Leibenstein-Stigler debate concerning X-Efficiency.

ing. Raith (2003) offers another approach. That paper considers incentives to invest in cost reduction, and shows two competing effects of competition— a business-stealing effect and a scale effect. The former implies that increased substitutability generates more incentive to invest, while operating on a smaller scale in a thicker market weakens those incentives. The stark result of that paper is that in the long run – i.e. allowing for equilibrium entry and exit – the business stealing effect of competition dominates the scale effect, assuring a positive correlation between competition and productivity.

The intuition that competition can drive within-firm productivity changes has been documented in several industry-level studies. Schmitz (2002, 2005) show how increased competition drove greater labor productivity in the US Iron Ore industry. In cement, which saw dramatic changes due to import penetration in the 1980’s, this took the form of renegotiation of work rules and contracts at the plant level, as documented in Dunne et al. (2010). This literature is surveyed in more detail in Holmes and Schmitz (2010).

2.2 Selection Effect

The selection effect hypothesis is a theoretical prediction of industry dynamics models that descend from Hopenhayn (1992). Establishments in this framework make a decision of whether to stay in the market or exit, and they do so based on their idiosyncratic cost type as well as the degree of competition they face from other firms. Those establishments may exit if their productivity type is such that the present discounted value of remaining is lower than their scrap value. The stationary distribution of firms’ types, then, is the ergodic distribution generated by equilibrium entry and left-truncation as low-productivity firms exit. These models have been successful at modeling turnover and market size (Asplund and Nocke, 2006), barriers to entry (Barseghyan and DiCecio, 2011), and have been famously extended to capture trade liberalization and establishments’ decisions to export (Melitz, 2003).

These models are also the basis of the selection effect hypothesis, which claims that the equilibrium exit threshold increases as a market becomes more competitive. As the threshold rises, the selected set of surviving firms is on average more productive, generating a correlation between competition and productivity without resort to a treatment effect. This prediction depends on a key assumption on stage-game profits, which has been called “re-allocation:” that more competitive markets are better at reallocating demand from less- to

more-productive firms. This assumption is standard in industry dynamics models, though it often takes different forms.⁴ In the empirical literature, this reallocation mechanism has been an important channel for productivity growth since Olley and Pakes (1996), and Boone (2008) goes still further, arguing that reallocation is *constitutive* of competition, not merely related to it. In Appendix B I offer a simple entry model to illustrate the subtle and important relationship between reallocation and selection and in Section 5 I test for the reallocation mechanism in ready-mix concrete directly.

2.3 Establishment-Level Productivity

The availability of comprehensive establishment-level input and output data sparked a vibrant literature on productivity analysis. Early contributions in this literature are reviewed in Bartelsman and Doms (2000). Subsequently Foster et al. (2008) studied the relationship between demand, productivity, and exit. They find that even in markets with homogenous products, idiosyncratic demand shocks are an important determinant of firm survival.

An important subset of this literature has focused on ready-mix concrete. In particular, this paper builds on the pioneering work of Syverson (2004), which studied the negative correlation of productivity *dispersion* with competition. That paper nested the selection effect hypothesis in an explicit entry model but observed that the relationship could also be driven by a treatment effect, leaving the empirical deconflation to future work.⁵ Also closely related is Collard-Wexler (2011), which studies the relationship between productivity dispersion and firms' decisions to exit the market. Finally, Collard-Wexler (2013) studies the role of demand shocks in local markets for ready-mix concrete, finding a substantial market expansion effect.

⁴The assumption is most transparent in Asplund and Nocke (2006), where it takes the form of a log-supermodularity assumption on the profit function in type and market size, however it also follows from the parametric assumptions in Syverson (2004) and Melitz and Ottaviano (2008). It is also related to the multiplicative separability assumption, Condition U2, of Hopenhayn (1992)

⁵From footnote 6 of Syverson (2004): "One can remain agnostic about the specific source of productivity gains when competition is intensified. One possibility is an effect on "slack" or X-efficiency. That is, competition-spurred productivity growth occurs because producers are forced to take costly action to become more efficient, as in Raith (2003), e.g. However, in the mechanism modeled here, productivity growth is instead achieved by selection across establishments with fixed productivity levels; less efficient producers are pushed out of the market. Both mechanisms are influenced by market competitiveness in theory, and both are likely to play a role in reality. Measuring the relative size of the contribution of each to determining productivity differences is beyond the scope of this paper, however."

3 Data and Measurement

3.1 Ready-Mix Concrete

I use US Census of Manufactures data for the ready-mix concrete industry (SIC 3273) from years 1982, 1987, and 1992. Ready-mix concrete is a mixture of cement, water, gravel, and chemical additives that is used in sidewalks, foundations, and roads, among other applications. These ingredients are combined at the plant and transported to the construction site in a large drum mounted on an even larger truck.

Two features make the industry particularly interesting for studying the role of market structure: first, there are markedly high transportation costs that make competition local in character. When the mixing truck is loaded the concrete begins to harden to the interior of the drum, wasting materials and incurring maintenance costs. Therefore construction sites are typically serviced by nearby plants. For my purpose, this motivates the definition of geographic market areas and affords cross-sectional variation in market structure.

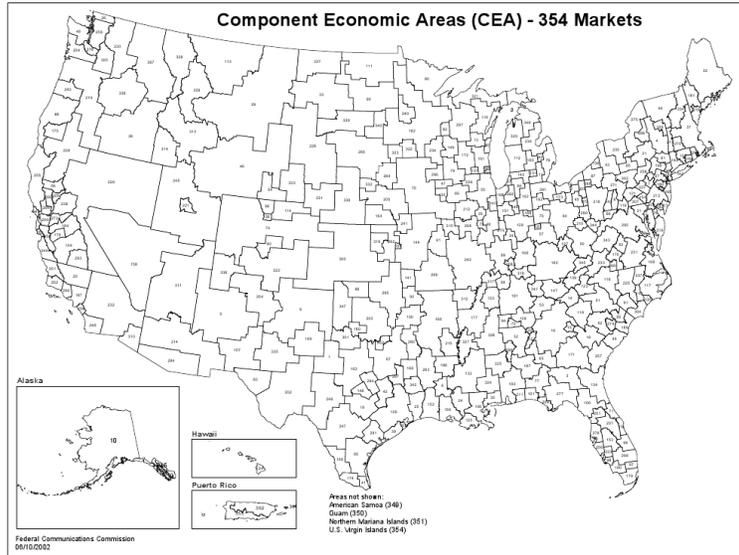
The second important feature of the industry is the homogeneity of the output. Though the composition of chemical additives may differ some by application, this generates little product differentiation. For this reason, in the years of my sample the Productivity Supplement to the Census of Manufactures collected output data in cubic yards, which obviates many of the concerns that would accompany the use of deflated revenue in estimating productivity for this application.

It is important to note that there is substantial entry and exit in this industry; in my sample roughly one third of plants disappear between each five-year census. Because the selection effect is predicated on differential exit, this tremendous amount of churn is favorable to the existence of such an effect. For a more extensive discussion of the ready-mix concrete industry, the interested reader is referred to Syverson (2008).

3.2 Market Definition

The empirical work that follows is identified off of cross-sectional variation in market competitiveness. This is motivated by the local character of competition, and therefore necessitates

Figure 1: Map of CEAs



Notes: This figure depicts Component Economic Areas (CEAs) as constructed by the Bureau of Economic Analysis, a complete partition of the counties of the US into market areas. See Johnson (1995) for details on construction.

careful market definition. I follow Syverson (2004) in using the 1995 Component Economic Areas (CEAs), which are constructed by the Bureau of Economic Analysis.

CEAs are a complete partition of the set of 3141 US counties into 348 market areas. They are constructed by assigning contiguous counties to nodes of economic activity (e.g. metro- or micropolitan areas). Assignment is based primarily on labor force commuting patterns from decennial census, and secondarily— for roughly 25% of non-nodal counties— by newspaper circulation data from the Audit Bureau of Circulations. See Johnson (1995) for more details on the construction of the 1995 CEA definitions.

Strictly speaking, the ideal market definition would be such that first, consumers in market A are *never* served by a firm in some other market B and second, all firms in each market have some strategic interaction. The use of CEAs is of course a compromise towards that end. While labor force commuting patterns are likely a good proxy for market areas when population is dense, a key weakness of this market definition is that rural market areas tend to be overly large; see Figure 1. For this reason I exclude from my sample all plants in the top decile of the CEA geographic size distribution.⁶

⁶See Appendix A for a complete discussion of my sample definition.

3.3 Competition and Demand

In order to exploit cross-sectional variation in competitiveness I require a competition index informed by the institutional features of the industry. I employ several variations on a count measure. My baseline is the number of ready-mix concrete establishments per square mile. The intuition behind this measure is the following: as the geographic density of establishments increases, their ability to differentiate from the nearest neighbor is diminished; in that sense establishments become more substitutable and therefore face greater competition.

The use of competition indices as a right-hand side variable in empirical work can be problematic. The main difficulty is summarized by Nickell (1996): when firms are heterogeneous, single-dimensional indices will seriously mis-measure competitiveness. For instance, in the case of a count measure, a market may have few firms because barriers to entry stifle competition; alternatively, it may have few firms because it is competitive and the existing firms are very efficient.⁷ This particular example is problematic for my application, as it suggests that the measurement error of the competition index is correlated with establishment productivity, my left-hand side-variable, a point raised by Holmes and Schmitz (2010)— i.e., the measurement error is non-classical.

To overcome this, I use demand shifters to isolate variation in long-run competitive structure.⁸ If establishment level heterogeneity is transient, a point on which the literature is in consensus, then we can integrate out over these short-run shocks by using an instrument for the long-run expected competitive structure. Isolating this source of variation rationalizes the use of a single-dimensional competition index; in particular, I use the expected number of firms conditional on exogenous demand.⁹ I introduce the following set of demand shifters as instruments for market size in ready-mix concrete: building permits issued, single-family building permits issued, and local government road and highway expenditures.¹⁰ These series

⁷Special thanks to James Schmitz for a colorful and helpful discussion on this point.

⁸By “long-run” here I mean the steady-state expectation with endogenous entry and exit.

⁹This appeal to long-run market fundamentals in order to understand the effect of market structure is not novel, and has been explored extensively in Sutton (1991). Falling barriers to trade and integration as European Community initiatives were rolled out created opportunities to study market expansion in Bottasso and Sembenelli (2001). Syverson (2004) regresses market-level productivity dispersion on construction employment, which proxies for demand. The market expansion effect of demand shifters for ready-mix concrete is studied more closely by Collard-Wexler (2013), which shows that entry responds to changes in construction employment. Beyond ready-mix concrete, trade economists have exploited changes in entry costs for international firms or trade barriers in order to capture changes in the competitive environment, and Kneller and McGowan (2014) study the relationship between demand for ethanol on productivity in the corn sector.

¹⁰The reason I diverge from prior work and eschew county-level construction employment as a demand

are generated at the county level by the US Census.¹¹

A weakness of the count measure as a competition index is that it ignores observed concentration, which may be informative if there are dimensions other than productivity on which firms are heterogeneous (e.g., their ability to secure contracts, c.f. Foster et al. (2008)). For this reason I use the number equivalent of the Herfindal-Hirschman Index (HHI) as an alternative measure.¹² As I will show in section 4.1, in some of my markets there are multiple establishments owned by the same firm. Without strong assumptions it is unclear whether the establishment count or the firm count is the better measure; therefore, in all of my results I present both.

3.4 Productivity

In all of my empirical analysis I treat establishment-level productivity residuals ω_{it} as data. These productivity residuals are the additive error in a log Cobb-Douglas production function with constant returns to scale:

$$\omega_{it} = y_{it} - \sum_k \alpha_{kt} x_{kit} \quad (1)$$

– where y_{it} is log physical output, x_{it} is a vector of inputs, and α_t is a vector of input elasticities, with $\sum_k \alpha_{kt} = 1$.¹³ The input elasticities α_t are consistently measured by input shares under the assumption that inputs are flexible. This is a reasonable assumption for ready-mix concrete– the amount of entry and exit in the industry has fostered a healthy

shifter for RMC is that I found that it consistently failed over-identification tests for exogeneity in my IV regressions. In some sense this might not be surprising because the measure of competition is so coarse– it may be that it is affecting the outcome variable through competition in a way that my measure does not capture. Alternatively, and more worryingly, it is likely that county-level construction employment is correlated with measurement error in productivity inputs due to local wage effects. By excluding it I find that my model passes over-identification tests easily; see the over-identification test results (Hansen J) from table 2. In the discussion of robustness in Section 4.4.2 I consider IV regressions with additional controls, including county construction employment.

¹¹Historical series for these variables are available from USA counties database online, however this resource is no longer updated.

¹²The number equivalent is simply the inverse of the HHI. Originally proposed by Adelman (1969), it can be interpreted as the number of symmetric firms that would generate the observed degree of concentration.

¹³The constant returns to scale assumption is critical because increasing returns to scale would generate a direct causal relationship between scale and (mis)measured productivity; I offer supplementary evidence for it as a robustness check in Section 4.4.1.

Table 1: Summary Statistics

	1982		1987		1992	
	mean	s.d.	mean	s.d.	mean	s.d.
No. Estab.	13.59	12.53	13.76	14.14.	13.52	13.92
No. Firms	9.51	7.45	10.46	9.11	10.43	9.23
HHI No. Estab.	6.84	5.28	6.99	6.48	7.03	6.40
HHI No. Firms	4.92	3.05	5.50	3.55	5.26	3.69
Building Permits	2645.51	5629.57	4291.51	7642.50	2951.70	4165.11
S.F. Building Permits	1419.62	2662.46	2851.15	4699.00	2443.78	3587.64
Road & Hwy \$	38994.95	59188.68	56030.91	93206.21	72173.89	121158.15

Notes: This table contains summary statistics for measures of competition $c_{m(i)t}$ as well as demand shifters for ready-mix concrete. All variables are aggregated to the year-CEA level.

secondary market for fungible used capital.¹⁴

As in prior work, my estimates of plant-level productivity exhibit a high degree of variance as well as intertemporal persistence. The standard deviation is 0.27 and the implied one-year autocorrelation of ω_{it} is 0.77. See Appendix A for details on the components and measurement of x_{it} .

4 Competition and Productivity

The techniques developed in this section are meant to measure the correlation between competition and productivity, document a causal relationship, and then to decompose that relationship into the portion attributable to the treatment and selection effects. Data characterized in Section 3 is summarized in Table 1.

4.1 Reduced-Form Analysis

I begin by estimating the following reduced-form relationship, which captures the correlation between competition and productivity:

¹⁴In fact, the particularly interested reader can purchase a ready-mix concrete truck on eBay.com.

$$\omega_{it} = \beta_t + \beta c_{m(i)t} + \varepsilon_{jt}. \quad (2)$$

There is an important source of bias in the OLS variant of this model: the presence of high-productivity establishments may deter entry by competitors, which would motivate a negative correlation. One can think of this as a source of non-classical measurement error; conditional on realizations of establishment-level productivity, the number of firms may be a bad proxy for the degree of competition, and the error is correlated with plant-level productivity.

In order to obtain an unbiased estimate I use exogenous demand shifters to instrument for long-run market structure. Intuitively, I am integrating out over transient, short-run productivity shocks to obtain the unconditional effect of the number of firms on expected establishment-level productivity. Pursuant to the discussion in Section 3.3, I use the geographic density of building permits, single-family residential building permits, and local government road and highway expenditures. These demand shifters are relevant insofar as demand has a market expansion effect in equilibrium, consistent with standard industry dynamics models and the findings of Collard-Wexler (2013) for ready-mix concrete in particular. To make the argument for exogeneity I observe that ready-mix concrete is but a small part of most construction budgets; price changes due to changes in market power are unlikely to drive reverse causality in my application. However, a reasonable concern is that these shifters may have a direct effect on productivity, if for no other reason than measurement error in the competition index. For this reason I also present results from over-identification tests (Hansen J).

Results for OLS and IV estimation of (2) are presented in Table 2. The coefficient on competition is stable across models (1)-(4) and (5)-(8), suggesting little dependence on the particular choice of competitive index. The most noticeable difference is across the OLS and IV specifications; as Nickell (1996) suggests, the bias introduced by correlation between productivity types and the measurement error in the competition index seems to be negative. Note as well that results from the over-identification tests consistently and safely fail to reject the hypothesis of exogeneity of the instruments. If demand had an indirect causal effect, whether due to mis-measurement of the competition index or some other hypothesis, e.g. agglomeration effects of market size, I should reject this hypothesis. This is a strong validation of the choice of measure of competition, and also helps to rule out alternative channels by which demand might affect productivity.

Table 2: OLS and IV Results

	Dependent Variable: ω_{it}				Instrumental Variables			
	OLS							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
No. Estab./mi. ²	0.0350*				0.0468*			
	(0.0056)				(0.0064)			
No. Firms /mi. ²		0.0318*				0.0494*		
		(0.0058)				(0.0072)		
HHI No. Estab./mi. ²			0.0302*				0.0481*	
			(0.0062)				(0.0066)	
HHI No. Firms /mi. ²				0.0214*				0.0561*
				(0.0060)				(0.0095)
Observations	7381	7381	7381	7381	7373	7373	7373	7373
R^2	0.1099	0.1075	0.1090	0.1039	0.1093	0.1057	0.1062	0.0935
First-Stage F					344.5243	196.9467	171.4296	53.0703
p Value					0.0000	0.0000	0.0000	0.0000
Hansen J Statistic					2.3297	1.3312	2.9247	0.8982
p Value					0.3120	0.5140	0.2317	0.6382

Notes: Here I present OLS and IV results for the effect of competition on productivity residuals, as discussed in Section 4.1. Models (1)-(4) and (5)-(8) use the four distinct competition measures indicated on the left. Year-specific constants are included but not reported and standard errors are clustered at the CEA level. Instruments for models (5)-(8) are the number of building permits per square mile, the number of single family building permits per square mile, and local government road and highway expenditure per square mile. First-stage F tests and Hansen J (over identification) test statistics are reported with associated p values. Standard errors are presented in parentheses. For reference, * signifies $p \leq 0.05$.

These results imply that there is an economically and statistically significant relationship between competition and productivity, and that relationship is *causal*. What they do not shed light on, however, is whether that causal effect is driven by the treatment effect of competition or the selection effect. The former effect is a within-establishment causal effect of competition—it behaves as if competition were an input to production. The latter is driven entirely by selective attrition of low-productivity establishments in more competitive markets. Separating these two stories is the objective of Sections 4.2 and 4.3.

4.2 Semi-parametric Disambiguation

The first identification strategy I propose is based on an explicit model of the firm’s exit choice, the problem driving the selection effect hypothesis. By modeling this I can construct a semi-parametric selection correction that will allow me to decompose the causal effect identified in the IV approach of Section 4.1 into a treatment component and a selection component. Then I quantify and compare their relative contributions.

4.2.1 Behavioral Model

Consider a firm’s exit decision in a model of Markovian industry dynamics cast after Ericson and Pakes (1995). There is a finite set of active firms in a market. \mathcal{S}_t denotes the state of the market in period t , which has three components: a vector of idiosyncratic productivity types for each active firm Φ_t ; a set of states X_t determined by firms’ dynamic choices, and a market-level demand shifter d_t . The game is public, so the information set of each establishment at time t corresponds to \mathcal{S}_t . In a Markov perfect equilibrium, firms make choices that depend only on \mathcal{S}_t , and in turn their actions determine Markov transition probabilities for the entire state of the game.

Establishments enter, exit, and make choices to maximize the expected discounted value of their profits. Stage game profits for establishment i are written $\pi_i(\mathcal{S}_t)$. This function abstracts from stage game decisions without dynamic effects, e.g. choice of price or quantity. Establishments also make dynamic choices denoted by a_t to affect the evolution of X_t , at a cost $c_i(a_i, \mathcal{S}_t)$.

The order of play determines the informational content of firms’ decisions, and is therefore

central to identification, a feature often overlooked but signature to the semi-parametric production function estimation literature. I assume a three-part period structure; this is my first identifying assumption:

Assumption 1. (*timing of play*)

$$\Phi_t \text{ evolves} \quad \rightarrow \quad \text{stage game } \pi \quad \rightarrow \quad \text{entry, exit, } a_t \text{ chosen}$$

At the beginning of the period, Φ_t is announced. Next, establishments produce and realize stage game profits according to π . Third and finally, potential entrants arrive and all establishments make exit and choose actions a_t (e.g., investment). Establishments only make choices in the third and final stage. What is significant about the timing structure is that the stage game precedes exit; intuitively, one might think that firms learn their new productivity draw *by producing*. The Bellman function of the active establishment, which chooses whether to exit and obtain a scrap value normalized to zero and, should they choose to persist, actions a_{it} subject to costs $c(a_{it}, X_t)$ can be written:

$$V_i(\mathcal{S}_t) = \max\{0, \max_{a_{it}}\{\delta\mathbb{E}[\pi_i(\mathcal{S}_{t+1}) + V_i(\mathcal{S}_{t+1})|a_{it}] - c_i(a_{it}, \mathcal{S}_t)\}\}. \quad (3)$$

This Bellman equation nests two decisions, but the one I am primarily interested in is the exit choice. Conditioning on equilibrium play by other agents, let $v_i(\mathcal{S}_t)$ denote the solution to the inner maximization problem. Now the establishment's optimal exit choice can be characterized by:

$$\chi_i(\mathcal{S}_t) = \begin{cases} 1 & \text{if } v_i(\mathcal{S}_t) \geq 0 \\ 0 & \text{else.} \end{cases} \quad (4)$$

Moreover, with standard monotonicity conditions on π it is possible to show that this exit choice can be written as a threshold rule. That is, there exists some $\phi_i^*(\mathcal{S}_t)$ such that the establishment chooses to exit if $\phi_{it} < \phi_i^*(\mathcal{S}_t)$. This exit threshold drives the selection effect hypothesis, per the discussion in Section 2.2.

Second, I make a standard assumption on the exogeneity of innovations in ϕ_{it} :

Assumption 2. (*exogeneity of innovations*)

$$p(\phi_{it+1}|\mathcal{S}_t) = p(\phi_{it+1}|\phi_{it}).$$

See Akerberg et al. (2006) for a fuller discussion of this assumption; as an economic assumption it implies that establishments do not make choices that affect the evolution of their idiosyncratic productivity type ϕ . Econometrically, the assumption is useful because it affords me a source of exogenous variation in the model.

4.2.2 Identification

Consider the naïve structural interpretation of the linear model I estimated in Section 4.1:

$$\omega_{it} = \beta_t + \beta_c c_{m(i)t} + \phi_{it}. \tag{5}$$

The difference with the earlier reduced-form model is that now the error term is interpreted as the underlying productivity type of the plant. In this light it is straightforward to see why the instrumental variables approach does not recover an unbiased estimate of β_c as the treatment effect. In particular, the sample is selected: according to my behavioral model I only observe plants that survived the prior period, i.e. such that $\phi_{it-1} \geq \phi_i^*(\mathcal{S}_t)$. Therefore $c_{m(i)t}$ and ϕ_{it} will be correlated, which biases the estimate of β_c . That bias *is* the selection effect.¹⁵ In this section I argue that they can be separated.

My solution to identifying β_c begins by restricting attention to surviving establishments—i.e., those for which we have data from period $t - 1$. The following exercise will make clear the value of this restriction. First, take the expectation of both sides of Equation (5) conditioning on the the state of the market at time $t - 1$ and the survival of the establishment:

$$\begin{aligned} \mathbb{E}[\omega_{it}|\mathcal{S}_{t-1}, \phi_{it-1} \geq \phi_i^*(\mathcal{S}_{t-1})] & \tag{6} \\ & = \beta_t + \beta_c \mathbb{E}[c_{m(i)t}|\mathcal{S}_{t-1}, \phi_{it-1} \geq \phi_i^*(\mathcal{S}_{t-1})] + \mathbb{E}[\phi_{it}|\mathcal{S}_{t-1}, \phi_{it-1} \geq \phi_i^*(\mathcal{S}_{t-1})]. \end{aligned}$$

¹⁵It is important to emphasize that the IV estimate is still *causal*, despite the fact that it conflates the two possible channels (treatment and selection).

Focusing on the last term,

$$\begin{aligned}
\mathbb{E}[\phi_{it} | \mathcal{S}_{it-1}, \phi_{it-1} \geq \phi_i^*(\mathcal{S}_{t-1})] &= \mathbb{E}[\phi_{it} | \mathcal{S}_{it-1}] \\
&= \mathbb{E}[\phi_{it} | \phi_{it-1}] \\
&= g(\phi_{it-1}).
\end{aligned} \tag{7}$$

The first equality follows from Assumption 1; the timing structure implies that the event $\phi_{it-1} \geq \phi_i^*(\mathcal{S}_{t-1})$ is fully determined by \mathcal{S}_{t-1} . The second equality follows directly from Assumption 2. Now define $\eta_{it} \equiv \phi_{it} - \mathbb{E}[\phi_{it} | \phi_{it-1}]$. From Assumption 2 we know that this innovation η_{it} is exogenous to all of the arguments in \mathcal{S}_t . Plugging this into (5) yields

$$\omega_{it} = \beta_t + \beta_c c_{m(i)t} + g_t(\phi_{it-1}) + \eta_{it} \tag{8}$$

— which can be evaluated by inverting the model in (5) and plugging it in for ϕ_{it-1} . This allows me to identify the treatment effect of competition on productivity, β_c . Note that $g_t(\cdot)$ is an unknown and potentially complicated function implied by the model; absent further assumptions it will be important to estimate this function flexibly, which is the sense in which my approach is semi-parametric.

4.2.3 Estimation

I estimate the model using nonlinear GMM with two sets of moments. The first is derived from Equation (8):

$$\mathbb{E}[\eta_{it} | Z_{it}] = 0. \tag{M1}$$

In order to construct η_{it} I use a third-order polynomial series to approximate $g_t(\cdot)$. This is the source of the nonlinearity— the argument of the polynomial series is the lagged residual from (5), which depends on the parameters of interest. To control for measurement error I continue to use my exogenous demand shifters as instruments for $c_{m(i)t}$.

Table 3: Semiparametric Results

	Dependent Variable: ω_{it}			
	(1)	(2)	(3)	(4)
No. Estab./mi. ²	0.0456*			
	(0.0090)			
No. Firms/mi. ²		0.0488*		
		(0.0118)		
HHI No. Estab./mi. ²			0.0503*	
			(0.0101)	
HHI No. Firms/mi. ²				0.0571*
				(0.0150)
Selection Coeff. α_c	0.0046	0.0049	0.0005	0.0043
	(0.0030)	(0.0065)	(0.0041)	(0.0085)
Observations	3064	3064	3064	3064

Notes: This table presents results for the semi-parametric selection correction procedure detailed in Section 4.2 using four different indices for competition. Year-specific constants are included but not reported. Standard errors are presented in parentheses. For reference, * signifies $p \leq 0.05$.

Moment (M1) is sufficient to estimate treatment effect of competition on productivity, but I would also like to be able to quantify the contribution of the selection effect. To see how I do this, consider the following thought experiment: given an estimate $\hat{\beta}_c$, one could go back and construct an estimate of productivity net of the treatment effect of competition. Taking that estimate as a dependent variable, next run the following regression:

$$\omega_{it} - \hat{\beta}_c c_{m(i)t} = \alpha_t + \alpha_c c_{m(i)t} + u_{it}. \quad (9)$$

Since I have subtracted off the treatment effect $\hat{\beta}_c c_{m(i)t}$ from the left-hand side, what α_c is capturing is the selection bias induced by correlation between $c_{m(i)t}$ and ϕ_{it} . If there is no selection effect, then I expect to find $\hat{\alpha}_c = 0$. In this sense it is useful because it allows me to quantify the relative contribution of the selection effect to the IV estimates obtained in Section 4.1.¹⁶ This informs the second set of moments I use in estimation, which implement the regression in Equation (9):

¹⁶Note, however, that α_c is not a primitive; it only measures the contribution of the selection effect in equilibrium.

$$\mathbb{E}[u_{it}|Z_{it}] = 0. \tag{M2}$$

I estimate the model using moments (M1) and (M2). Results are presented in Table 3. As a simple diagnostic check, since this is a decomposition of the causal effect of competition on productivity, the sum of β_c and α_c should be roughly equal to my IV estimate from Table 2 for the corresponding competition index. None of my specifications reject this hypothesis.¹⁷ With respect to the results themselves, the stark finding across all specifications is that the treatment effect seems to be driving almost all of the causal effect of competition on productivity. Where competition is measured by the number of RMC establishments per square mile, my estimates imply that doubling the number of establishments will raise output by 4.56% due to within-establishment responses to competition and 0.46% due to market selection on productivity type. Importantly, for none of my specifications is the estimate of the selection effect bias statistically different from zero.

This is the main result of the paper— that the effect of competition on productivity seems to be driven by a within-firm response — a treatment effect — instead of market selection driven by reallocation of demand — the selection effect. To strengthen this result in Section 4.3 below I adopt an even less model-dependent approach to looking for evidence of a selection effect, and in Section 5 I look directly for evidence of a reallocation mechanism responding to competition, the mechanism which is the theoretical basis of the selection effect hypothesis.

4.3 Quantile Analysis

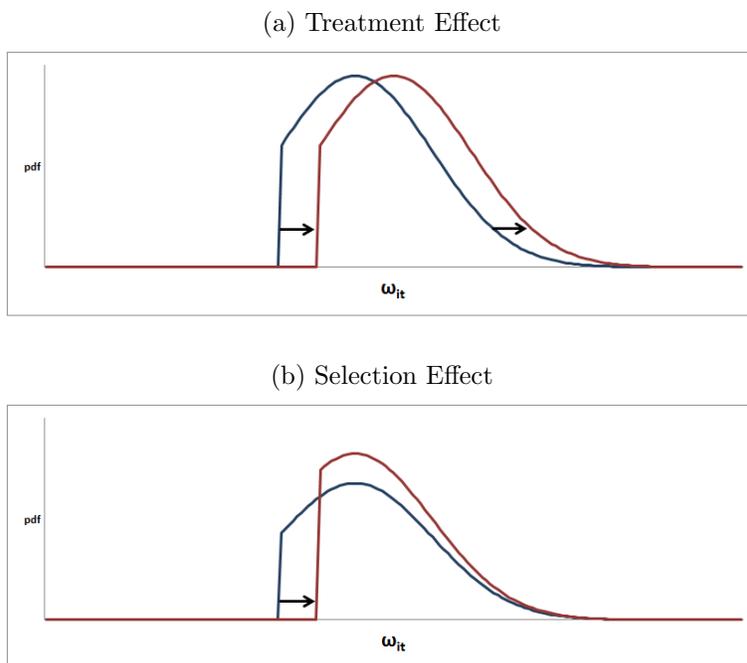
Industry dynamics models, in particular those on which the selection effect hypothesis relies, make predictions not merely for the first moment of productivity or the second, but the shape of the entire distribution. In this section I use a grouped quantile IV regression to see whether the response of the empirical productivity distribution to competition is consistent with the selection effect hypothesis.

¹⁷Note that the reason the correspondence is imprecise is because of the change of sample; the semi-parametric selection correction approach uses only observations from 1987 and 1992.

4.3.1 Identification

The identification strategy in this section hinges on the distinct predictions of the treatment effect and the selection effect for changes in the distribution of productivity types among active establishments in equilibrium as the market becomes more competitive. A visual motivation for the distinction is presented in Figure 2. The distribution of productivity types is assumed to be left-truncated, which reflects a threshold rule for exit common in industry dynamics models. Figure 2a depicts an additive shift in the entire distribution, consistent with the linear model described in (2). Alternatively, Figure 2b depicts a positive shift that is driven by an increase in the left-truncation point. The key difference here is that in the latter case, the increase in average productivity is driven entirely by an increase in the distribution at lower quantiles. Consistent with the selection effect, the higher end of the productivity distribution, which is determined by technological primitives, is invariant. In other words, the selection effect predicts that the marginal effect of competition should be declining in the quantile of the market-level productivity distribution.

Figure 2: Treatment Effect vs. Selection Effect on Distribution of Productivity



Notes: This figure illustrates two changes in the distribution of productivity residuals; in panel (a) a linear shift of the entire distribution, and in panel (b) a shift in the left truncation point. These correspond to the treatment effect and selection effect hypotheses, respectively.

It is important to note that the common shift depicted in Figure 2a, meant to represent the treatment effect hypothesis, is entirely driven by the parametric assumption of a common, constant treatment effect in model (2). Indeed, if the treatment effect is motivated by bankruptcy aversion then the lower tail of the distribution should be more responsive to competitive pressures. For this reason we cannot use this approach to credibly test whether the treatment effect is at play. However, the truncation interpretation of the selection effect does not depend on parametric form, and therefore we can test whether the effect of competition on productivity is monotonically declining to zero as we look at the effect on higher quantiles.¹⁸ This test is less decisive than the semi-parametric approach in Section 4.2, but has the advantage that it requires correspondingly fewer assumptions and offers a simple visual interpretation.

4.3.2 Estimation

I begin by aggregating my data to the market level. Let $\rho_{mt}^{(k)}$ be the k th decile of the market-level distribution of ω_{jt} . Now I am interested in the following empirical model:

$$\rho_{mt}^{(k)} = \beta_t^{(k)} + \beta_c^{(k)} c_{mt} + \nu_{mt}. \quad (10)$$

Aggregating to the market level is important here; this is the sense in which it is a “grouped quantile” approach. Alternatively, one could pool across markets and run a standard IV quantile regression. However, identification in that model depends on a rank similarity condition— that the error term is identically distributed across markets (Chernozhukov and Hansen, 2005). This is contradicted both by the theory motivating the selection effect hypothesis and also, more importantly, by the likely existence of other market-level factors affecting productivity. The advantage of the grouped quantile IV regression approach is that, by making the group quantile the dependent variable, it nets out the common, group-level effects.¹⁹

There is, however, a subtle source of bias in (10) as an empirical model: as we see from

¹⁸Note that this argument can be weakened substantially— in a more complex model the truncation point may depend on other attributes of the market or of the establishment. I handle market effects explicitly with my grouped quantile approach, which allows for market-specific random effects on productivity. Establishment-specific determinants of exit will make exit appear probabilistic given limited data, but none the less the left tail of the productivity distribution should exhibit more sensitivity.

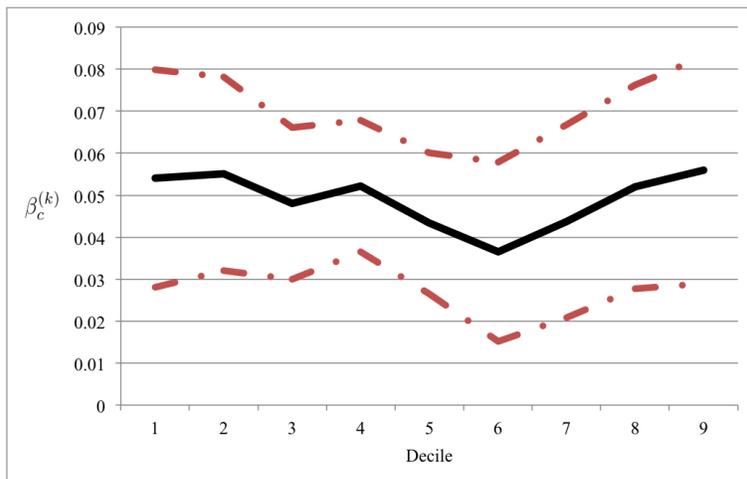
¹⁹See Chetverikov et al. (2013) for further discussion of the grouped IV quantile regression approach.

Table 4: Quantile IV Results

	Dependent Variable: $\rho_{mt}^{(k)}$								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
No. Estab./mi. ²	0.0540* (0.0132)	0.0551* (0.0118)	0.0480* (0.0092)	0.0522* (0.0080)	0.0433* (0.0086)	0.0365* (0.0109)	0.0438* (0.0117)	0.0520* (0.0124)	0.0559* (0.0138)
$H_0: p$ Value	0.0000	0.0157	0.7378	0.6990	0.6891	0.7322	0.1296	0.0069	0.0000
No. Firms/mi. ²	0.0544* (0.0132)	0.0552* (0.0117)	0.0483* (0.0093)	0.0524* (0.0080)	0.0435* (0.0086)	0.0369* (0.0109)	0.0443* (0.0116)	0.0518* (0.0124)	0.0554* (0.0137)
$H_0: p$ Value	0.0000	0.0130	0.5339	0.8455	0.9857	0.8941	0.2471	0.0097	0.0000
HHI No. Estab./mi. ²	0.0502* (0.0126)	0.0516* (0.0112)	0.0446* (0.0086)	0.0488* (0.0076)	0.0402* (0.0081)	0.0334* (0.0102)	0.0400* (0.0110)	0.0487* (0.0116)	0.0530* (0.0130)
$H_0: p$ Value	0.0000	0.0098	0.2295	0.6351	0.9271	0.9840	0.5420	0.0236	0.0000
HHI No. Estab./mi. ²	0.0536* (0.0132)	0.0546* (0.0118)	0.0476* (0.0093)	0.0517* (0.0081)	0.0428* (0.0086)	0.0360* (0.0108)	0.0434* (0.0115)	0.0511* (0.0122)	0.0550* (0.0136)
$H_0: p$ Value	0.0000	0.0005	0.0048	0.0111	0.1224	0.5036	0.5628	0.0080	0.0000
Observations	917	917	917	917	917	917	917	917	917

Notes: This table contains IV regression results for the effect of competition on deciles of the productivity residual distribution at the year-CEA level of aggregation, as discussed in Section 4.3. Year-specific constants are included but not reported and standard errors are clustered at the CEA level. Every cell represents an independent IV regression. By column, Model (k) corresponds to IV regressions with the k^{th} decile of the productivity residual distribution as a dependent variable. By row, regressions use the competition measure reported on the left. Instruments for all regressions are the number of building permits per square mile, the number of single family building permits per square mile, and local government road and highway expenditure per square mile. Standard errors are presented in parentheses. For reference, * signifies $p \leq 0.05$. For each regression I also present a p Value for H_0 , the null hypothesis that the coefficients on the polynomial series in the number of establishments are jointly zero. A p value less than 0.05 can be interpreted as evidence that the bias correction terms are estimated to be statistically different from zero.

Figure 3: Effect of the (log) No. Establ./mi.² on Productivity by Decile



Notes: Here I present graphically the results from Table 4 for the case where the density of ready-mix establishments is taken as the competition index. The equivalent diagram for alternative competition indices is so similar that I omit them.

Table 1, in many of my markets there are only a handful of establishments. If the number of establishments is small relative to the number of quantiles, then the higher (lower) quantiles will be equal to the highest (lowest) order statistic of the distribution. Those order statistics are mechanically correlated with the number of draws from which they are taken, and therefore mechanically correlated with count measures of competition. I call this “order statistic bias.” In this application it implies that the $\hat{\beta}_c^{(k)}$ will be biased downward for small k and upward for large k , where the productivity effect of competition is conflated with the mechanical properties of order statistics. To control for this effect I instead estimate the following regression:

$$\rho_{mt}^{(k)} = \beta_t^{(k)} + \beta_c^{(k)} c_{mt} + g^{(k)}(n_{mt}) + \nu_{mt} \quad (11)$$

— where n_{mt} is the number of establishment observations which are aggregated up to the market level for this regression, and $g^{(k)}(\cdot)$ is an unknown function which I approximate with a third-order polynomial series, independently for each k . This term functions as a semi-parametric correction for the order statistic bias described above, and implies that we are identifying $\beta_c^{(k)}$ of changes in competitive *density* alone. Recall from section 4.3.1 the prediction of the selection effect hypothesis: $\beta_c^{(k)}$ should be positive and declining to zero as k

goes from 1 to 9. As before, I instrument for c_{mt} using my market level demand shifters: the geographic density of building permits, single-family building permits, and local government road and highway expenditures.

Results for the four models — each using a distinct competition measure — for deciles $k = 1, \dots, 9$ are presented in Table 4 and visually in Figure 3. There is some evidence of a declining effect in the left tail, consistent with the selection effect. However the stark and surprising result, for each of these models, is that the highest productivity establishments seem to enjoy not only a nonzero effect, but a particularly large effect of competition on productivity. This is inconsistent with the selection effect hypothesis.

Table 4 also presents p values for the hypothesis that the coefficients on the polynomial series approximating $g^{(k)}(\cdot)$ are jointly zero. This is a test of the null that the order statistic bias is negligible. Consistent with the understanding above, I find that we can reject the null hypothesis for k very small or k very large, and further that the order statistic bias plays a larger role for HHI numbers-equivalent measures of competition. This is intuitive— the two non-HHI measures of competition are more mechanically correlated with the number of establishments in a market.

It is important to observe that results from the quantile approach do not decisively rule out the existence of a selection effect. I am only able to identify and quantify the contribution of the selection effect, as I did in Section 4.2, by supplementing the econometrics with assumptions about the establishments’ exit decision problem. In contrast with that approach, which found zero role for the selection effect, what the quantile analysis tells us is that the selection effect is insufficient to explain the features of the data— in particular, that establishments in the upper quantiles of the productivity distribution, those theory tells us should be least likely to exit, exhibit productivity gains on par with those of the lower quantiles when competition increases. While this does not rule out the selection effect entirely, it does suggest that the treatment effect is operative for some establishments.

4.4 Robustness Analysis

Both identification strategies strongly support the treatment effect hypothesis. However, one may be concerned that sample selection, my instruments, or alternative hypotheses are driving the existence of a correlation in the first place. The most concerning of these is

returns to scale— if there are increasing returns to scale in ready-mix concrete, then the demand shifters may be having a treatment effect on measured productivity independent of competition. I offer supplementary evidence for the constant returns to scale assumption in Section 4.4.1. Section 4.4.2 continues with an array of alternative specifications to show that the causal relationship is robust.

4.4.1 Returns to Scale

In Section 3.4 I assumed constant returns to scale in order to use input shares to measure productivity. By that construction, I obtained elasticities α and used them to derive productivity residuals ω . Consistent with this, define \hat{y}_{it} to be the constant returns to scale predicted output:

$$\hat{y}_{it} = \sum_k \alpha_k x_{kit}. \quad (12)$$

Suppose, however, that the true model involves returns to scale of order γ . Now,

$$y_{it} = (1 + \gamma)\hat{y}_{it} + \tilde{\omega}_{it} \quad (13)$$

— where $\tilde{\omega}_{it}$ is the true productivity term. Rearranging Equation (13) and taking expectations the relationship to productivity, as measured in Section 3.4, becomes clearer:

$$\mathbb{E}[y_{it} - \hat{y}_{it}] = \mathbb{E}[\tilde{\omega}_{it}] + \gamma\hat{y}_{it}. \quad (14)$$

The left-hand side of this equation corresponds to actual output net of predicted output under CRS. It suggests a simple approach to measuring γ : regress ω_{it} , as constructed in Section 3.4, on CRS predicted output. Predicted output is of course endogenous, but my market demand shifters from Section 3.3 are natural instruments for scale. Model (1) of Table 5 presents results for this IV regression. At first glance, the coefficient of $\hat{\gamma} = 0.0985$ suggests increasing returns to scale, which would trivialize the correlation between productivity and competition. However, note well the results for the over-identification test, which strongly reject the hypothesis that the instruments are not having a treatment effect on the outcome

Table 5: Returns to Scale Results

Dependent Variable: ω_{it}					
	(1)	(2)	(3)	(4)	(5)
Scale Coeff. (\hat{y}_{it})	0.0985*	-0.0273	-0.0139	-0.0392	-0.0131
	(0.0172)	(0.0311)	(0.0287)	(0.0353)	(0.0319)
No. Estab./mi. ²		0.0556*			
		(0.0121)			
No. Firms/mi. ²			0.0539*		
			(0.0122)		
HHI No. Estab./mi. ²				0.0615*	
				(0.0137)	
HHI No. Firms/mi. ²					0.0613*
					(0.0161)
Observations	7373	7373	7373	7373	7373
R^2	0.0000	0.1122	0.1094	0.1018	0.0947
Hansen J Statistic	20.5482	1.5197	1.0716	1.3913	0.6728
p Value	0.0000	0.2177	0.3131	0.2382	0.4121

Notes: Here I present IV results for the effect of scale and competition on productivity residuals, as discussed in Section 4.4.1. In model (1) scale is the only dependent variable; models (2)-(5) include the four alternative competition indices as well. Year-specific fixed effects are included but not reported, and standard errors are clustered at the CEA level. Hansen J (over identification) test statistics are reported with associated p values. Standard errors are presented in parentheses. For reference, * signifies $p \leq 0.05$.

variable. Indeed, this is consistent with the maintained hypothesis of the paper— if our demand shifters are having an effect on productivity via competition, this would invalidate them as instruments for the regression of ω on predicted output. In models (2) – (5) I resolve this problem by adding $c_{m(i)t}$ as a regressor, with the result that the returns to scale parameter disappears.

This result is of independent interest— it suggests that if we take a purely physical approach to measuring productivity, ignoring market conditions, we might mistake the productivity effects of competition for increasing returns to scale; a production-side analogue of the scale estimator bias discussed in Klette and Giliches (1996).

Table 6: Alternative Specification Results

	Dependent Variable: ω_{it}					
	(1)	(2)	(3)	(4)	(5)	(6)
No. Estab./mi. ²	0.0434*	0.0446*	0.0409*	0.0625	0.0431*	0.0443*
	(0.0071)	(0.0078)	(0.0062)	(0.0849)	(0.0068)	(0.0058)
No. Firms/mi. ²	0.0467*	0.0459*	0.0425*	0.0569	0.0454*	0.0471*
	(0.0079)	(0.0086)	(0.0069)	(0.0876)	(0.0077)	(0.0065)
HHI No. Estab./mi. ²	0.0435*	0.0434*	0.0392*	0.0747	0.0443*	0.0446*
	(0.0071)	(0.0075)	(0.0060)	(0.0895)	(0.0069)	(0.0059)
HHI No. Firms/mi. ²	0.0518*	0.0492*	0.0418*	0.0643	0.0517*	0.0518*
	(0.0097)	(0.0105)	(0.0080)	(0.0832)	(0.0096)	(0.0084)
Observations	4958	5861	7098	7373	4202	7373

Notes: Here I consider alternative specification results for the effect of competition on productivity residuals. Note that each cell represents an independent regression. Model (1) uses five-year lagged instruments; Model (2) uses only CEAs with 10 or more plants; Model (3) uses additional instruments (see text); Model (4) is the within-firm estimator; Model (5) is the between-county estimator, and finally Model (6) uses a revenue-based productivity residual. Year-specific constants are included but not reported and standard errors are clustered at the CEA level. Standard errors are presented in parentheses. For reference, * signifies $p \leq 0.05$.

4.4.2 Alternative Specifications

In this section I present a number of alternative specifications for the instrumental variables results from Table 2. Results for the alternative specifications are presented in Table 6 for each of the four measures of competition; note that each coefficient estimate reported is from a distinct regression.

In Model (1) I report IV results for the same regression with 5-year lagged instruments. This is meant to address the concern that contemporaneous demand shifters may be poor instruments for long-run demand shocks if convergence is slow. Alternatively, Model (2) restricts attention to CEAs in which there are 10 or more ready-mix concrete plants, to guarantee that the results are not driven by low-density markets. Model (3) uses additional instruments: county employment in construction-related industries per square mile, population per square mile, and the number of 5+ family building permits per square mile. For each of these specifications, (1)-(3) I find no substantive deviation from the main results of Table 2.

In order to more carefully examine the source of variation, model (4) reports the within-plant regression. This is identified exclusively off of the within-market time series variation;

the value of the coefficient is consistent with my other results, but there is insufficient variation to obtain a statistically significant estimate. This suggests that most of the identifying variation in my main regressions is between county. Model (5) reports results for the between-county regression (i.e., with county fixed effects) that are consistent with my earlier results.

Finally, model (6) reports results for a specification where ω_{it} is calculated using plant revenues instead of physical output. This is important because the availability of physical output data is special to ready-mix concrete and a handful of other industries; the fact that the result is similar offers some reassurance about the extension of these methods to other industries.

5 Reallocation and Productivity

The selection effect hypothesis, for which I find no evidence in my empirical results above, is corollary to an important idea in the productivity and trade literatures: that more competitive markets reallocate demand from less productive firms to more productive firms. Note that selection is the only manifestation of reallocation in the results above; my regressions treat small plants and large plants identically. This means I capture only the extensive margin of reallocation (exit), and not the intensive margin (reallocation of output among surviving establishments). In this section I use a market-level static decomposition of output-weighted productivity to take a closer look at why the selection effect is not salient and whether there is evidence for the reallocation mechanism in ready-mix concrete.

Despite innovations in the design of productivity decompositions, the state of the art in *analysis* of the decomposed elements of aggregate productivity amounts to carefully eyeballing trends at a handful of data points in a time series. This is, of course, due to limitations of the data. In this sense ready-mix concrete offers a unique opportunity—using cross-sectional variation in market structure I can formally identify causal relationships between the decomposed elements of productivity and my competition index, instrumented by local demand shifters.

To construct my left-hand side variables I borrow from Olley and Pakes (1996), who develop a static decomposition of output-weighted average productivity, denoted p_{mt} , into a mean term and a covariance term:

$$\begin{aligned}
p_{mt} &\equiv \sum_{\{i:m(i)=m\}} s_{it}\omega_{it} \\
&= \sum_{\{i:m(i)=m\}} \bar{s}_t\bar{\omega}_t \\
&\quad + \underbrace{\bar{s}_t \sum_{\{i:m(i)=m\}} (\omega_{it} - \bar{\omega}_t)}_{=0} + \bar{\omega}_t \underbrace{\sum_{\{i:m(i)=m\}} (s_{it} - \bar{s}_t)}_{=0} \\
&\quad + \sum_{\{i:m(i)=m\}} (s_{it} - \bar{s}_t)(\omega_{it} - \bar{\omega}_t) \\
&= \bar{\omega}_t + \underbrace{\sum_{\{i:m(i)=m\}} (s_{it} - \bar{s}_t)(\omega_{it} - \bar{\omega}_t)}_{\equiv \Gamma_t}.
\end{aligned} \tag{15}$$

Table 7: Decomposition Results

		1982	1987	1992
Output-weighted Productivity ($\tilde{\omega}_t$)	mean	1.00	1.13	1.19
	s.d.	0.17	0.19	0.17
Unweighted Average Productivity ($\bar{\omega}_t$)	mean	0.97	1.11	1.17
	s.d.	0.16	0.15	0.17
OP Covariance (Γ_t)	mean	0.03	0.02	0.03
	s.d.	0.08	0.10	0.07

Notes: This table contains decomposition results for the OP static decomposition at the CEA-Year level. See equation (15) for construction of terms of the decomposition.

The term $\bar{\omega}_{mt}$ is unweighted average productivity. The relationship between $\bar{\omega}_{mt}$ and competition was the focus of Section 4. The second term in the decomposition, Γ_{mt} , is the covariance of establishment-level productivity and market share. Though the reallocation mechanism affects $\bar{\omega}_{mt}$ via the selection effect, it also affects *weighted* average productivity via this additional channel. Results for the decomposition, which I conduct separately at the year-CEA level, are presented in Table 7. Two features stand out— first, note that there is substantial variation in $\bar{\omega}_{it}$ at the year-CEA level. This offers some ex-post motivation for aggregating to the CEA level in order to allow market-level random effects in our quantile regression analysis of Section 4.3. Second, there is substantial productivity growth in the period which seems to be mostly driven by *unweighted* average productivity changes. I observe little change in Γ_{mt} on average, but note that there is substantial variation across

Table 8: Decomposition Regression Results

	Dependent Variable: $\tilde{\omega}_t$			
	(1)	(2)	(3)	(4)
No. Estab./mi. ²	0.0463*** (0.0086)			
No. Firms/mi. ²		0.0482*** (0.0092)		
HHI No. Estab./mi. ²			0.0459*** (0.0087)	
HHI No. Firms/mi. ²				0.0513*** (0.0103)
Observations	847	847	847	847
R^2	0.2026	0.2012	0.1931	0.1783
	Dependent Variable: $\bar{\omega}_t$			
	(5)	(6)	(7)	(8)
No. Estab./mi. ²	0.0434*** (0.0078)			
No. Firms/mi. ²		0.0455*** (0.0083)		
HHI No. Estab./mi. ²			0.0426*** (0.0079)	
HHI No. Firms/mi. ²				0.0483*** (0.0093)
Observations	847	847	847	847
R^2	0.2467	0.2456	0.2415	0.2261
	Dependent Variable: Γ_t			
	(9)	(10)	(11)	(12)
No. Estab./mi. ²	0.0029 (0.0036)			
No. Firms/mi. ²		0.0026 (0.0037)		
HHI No. Estab./mi. ²			0.0034 (0.0036)	
HHI No. Firms/mi. ²				0.0030 (0.0041)
Observations	847	847	847	847
R^2	0.0018	0.0016	0.0001	0.0000

Notes: This table presents IV results for the effect of competition on components of the OP weighted average productivity decomposition. Models (1)-(4) use p_t , weighted average productivity at the CEA-year level; (5)-(8) use unweighted average productivity at the CEA-year level, and (9)-(12) use the covariance of market share and productivity at the CEA-year level. Year-specific constants are included but not reported and standard errors are clustered at the CEA level. Instruments include the number of building permits per square mile, the number of single family building permits per square mile, and local government road and highway expenditure per square mile. For reference, * signifies $p \leq 0.05$.

markets which may be correlated with competition.²⁰

In order to test the reallocation hypothesis I regress p_{mt} , $\bar{\omega}_{mt}$, and Γ_{mt} on my measures of competition, instrumented using demand density measures as before. Results for these instrumental variables regression are presented in Table 8. I find that the effect of competition on unweighted productivity tracks the effect of competition on output-weighted productivity across all specifications. Most importantly, I find no effect of competition on the covariance of shares and output—in other words, for ready-mix concrete there is no evidence that competition reallocates shares from less- to more-productive establishments. Since reallocation is the mechanism which drives the selection effect hypothesis, this explains why we find no evidence for the selection effect in Section 4.

6 Discussion

The objective of this paper has been to test two competing hypotheses that could explain the positive correlation between productivity and competition. Nested in each of these hypotheses — the treatment effect and the selection effect of competition — is a vast theoretical literature, so it is a first-order question to know which is consistent with the data. I have offered three pieces of evidence on this point: a semi-parametric selection correction, a grouped IV quantile approach, and a direct test of the reallocation hypothesis in ready-mix concrete. All of my evidence points to the conclusion that competition has a treatment effect on measured productivity.

One can and should wonder whether these results would hold in other industries, beyond the special case of ready-mix concrete. To this point I would note that ready-mix concrete exhibits tremendous churn: roughly one third of observed establishments exit between each census. As exit is the operative margin of the selection effect, ready-mix concrete seems to

²⁰Many papers in the productivity literature follow Baily et al. (1992) in decomposing productivity changes into four parts: changes among surviving firms, reallocation among surviving firms, exiting firms, and entrants. Collard-Wexler and De Loecker (2013) adopt this approach (but note the problems below), and Melitz and Polanec (2013) use it to develop a dynamic version of the OP decomposition. Unfortunately (and somewhat surprisingly), distinguishing between surviving, exiting, and entrant firms may confuse rather than clarify the treatment versus selection effect question. The treatment effect is not confined to surviving firms; instead, it will affect exiting and entrant firms as well—differentially insofar as it affects the exit threshold. Similarly, the selection effect is not confined to exiting firms; it will also bias the sample of surviving firms in favor of those that have positive productivity changes, as those with negative productivity changes are more likely to cross the (higher) exit threshold. What is unique to the selection effect hypothesis is *reallocation*, which is most clearly captured by the covariance term of the OP decomposition I use here.

be a particularly likely industry to find it. However, I find no evidence for reallocation or selection.

Taking for granted the external validity of my finding — that the treatment effect is the primary channel through which competition affects productivity — the first-order question becomes understanding its source. This suggests that research on the within-firm determinants of productivity, e.g. human resources practices (Ichniowski et al., 1997), technology adoption (Van Biesebroeck, 2003), and the organization of the firm (Bloom and Van Reenen, 2010), as well as the response of these factors to competition, is of first-order significance for understanding productivity dynamics.

References

- Akerberg, D., Caves, K., and Frazer, G. (2006). Structural identification of production functions. Working Paper.
- Adelman, M. A. (1969). Comment on the “h” concentration measure as a numbers-equivalent. *Review of Economics and Statistics*, 51(1):99–101.
- Asplund, M. and Nocke, V. (2006). Firm turnover in imperfectly competitive markets. *Review of Economic Studies*, 73(2):295–327.
- Baily, M., Hulten, C., and Campbell, D. (1992). Productivity dynamics in manufacturing plants. *Brookings Papers on Economic Activity: Microeconomics*, 4:187–267.
- Barseghyan, L. and DiCecio, R. (2011). Entry costs, industry structure, and cross-country income and tfp differences. *Journal of Economic Theory*, 146(5):1828–1851.
- Bartelsman, E. J. and Doms, M. (2000). Understanding productivity: Lessons from longitudinal microdata. *Journal of Economic Literature*, 38(3):569–594.
- Bloom, N. and Van Reenen, J. (2010). Why do management practices differ across firms? *Journal of Economic Perspectives*, 24(1):203–224.
- Boone, J. (2008). A new way to measure competition. *The Economic Journal*, 118(531):1245–1261.
- Bottasso, A. and Sembenelli, A. (2001). Market power, productivity, and the eu single market program: Evidence from a panel of italian firms. *European Economic Review*, 45(1):167–186.
- Caves, R. and Barton, D. (1990). *Efficiency in US Manufacturing Industries*. MIT Press.
- Chernozhukov, V. and Hansen, C. B. (2005). An iv model of quantile treatment effects. *Econometrica*, 73(1):245–261.
- Chetverikov, D., Larsen, B., and Palmer, C. (2013). Iv quantile regression for group-level treatments, with an application to the effects of trade on the distribution of wages. Working Paper.
- Collard-Wexler, A. (2011). Productivity dispersion and plant selection in the ready-mix concrete industry. Working Paper.

- Collard-Wexler, A. (2013). Demand fluctuations in the ready-mix concrete industry. *Econometrica*, 81(3):1003–1037.
- Collard-Wexler, A. and De Loecker, J. K. (2013). Reallocation and technology: Evidence from the us steel industry. Working Paper.
- Dertouzos, M. L., Lester, R., and Solow, R. M. (1989). *Made in America: Regaining the Productive Edge*. MIT Press, Cambridge, MA.
- Dunne, T., Klimek, S., and Schmitz, J. (2010). Competition and productivity: Evidence from the post wii us cement industry. Working Paper.
- Ericson, R. and Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies*, 62(1):53–82.
- Fabrizio, K. R., Rose, N. L., and Wolfram, C. D. (2007). Do markets reduce costs? assessing the impact of regulatory restructuring on us electric generation efficiency. *American Economic Review*, 97(4):1250–1277.
- Foster, L., Haltiwanger, J., and Syverson, C. (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, 98(1):394–425.
- Graham, D. R., Kaplan, D. P., and Sibley, D. S. (1983). Efficiency and competition in the airline industry. *Bell Journal of Economics*, 14:118–38.
- Green, A. and Mayes, D. (1991). Technical inefficiency in manufacturing industries. *The Economic Journal*, 101(406):523–538.
- Hay, D. A. and Liu, G. S. (1997). The efficiency of firms: What difference does competition make? *The Economic Journal*, 107(442):597–617.
- Hicks, J. R. (1935). Annual survey of economic theory: The theory of monopoly. *Econometrica*, 3(1):1–20.
- Holmes, T. J. and Schmitz, J. A. (2010). Competition and productivity: A review of evidence. *Annual Review of Economics*, 2:619–642.
- Holmström, B. (1982). Moral hazard in teams. *Bell Journal of Economics*, 13(2):324–340.
- Hopenhayn, H. A. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica*, 60(5):1127–1150.

- Ichniowski, C., Shaw, K., and Prennushi, G. (1997). The effects of human resource management practices on productivity: A study of steel finishing lines. *American Economic Review*, 87(3):291–313.
- Johnson, K. P. (1995). Redefinition of the bea economic areas. *Survey of Current Business*, 75:75–81.
- Klette, T. J. and Giliches, Z. (1996). The inconsistency of common scale estimators when output prices are unobserved and endogenous. *Journal of Applied Econometrics*, 11:343–361.
- Kneller, R. and McGowan, D. (2014). Productivity-enhancing demand shocks: Technology adoption during the u.s. ethanol boom: Technology adoption during the u.s. ethanol boom. Working Paper.
- Leibenstein, H. (1966). Allocative efficiency vs. “x-efficiency”. *American Economic Review*, 56(3):392–415.
- Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, 70(2):317–342.
- Melitz, M. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725.
- Melitz, M. and Ottaviano, G. (2008). Market size, trade, and productivity. *Review of Economic Studies*, 75:295–316.
- Melitz, M. and Polanec, S. (2013). Dynamic olley-pakes productivity decomposition with entry and exit. Working Paper.
- Nickell, S. J. (1996). Competition and corporate performance. *Journal of Political Economy*, 104(4):724–746.
- Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297.
- Pavcnik, N. (2002). Trade liberalization, exit, and productivity improvements: Evidence from chilean plants. *Review of Economic Studies*, 69(1):245–276.
- Perelman, M. (2011). Retrospectives: X-efficiency. *Journal of Economic Perspectives*, 25(4):211–222.

- Porter, M. E. (1990). *The Competitive Advantage of Nations*. Macmillan, London.
- Raith, M. (2003). Competition, risk, and managerial incentives. *American Economic Review*, 93(4):1425–1436.
- Schmidt, K. M. (1997). Managerial incentives and product market competition. *Review of Economic Studies*, 64(2):191–213.
- Schmitz, J. (2002). Competitive pressure and labor productivity: World iron ore markets in the 1980s. *American Economic Review*, 92(4):745–760.
- Schmitz, J. (2005). What determines productivity? lessons from the dramatic recovery of the u.s. and canadian iron ore industries following their early 1980s crisis. *Journal of Political Economy*, 113(3):582–625.
- Sivadasan, J. (2009). Barriers to competition and productivity: Evidence from india. *The B.E. Journal of Economic Analysis & Policy*, 9(1):1–66.
- Stigler, G. J. (1976). The existence of x-efficiency. *American Economic Review*, 66(1):213–216.
- Sutton, J. (1991). *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*. MIT Press.
- Syverson, C. (2004). Market structure and productivity: A concrete example. *Journal of Political Economy*, 112(6):1181–1222.
- Syverson, C. (2008). Markets: Ready-mixed concrete. *Journal of Economic Perspectives*, 22(1):217–234.
- Van Biesebroeck, J. (2003). Productivity dynamics with technology choice: An application to automobile assembly. *Review of Economic Studies*, 70(1):167–198.
- Williamson, O. E. (1968). Economies as an antitrust defense: The welfare tradeoffs. *American Economic Review*, 58(1):18–36.

Appendices

A Data

There are over five thousand ready-mix concrete establishments observed by the Census of Manufactures in each year of my sample. Unfortunately, roughly one-third of these establishments are “administrative records” establishments; that is, small enough to be exempt from completing long-form census surveys. Census data for these establishments do exist, but they are generated from a combination of administrative records from other agencies and imputation, the latter of which makes them unusable for establishment-level productivity analysis. I therefore exclude all establishments that fall into this category.

A second constraint is that a small handful of establishments in my sample are extensively diversified and operate in multiple SIC codes. This makes it difficult to construct a productivity residual for the ready-mix concrete portion of their business because data on their inputs are pooled at the establishment level, across all lines of production. I deal with this by excluding establishments for which less than fifty percent of their total sales is from ready-mix concrete. For diversified establishments that survive this exclusion, inputs devoted to ready-mix concrete are approximated by scaling the conflated input variable by the fraction of sales revenue from ready-mix concrete.

Finally, I measure the establishment-level price of a cubic yard of concrete by dividing sales by quantity. I observe a small number of firms with extremal values, presumably generated by misreporting, and exclude them from the sample.

It is important to note that while these establishments are excluded for regressions that depend on estimates of the productivity residuals, they are not excluded in the calculation of market-level variables; in particular, the competition indexes discussed below in Section 3.3 use data for all establishments.

In order to construct productivity estimates as described in Section 3.4, I use data from the Census of Manufactures at the establishment level for inputs including: energy, materials, equipment capital, structural capital, and labor. Energy and materials inputs are captured by expenditures reported by the Census of Manufacturers divided by two-digit deflators from the NBER productivity database. Equipment and structural capital are the reported

book value multiplied by two-digit capital rental rates from the Bureau of Labor statistics. Finally, labor is taken as the number of production labor hours multiplicatively adjusted by the ratio of total wages to production wages.

B Theoretical Motivation

Consider the following variation on an entry game. In the first stage of the game there is an infinite mass of potential entrants. They are ex-ante identical; they know nothing of their type. If they enter, they pay a cost of entry c_e . Let λ denote the endogenous mass of entrants in state 1. In the second stage entrants learn their idiosyncratic types, denoted ϕ , which are distributed i.i.d. according to a continuous distribution G on $[0, 1]$. At this point they make a second choice. They may exit or stay active; if they exit they receive a payoff normalized to zero. Let μ denote the measure of active, non-exiting firms on the type space $[0, 1]$. Those firms make obtain profits given by $\pi(\phi, C(\mu), D)$, where D is a demand shifter and C is a continuous function mapping measures μ into \mathbb{R} . C represents an ideal competition index. I assume that C is increasing in the following partial order: if $\mu \geq \mu'$ on $[0, 1]$ and $\mu' > \mu$ for some open set in $[0, 1]$ then $C(\mu') > C(\mu)$. I also assume that π is continuously differentiable in all of its arguments, strictly increasing in ϕ and D , and strictly decreasing in C .

This model is in the spirit of Hopenhayn (1992) and Asplund and Nocke (2006), however I have abstracted away from dynamics for simplicity. Consistent with those models, monotonicity of π implies that the exit decision in the second stage follows a threshold rule; let us call it $\bar{\phi}$. An equilibrium is a pair $\langle \lambda, \bar{\phi} \rangle$ such that:

$$\int_{\bar{\phi}^*}^1 \pi(\phi, C(\mu), D) dG(\phi) = c_e \tag{E1}$$

$$\pi(\bar{\phi}, C(\mu), D) dG(\phi) = 0 \tag{E2}$$

Equilibrium condition (E1) reflects optimal choice by potential entrants ex ante, while condition (E2) reflects the ex post exit choice of entrant firms. First, note that an equilibrium exists and is unique under mild conditions.

Proposition 1. *Let $\psi \equiv \{\phi : \pi(\phi, C(0), D) = 0\}$. If*

1. $\int_{\psi}^1 \pi(\phi, C(0), D) dG(\phi) > c_e$, and
2. there exists a finite measure $\tilde{\mu}$ on $[0, 1]$ such that $\pi(1, C(\tilde{\mu}), D) < 0$,

then there exists a unique equilibrium with a nonempty market (i.e., such that $\mu \neq 0$).

Proof. I omit the proof of existence, which follows directly from Conditions 1 and 2 and application of the Schauder Fixed Point theorem.

To see uniqueness, suppose by way of contradiction that there were two equilibria at $\langle \lambda, \bar{\phi} \rangle$ and $\langle \lambda', \bar{\phi}' \rangle$. From (E2), we have $\bar{\phi}' > \bar{\phi} \Leftrightarrow \lambda' > \lambda$. Suppose without loss of generality that $\bar{\phi}' > \bar{\phi}$. However, now $\int_{\bar{\phi}'}^1 \pi(\phi, C(\mu'), D) < \int_{\bar{\phi}}^1 \pi(\phi, C(\mu), D) dG(\phi) = c_e$, which contradicts the claim that $\langle \lambda', \bar{\phi}' \rangle$ is an equilibrium. \square

The key result is a comparative static in D , which stands in for demand shifters. I am interested in showing how the optimal threshold $\bar{\phi}$ moves as the market size grows. In order to prove this I need two more assumptions. The first is an innocuous assumption that says that when the market grows, the profits of high type firms grow no less than proportionately.

Assumption 3. For $\phi' > \phi$, $D' > D$, and any C ,

$$\frac{\pi(\phi', C, D')}{\pi(\phi, C, D')} \geq \frac{\pi(\phi', C, D)}{\pi(\phi, C, D)}.$$

The second assumption is more critical, and embodies the reallocation hypothesis:

Assumption 4. For $\phi' > \phi$, $C' > C$, and any D ,

$$\frac{\pi(\phi', C', D)}{\pi(\phi, C', D)} > \frac{\pi(\phi', C, D)}{\pi(\phi, C, D)}.$$

This assumption is critical to the selection hypothesis. On the theory side, Asplund and Nocke (2006) argue that it is consistent with many standard models. Going a step further, Boone (2008) argues that this is in fact constitutive of our very idea of competition. Finally this mechanism has become central in trade and productivity analysis literatures as well. In Proposition 2 I show that it implies the selection effect hypothesis.

Proposition 2. If $D' \geq D$, then $\bar{\phi}' \geq \bar{\phi}$

Proof. First, note that $C(\mu') \geq C(\mu)$. Suppose, by way of contradiction, otherwise. Then (E2) implies that $\bar{\phi}' < \bar{\phi}$. Now, $\int_{\bar{\phi}'}^1 \pi(\phi, C(\mu'), D') > \int_{\bar{\phi}^*}^1 \pi(\phi, C(\mu), D) dG(\phi) = c_e$, which generates a contradiction.

Next observe that, for (E1) to hold, there must exist $\tilde{\phi}$ such that $\pi(\tilde{\phi}, C(\mu'), D') = \pi(\tilde{\phi}, C(\mu), D) > 0$. Moreover, for $\phi' < \tilde{\phi}$,

$$\begin{aligned} \frac{\pi(\phi', C(\mu'), D')}{\pi(\tilde{\phi}, C(\mu'), D')} &< \frac{\pi(\phi', C(\mu), D)}{\pi(\tilde{\phi}, C(\mu), D)} \\ \Rightarrow \pi(\phi', C(\mu'), D') &< \pi(\phi', C(\mu), D). \end{aligned} \tag{16}$$

The first line follows from log increasing differences and the complementarity of D and $C(\mu)$. The second line follows from $\pi(\tilde{\phi}, C(\mu'), D') = \pi(\tilde{\phi}, C(\mu), D) > 0$, and implies that $\bar{\phi}' < \bar{\phi}$. \square