Why is Productivity Correlated with Competition?*

Matthew Backus†

June 30, 2020

Abstract

The correlation between productivity and competition is an oft–observed but incompletely–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, using a standard decomposition approach, I look for evidence of greater reallocation of demand to productive plants in more competitive markets. Second, I model the establishment exit decision and construct a semi–parametric selection correction to quantify the empirical significance of treatment and selection. Finally, I use a grouped IV quantile regression to test the distributional predictions of the selection hypothesis. I find no evidence for greater selection or reallocation in more competitive markets; instead, all three results suggest that measured productivity responds directly to competition. Potential channels include specialization and managerial inputs.

*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 Ackerberg, 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 numerous seminar participants 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.

†Columbia University, CEPR, and NBER. E-mail: matthew.backus@columbia.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 highlights 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 different quantities of output.

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. I refer to them as, 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 though selective attrition of low–productivity establishments in highly competitive markets; to wit, it posits 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 construct a way to test how competition affects productivity. Rather than looking at a handful of changes over time, it uses variation across hundreds of markets in the ready–mix concrete industry. And rather than using point estimates from an accounting decomposition, it uses the covariance of plant–level productivity estimates and the extent of competition in the local market. This allows me to show how these two hypotheses—the treatment effect and the selection effect of competition—are separable in the data. I develop three alternative approaches, two novel, all of which yield the same conclusion.

First, I use a standard decomposition of output–weighted productivity at the market level into average and compositional effects. The latter is informative about the intensive margin, reallocation, but not the extensive margin, i.e. selection via firm exit.

Next, 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 the literature on semi–parametric production function estimation.[1]

Finally, I use a grouped instrumental variables quantile regression approach at the market

[1] In particular, my identification result depends on timing assumptions related to those advanced in Olley and Pakes (1996); Levinsohn and Petrin (2003), and Ackerberg et al. (2015).

level; this allows me to identify the marginal effect of competition on productivity at all quantiles of the productivity distribution. I use this to test the hypothesis, implied by selective attrition of low–productivity establishments, 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 weak: 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 them 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 total factor productivity (TFP) more directly. Following the language of Foster et al. (2008), I therefore use TFPQ (Q for quantity) rather than TFPR (R for revenues). This averts the issue raised in Klette and Giliches (1996), that revenue–based output measures include equilibrium markups and are therefore mechanically correlated with competition, a concern that would be particularly problematic for my application.[2]

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, nor do I observe greater reallocation in more competitive markets. 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.[3] My results show that the *degree* of market selection on physical productivity is not driven by differences in competition—more precisely, not enough

[2]The availability of physical output data handles the problem of output price bias, but there remains a possibility of input price bias if input prices are systematically different across firms. See De Loecker and Goldberg (2014) for a recent discussion of input and output price bias.

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

to explain the observed correlation between competition and productivity.

These results suggest productive directions for future research. If reallocation and selection does not drive the correlation, what are firms doing differently in more competitive markets? I am able to offer some limited evidence: the data here suggest greater specialization, greater managerial inputs, and are not consistent with a story based on capacity utilization. Taken together, this suggests that, in light of the growing policy interest in understanding the effects of concentration and market power, a productive direction is to look within the firm, and not exclusively at the allocative and re–allocative effects of competition. Within–firm adjustment mechanisms may be of first–order importance for understanding productive efficiencies from competition, gains from trade, and establishment–level productivity dispersion.

Section 2 briefly summarizes related literature to situate 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 reduced–form results: I replicate the standard finding of a correlation between productivity and competition, introduce my IV strategy, and consider the role of "reallocation." Section 5 contains the main methodological contributions of the paper: the structurally-motivated selection correction and the quantile IV results. In Section 6, I address several robustness concerns and offer some evidence on mechanisms. Section 7 concludes.

2 Related Literature

This paper draws on several bodies of work. The first elaborates on the treatment effect hypothesis: this includes empirical documentation and a theoretical literature offering a rigorous foundation. The second is related to the empirical and theoretical basis of the selection effect. Finally, I briefly discuss recent work 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) foreshadows the notion of managerial slack when he notes 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.[4] 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?

Subsequently, a variety of explanations were proposed. Nalebuff and Stiglitz (1983), building on Holmström (1982), develops a model of compensation in which competition generates informational externalities. In their model, more competition may enable the principal to extract more effort from the agent. Alternatively, in Schmidt (1997), the possibility of bankruptcy offers a direct check on shirking. Finally, Raith (2003) considers an alternative approach, directly modeling the returns to investment in cost reduction in an equilibrium with entry and exit. Any of these models could rationalize the finding of a treatment effect of competition on productivity.

Moreover, there is a body of empirical work that documents within–firm productivity changes in response to competition. 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 greater 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 decide 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

[4]See Perelman (2011) for a summary of the Leibenstein-Stigler debate concerning X–Efficiency.

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 a treatment effect. This prediction depends on a key assumption on stage–game profits, which has been called "reallocation:" 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.[5] In the empirical literature, this reallocation mechanism has been considered an important channel for productivity growth since Olley and Pakes (1996); Boone (2008) goes still further, arguing that reallocation is constitutive of competition, not merely related to it. In Online Appendix B, I offer a simple entry model to illustrate the relationship between reallocation and selection.

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). A recent contribution by Bloom et al. (2016) revisits the trade shock literature using firm–level data from Europe in the context of China's entry into the World Trade Organization and finds substantial productivity effects, but through a different mechanism: technological upgrading and the release of "trapped" factors of production, with a distinctly Ricardian flavor.

Closest to this work is a subset of the literature that has focused on industries where we can measure TFPQ rather than TFPR. Notable contributions include Foster et al. (2008) using data on several industries where measures of physical output are available, Collard-Wexler and De Loecker (2015) on steel, and Asker et al. (2019) on oil production. The leading example of such an industry, however, is ready–mix concrete. This paper builds on

[5]The 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)

the pioneering work of Syverson (2004), which studied the negative correlation of productivity dispersion with competition. That paper sidestepped the question of mechanisms and assumes selection; this one steps back and reconsiders the treatment effect.[6] 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.

3 Data and Measurement

3.1 Ready–Mix Concrete

I use US Census of Manufacturers (CMF) data for the ready–mix concrete industry (SIC 3273) from the 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

[6]From 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), for example. 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."

the CMF 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.[7] 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 primarily by cross–sectional variation in market competitiveness. This is motivated by the local character of competition, and therefore necessitates 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 the 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.

The use of CEAs as market definition is motivated by the fact that they are based on real economic behavior. Labor force commuting patterns are a good proxy for market areas when population is dense, however in rural market areas they are sometimes implausibly large.For this reason I exclude from my sample all plants in the top decile of the CEA geographic size distribution.[8]

[7] Consistent with Dunne et al. (1988), entry and exit within a market are highly correlated for ready–mix concrete (regression results available on request), suggesting that most entry and exit is related to churn rather than growing or expanding markets.

[8] See Online Appendix A.2 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. My baseline and preferred specification is the number of ready–mix concrete establishments per square mile. This can be derived from a formal model of travel costs and geographic differentiation as in Syverson (2004), but the intuition is simple and depends on a critical feature of the industry: that ready–mix concrete plants are only differentiated geographically, by the transit costs of serving a customer. Therefore as demand grows, there are returns to opening new plants rather than merely expanding capacity. Then, as demand grows and more firms enter a finite geographic area, the geographic density of establishments increases and they are correspondingly less differentiated. In that sense, establishments become more substitutable and face greater competition.

As the number of competing plants is endogenous to plant productivity, I instrument using demand shifters to isolate variation in long–run competitive structure.[9] If establishment–level productivity shocks are transient, a point for which there is extensive evidence (Baily et al., 1992; Foster et al., 2008), then we can integrate out over these short–run shocks by using an instrument for the long–run expected competitive structure.[10] 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.[11] In practice, these enter my regressions as log densities, i.e. divided by

[9]By "long–run" here I mean the steady-state expectation with endogenous entry and exit.

[10]This 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.

[11]The reason I diverge from prior work and eschew county–level construction employment as a demand shifter for ready–mix concrete is that it consistently failed overidentification tests for exogeneity in my IV regressions. 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 overidentification tests easily; see the overidentification test results (Hansen J) from Table 2. In the discussion of robustness in Section 6.2 I consider IV regressions with additional controls, including county construction employment.

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
log density	-6.02	0.85	-6.04	0.86	-6.04	0.86
No. Firms	9.51	7.45	10.46	9.11	10.43	9.23
log density	-6.23	0.83	-6.28	0.84	-6.35	0.85
HHI No. Estab.	6.84	5.28	6.99	6.48	7.03	6.40
log density	-6.61	0.88	-6.63	0.90	-6.61	0.90
HHI No. Firms	4.92	3.05	5.50	3.55	5.26	3.69
log density	-6.86	0.85	-6.92	0.85	-6.95	0.88
Building Permits	2645.51	5629.57	4291.51	7642.50	2951.70	4165.11
log density	-1.25	1.23	-0.89	1.62	-0.96	1.33
S.F. Building Permits	1419.62	2662.46	2851.15	4699.00	2443.78	3587.64
log density	-1.84	1.28	-1.22	1.55	-1.17	1.33
Road & Hwy $	38994.95	59188.68	56030.91	93206.21	72173.89	121158.15
log density	1.72	1.07	2.05	1.09	2.30	1.11

Notes: This table contains summary statistics for measures of competition $c_{m(i)t}$ as well as demand shifters for ready–mix concrete plants in CMF years 1982, 1987, and 1992. All variables are aggregated to the year–CEA level. For each variable I also report the log of the ratio of the variable to the area (in square miles) of the CEA, which is the form in which they enter the analysis that follows.

the number of square miles in the CEA and logged. The demand shifter data are generated at the county level by the US Census.[12]

I supplement my preferred measure, the number of ready–mix concrete establishments per square mile, with two additional variations as robustness checks. The first is the number equivalent of the Herfindal-Hirschman Index (HHI), i.e. HHI^{-1}. Originally proposed by Adelman (1969), it can be interpreted as the number of symmetric firms that would generate the observed HHI. A potential criticism of my preferred measure is that it ignores variation in market share, 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)). The numbers equivalent HHI is a conventional, if arbitrary, way to incorporate that. Second, in some of my markets there are multiple establishments owned by the same firm. Without additional 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.

Summary statistics for my competition variables as well as my market size instruments are

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

presented in Table 1. They are described both in raw form as well as the log density form in which they enter the regressions below. For my preferred specification, the count of establishments, we see that a standard deviation increase implies an approximately eighty-five percent increase in the number of ready–mix concrete plants per square mile.

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:

$$q_{it} = \alpha_{Lt} l_{it} + \alpha_{K^E t} k^E_{it} + \alpha_{K^S t} k^S_{it} + \alpha_{Mt} m_{it} + \alpha_{Et} e_{it} + \omega_{it}, \tag{1}$$

where q_{it} is log physical output; l_{it} is labor, k^E_{it} and k^S_{it} are equipment and structural capital, respectively; m_{it} is expenditure on materials; e_{it} is energy expenditures, and the coefficients α are input elasticities, with $\sum_k \alpha_{kt} = 1$.[13,14] The input elasticities α_t are consistently measured by input shares under the assumption that all inputs are flexible and homogenous, have common prices (excluding labor), that there are constant returns to scale, and that purchasers do not have market power.[15]

Estimation of the year–specific industry–wide input cost shares, which are taken as estimates of the input elasticities, follows the approach of Foster et al. (2008), detailed in Online Appendix A.1.[16]

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

[13]The 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 the assumtpion as a robustness check in Section 6.1.

[14]For ready–mix concrete, trucks are the salient example of equipment capital while buildings are the salient example of structural capital. See the Annual Capital Expenditures Survey Instructions, Definition, and Codes List for a detailed list of examples of the distinction.

[15]Both the concerned and the interested reader will find a healthy secondary market for ready–mix concrete capital on eBay.com, among other online auction platforms. With respect to input costs, see Atalay (2014) for a discussion of potential bias. To address this, much of the analysis is re-done in terms of labor productivity in Online Appendix Section A.2.3.

[16]Thanks to Chad Syverson for sharing these estimates.

autocorrelation of ω_{it} is 0.77. See Online Appendix A.2 for details on the components and measurement of x_{it}.

4 Competition and Productivity

In this section I present my reduced–form evidence on the correlation between competition and productivity, as well as a decomposition analysis that sheds some light on the response of reallocation to greater competition.

4.1 Reduced–Form Analysis

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

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

In this expression, ω_{it} is plant–level productivity, as defined in Section 3.4, and $c_{m(i)t}$ is a measure of competitiveness, as defined in Section 3.3. 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. The idea here is to integrate 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, my instruments for long–run demand are 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. To address this concern, I also present results from overidentification 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. Based on specification (5), a one standard deviation increase in the number of establishments per square mile (approximately 0.85, from Table 1), would predict a 3.98% increase in TFPQ. 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 overidentification tests consistently and safely fail to reject the hypothesis of exogeneity of the instruments. If demand had an indirect causal effect, whether due to mismeasurement of the competition index or some other hypothesis, e.g. agglomeration effects of market size, I should reject this hypothesis. The results support the choice of measure of competition, and also help to rule out alternative channels by which demand might affect productivity.

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 5.1 and 5.2.

4.2 Decomposition Approach

There is substantial variation in plant size, and so it is natural to wonder whether the average effects reported in Table 2 would be diminished if we weighted plants according to their output. And if we found such a difference, we should ask whether there is greater *reallocation* of demand to more productive plants in more competitive markets. To structure

Table 2: Competition and Productivity

	Dependent Variable: TFPQ (ω_{it})							
	OLS				Instrumental Variables			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log(No. Estab./mi.^{2})	0.0350*				0.0468*			
	(0.0056)				(0.0064)			
log(No. Firms /mi.^{2})		0.0318*				0.0494*		
		(0.0058)				(0.0072)		
log(HHI No. Estab./mi.^{2})			0.0302*				0.0481*	
			(0.0062)				(0.0066)	
log(HHI No. Firms /mi.^{2})				0.0214*				0.0561*
				(0.0060)				(0.0095)
Observations (rounded)	7400	7400	7400	7400	7400	7400	7400	7400
Clusters (rounded)	300	300	300	300	300	300	300	300
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, for ready–mix concrete plants in CMF years 1982, 1987, and 1992. 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 (in parentheses) 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. For reference, * signifies $p \leq 0.05$.

Table 3: Summary of Decomposition

		1982	1987	1992
Output–Weighted Productivity (p_{mt})	mean	1.00	1.13	1.19
	s.d.	0.17	0.19	0.17
Unweighted Average Productivity ($\bar{\omega}_{mt}$)	mean	0.97	1.11	1.17
	s.d.	0.16	0.15	0.17
Productivity–Output Covariance (Γ_{mt})	mean	0.03	0.02	0.03
	s.d.	0.08	0.10	0.07

Notes: This table contains decomposition results for the Olley and Pakes (1996) static decomposition at the CEA–year level for ready–mix concrete plants in CMF years 1982, 1987, and 1992. See equation (3) for construction of terms of the decomposition.

this exercise, I borrow a static decomposition of productivity from Olley and Pakes (1996) and write output–weighted average productivity, denoted p_{mt}, in terms of an unweighted mean and a covariance term:

$$\begin{aligned} p_{mt} &\equiv \sum_{\{i:m(i)=m\}} s_{it}\omega_{it} \\ &= \bar{\omega}_{mt} + \underbrace{\sum_{\{i:m(i)=m\}} (s_{it} - \bar{s}_{mt})(\omega_{it} - \bar{\omega}_{mt})}_{\equiv \Gamma_{mt}}. \end{aligned} \tag{3}$$

The term $\bar{\omega}_{mt}$, which appears on the last line, is unweighted average productivity. The second term in the final line, Γ_{mt}, is the covariance of establishment–level productivity and market share. This, in the literature on productivity, is taken as a direct measure of reallocation. I compute these terms at the CEA level for my dataset, and results are presented in Table 3. Two features stand out—first, note that there is substantial variation in $\bar{\omega}_{it}$ at the year–CEA level. This will be important for motivating the grouped quantile IV approach that I adopt in Section 5.2 below. 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 markets which may be correlated with competition.[17]

[17] Many 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 (2015) adopt this approach (but note the problems below), and Melitz

Typically, once the econometrician aggregates to the year–market level as the decomposition requires, only time series variation remains.[18] Here, however, the local character of ready–mix concrete markets offers cross–sectional variation as well, allowing us to put these decomposition terms in a regression framework. This means that we can do statistical, rather than eyeball, inference. I run the following simple regression:

$$y_{mt} = \beta_t + \beta c_{mt} + \varepsilon_{mt}, \tag{4}$$

which differs from equation 2 in two ways: first, the unit of observation is the market rather than the plant, and second, the left-hand side variable is $y_{mt} \in \{p_{mt}, \bar{\omega}_{mt}, \Gamma_{mt}\}$. As before, I include year fixed effects and standard errors are clustered at the CEA level. Results for variants of these regressions are presented in Table 4.

Reassuringly, the coefficients in Models (5)–(8) are close to those we found in the IV results of Table 2. Also, mechanically, the coefficients in, e.g., Models (5) and (9) add up to the coefficient in Model (1). Therefore the striking similarity of the coefficients in Models (1)–(4) and (5)–(8) implies that the coefficients in (9)–(12) must be very close to zero, and they are. This is the most striking feature of these results: that there is no statistically significant correlation between competition and Γ_{mt}. Even as ready–mix plants become more densely situated and therefore more substitutable, we are not able to detect a reallocation of demand to more productive plants. According to Foster et al. (2001), changes in Γ_{mt} made up approximately one third of average industry change in multifactor productivity between 1977 and 1987. At least in ready–mix concrete, this does not appear to be a response to changes in the competitive environment.

Already, this result bodes poorly for the selection hypothesis. Models that generate selection effects of competition hinge on an assumption that as a market becomes more competitive, demand is reallocated to more productive plants. If that assumption fails, then we may see

and Polanec (2015) use it to develop a dynamic version of the decomposition. Unfortunately, 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 decomposition I use here.

[18] This is the case in, e.g., Olley and Pakes (1996), where an increase in Γ_t following the divestiture of AT&T is evidence that reallocation was an important channel for productivity gains.

Table 4: Competition and Decomposed Productivity

	DV: Weighted Avg. Prod. (p_{mt})			
	(1)	(2)	(3)	(4)
log(No. Estab./mi.^{2})	0.0463*			
	(0.0086)			
log(No. Firms/mi.^{2})		0.0482*		
		(0.0092)		
log(HHI No. Estab./mi.^{2})			0.0459*	
			(0.0087)	
log(HHI No. Firms/mi.^{2})				0.0513*
				(0.0103)
Observations (rounded)	800	800	800	800
Clusters (rounded)	300	300	300	300
R^2	0.2026	0.2012	0.1931	0.1783
	DV: Unweighted Avg. Prod. ($\bar{\omega}_{mt}$)			
	(5)	(6)	(7)	(8)
log(No. Estab./mi.^{2})	0.0434*			
	(0.0078)			
log(No. Firms/mi.^{2})		0.0455*		
		(0.0083)		
log(HHI No. Estab./mi.^{2})			0.0426*	
			(0.0079)	
log(HHI No. Firms/mi.^{2})				0.0483*
				(0.0093)
Observations (rounded)	800	800	800	800
Clusters (rounded)	300	300	300	300
R^2	0.2467	0.2456	0.2415	0.2261
	DV: Prod.-Output Cov. (Γ_{mt})			
	(9)	(10)	(11)	(12)
log(No. Estab./mi.^{2})	0.0029			
	(0.0036)			
log(No. Firms/mi.^{2})		0.0026		
		(0.0037)		
log(HHI No. Estab./mi.^{2})			0.0034	
			(0.0036)	
log(HHI No. Firms/mi.^{2})				0.0030
				(0.0041)
Observations (rounded)	800	800	800	800
Clusters (rounded)	300	300	300	300
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 for ready–mix concrete plants in CMF years 1982, 1987, and 1992. Models (1)–(4) use p_{mt}, 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 (in parentheses) 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 $p \leq 0.05$.

no effect or even the opposite correlation. See Online Appendix B for a formal discussion on this point. However, the results in Table 4 offer an incomplete argument against selection. First, while Γ_{mt} has been advanced by empiricists as an intuitive counterpart of the reallocation mechanism, it does not correspond exactly to the assumptions required in the theory literature, which characterize profits directly, rather than output. Second, if the selection effect is operative, then it will positively bias the estimate of Models (1)–(4) and (5)–(8). For intuition, it is helpful to think of the effect of competition on Γ_{mt}, see Models (9)–(12), as a measure of the average infra–marginal effect of reallocation, whereas the selection effect is the extensive margin, driving firm exit. Without restrictive assumptions about how reallocation affects firms with different levels of productivity, measuring one does not offer us direct access to the other.

5 Treatment and Selection Effects of Competition

In this section I offer two perspectives on the conflation of the treatment and selection effects in the IV results of Table 2: one, more model–dependent, that allows me to quantitatively decompose treatment and selection, and a second, less so, that allows me to test the claim that the selection effect primarily generates the correlation.

5.1 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, i.e. a control function, 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 the relative contributions.

5.1.1 Behavioral Model

Consider a firm's exit decision in a model of Markov perfect 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 the entire history of $\mathcal{S}_t$. 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 $c_i(a_i, \mathcal{S}_t)$. All active establishments also make an exit decision (with scrap value normalized to zero), and potential establishments make an entry decision.

The order of play determines the informational content of firms' decisions, and is therefore central to identification, a feature which is 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)*

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

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 then all establishments make exit choices 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; the model is consistent with the notion that firms learn their new productivity draw *by producing.*

My second identifying assumption is the exogeneity of innovations in firms' productivity types, ϕ_{it}:

Assumption 2. *(exogeneity of innovations)*

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

This assumption implies that establishments do not make choices that affect the evolution of their idiosyncratic productivity type ϕ. It rules out, for instance, endogenous productivity due to research and development (see Doraszelski and Jaumandreu (2013) for a relaxation of the assumption in that spirit). Econometrically, the assumption is useful because it affords me a source of exogenous variation in the model.[19]

In this environment, the Bellman function of the active establishment, which chooses whether to exit and, should they choose to persist, actions a_{it} subject to costs $c(a_{it}, X_t)$, can be written:

$$V_i(\mathcal{S}_t) = \max\{0, \max_{a_{it}}\{\delta\mathbb{E}[\pi_i(\mathcal{S}_{t+1}) + V_i(\mathcal{S}_{t+1})|a_t] - c_i(a_{it}, \mathcal{S}_t)\}\}. \tag{5}$$

This Bellman equation nests two decisions, but the one I am 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:

$$\chi_i(\mathcal{S}_t) = \begin{cases} 1 & \text{if } v_i(\mathcal{S}_t) \geq 0 \\ 0 & \text{else.} \end{cases} \tag{6}$$

Moreover, in general it is possible to show that this exit choice can, in equilibrium, be written as a threshold rule.[20] 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.

5.1.2 Identification

As a straw man, consider the naïve structural interpretation of the linear model I estimated in Section 4.1:

$$\omega_{it} = \beta_t + \beta_c c_{m(i)t} + \phi_{it}. \tag{7}$$

[19]See Ackerberg et al. (2015) for a fuller discussion of this assumption.

[20]See Doraszelski and Satterthwaite (2010) for a discussion of the purification arguments required to guarantee existence of a threshold strategy in Markov games based on Ericson and Pakes (1995).

All that has changed from the earlier reduced–form model in (2) is notation and interpretation: the error term ε_{it} has become ϕ_{it}, the primitive, underlying productivity type of the plant, and β has become β_c, a treatment effect of competition on productivity. Viewed this way it is easy to see why the instrumental variables approach does not identify β_c. 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.[21] My identification argument shows that, subject to the assumptions of my model, the bias can be estimated, and β_c recovered.

The first step is to characterize the bias. First, take the expectation of both sides of equation (7) conditioning on the state of the market at time $t-1$ and the survival of the establishment:

$$\begin{aligned} &\mathbb{E}[\omega_{it}|\mathcal{S}_{t-1}, \phi_{it-1} \geq \phi_i^*(\mathcal{S}_{t-1})] \\ &\quad = \beta_t + \beta_c \mathbb{E}[c_{m(i)t}|\mathcal{S}_{t-1}, \phi_{it-1} \geq \phi^*(\mathcal{S}_{t-1})] + \mathbb{E}[\phi_{it}|\mathcal{S}_{t-1}, \phi_{it-1} \geq \phi_i^*(\mathcal{S}_{t-1})]. \end{aligned} \tag{8}$$

Focusing on the last term,

$$\begin{aligned} \mathbb{E}[\phi_{it}|\mathcal{S}_{it-1}, \phi_{it-1} \geq \phi_i^*(\mathcal{S}_{t-1})] &= \mathbb{E}[\phi_{it}|\mathcal{S}_{it-1}] \\ &= \mathbb{E}[\phi_{it}|\phi_{it-1}] \\ &= g(\phi_{it-1}). \end{aligned} \tag{9}$$

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 $\mathcal{S}_t$. Isolating ϕ_{it} and plugging this into (7) yields

$$\omega_{it} = \beta_t + \beta_c c_{m(i)t} + g_t(\phi_{it-1}) + \eta_{it}, \tag{10}$$

[21] The IV estimate is still *causal*, despite the fact that it conflates the two possible channels (treatment and selection), however it is uninterpretable in terms of primitives.

which can be evaluated by inverting the model in (7) and plugging it in for ϕ_{it-1}. This allows me to identify the treatment effect of competition on productivity, β_c. Note that the control function $g_t(\cdot)$ 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.

5.1.3 Estimation

I estimate the model using nonlinear GMM with two sets of moments. The first implements the regression in equation (10):

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

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 (7), which depends on the parameters of interest. To control for measurement error I continue to use my exogenous demand shifters as instruments for $c_{m(i)t}$.

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:

$$\underbrace{\omega_{it} - \hat{\beta}_c c_{m(i)t}}_{=\phi_{m(i)t}} = \alpha_t + \alpha_c c_{m(i)t} + u_{it}. \tag{11}$$

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 $\hat{\alpha}_c = 0$. Estimating this term allows me to quantify the relative contribution of the selection effect to the IV estimates obtained in Section 4.1.[22] This

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

Table 5: Treatment and Selection Effects of Competition

	Dependent Variable: TFPQ (ω_{it})			
	(1)	(2)	(3)	(4)
log(No. Estab./mi.^{2}) ($\hat{\beta}_c$)	0.0456*			
	(0.0090)			
log(No. Firms/mi.^{2}) ($\hat{\beta}_c$)		0.0488*		
		(0.0118)		
log(HHI No. Estab./mi.^{2}) ($\hat{\beta}_c$)			0.0503*	
			(0.0101)	
log(HHI No. Firms/mi.^{2}) ($\hat{\beta}_c$)				0.0571*
				(0.0150)
Selection Coeff. ($\hat{\alpha}_c$)	0.0046	0.0049	0.0005	0.0043
	(0.0030)	(0.0065)	(0.0041)	(0.0085)
Observations (rounded)	3100	3100	3100	3100
Clusters (rounded)	300	300	300	300

Notes: This table presents results for the semi–parametric selection correction procedure detailed in Section 5.1 using four different indices for competition using ready–mix concrete plants in CMF years 1982, 1987, and 1992. Year–specific constants are included but not reported and standard errors (in parentheses) are clustered at the CEA level. For reference, * signifies $p \leq 0.05$.

motivates the second set of moments I use in estimation, which implement the regression in equation (11):

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

I estimate the model using moments (M1) and (M2). Though they are estimated jointly, (M1) identifies the estimate of the treatment effect ($\hat{\beta}_c$) while (M2) identifies the estimate of the selection effect ($\hat{\alpha}_c$). Results are presented in Table 5. As a simple diagnostic check, since this is a decomposition of the causal effect of competition on productivity, the sum of $\hat{\beta}_c$ and $\hat{\alpha}_c$ should be roughly equal to my IV estimate from Table 2 for the corresponding competition index. None of my specifications reject this hypothesis.[23] 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 ready–mix concrete establishments per square mile, my estimates imply that doubling the number of establishments will raise output

[23]Note 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.

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.

In summary, the effect of competition on productivity seems to be driven by a within–firm response—the treatment effect—instead of market selection driven by reallocation of demand—the selection effect. However, the argument here is subject to the very strong, if conventional, assumptions of the model: timing and exogeneity of productivity innovations. To strengthen this result, in Section 5.2 below I adopt a less model–dependent approach to looking for evidence of a selection effect.

5.2 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.

5.2.1 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 1. The distribution of productivity types is assumed to be left–truncated, which reflects a threshold rule for exit common in industry dynamics models. Figure 1a depicts an additive shift in the entire distribution, consistent with the linear model described in (2). Alternatively, Figure 1b 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 1: Treatment Effect vs. Selection Effect on Distribution of Productivity

(a) Treatment Effect

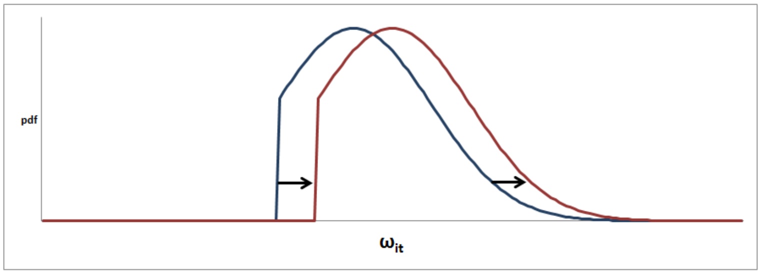

(b) Selection Effect

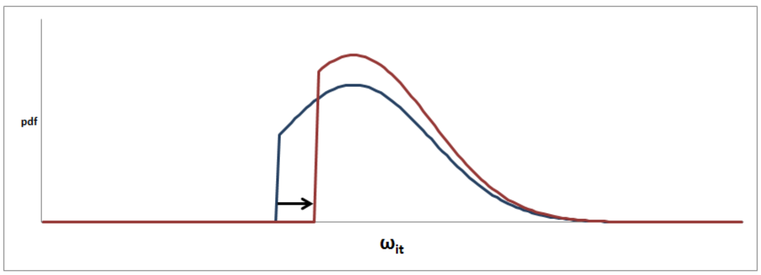

Notes: This figure illustrates two changes in the distribution of productivity residuals; in panel (a) a linear shift of the entire distribution, and in panel (b) a shift in the left truncation point. These correspond to the treatment effect and selection effect hypotheses, respectively.

It is important to note that the common shift depicted in Figure 1a, 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.[24] This test is less decisive than the semi–parametric approach in Section 5.1, but has the advantage that it requires correspondingly fewer assumptions and offers a simple visual interpretation.

[24]Note 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 nonetheless the left tail of the productivity distribution should exhibit more sensitivity.

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

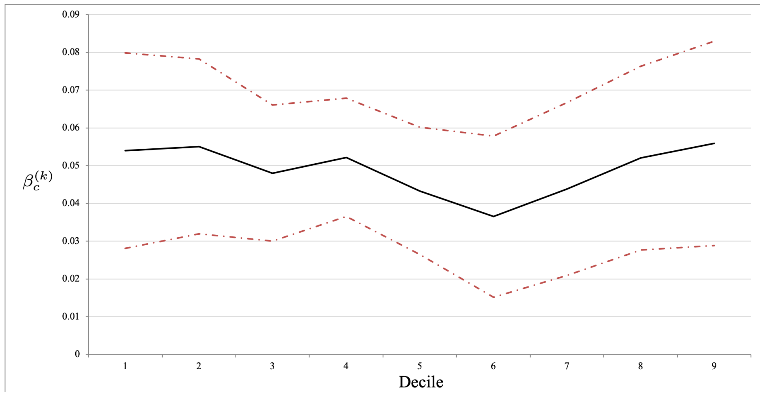

Notes: Here I present graphically the results from Table 6 for the case where the density of ready–mix establishments is taken as the competition index, for ready–mix concrete plants in CMF years 1982, 1987, and 1992. Equivalent diagrams for alternative competition indices are qualitatively similar. The dashed lines represent a 95% confidence interval.

5.2.2 Estimation

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

$$\rho_{mt}^{(k)} = \beta_t^{(k)} + \beta_c^{(k)} c_{mt} + \nu_{mt}. \tag{12}$$

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.[25]

There is, however, a subtle source of bias in (12) as an empirical model: as we see from Table 1, in many of my markets there are only a handful of establishments. This introduces an incidental parameters problem (Neyman and Scott, 1948). Fortunately, we can characterize

[25] See Chetverikov et al. (2016) for further discussion of the grouped IV quantile regression approach.

Table 6: Competition and Productivity by Grouped Quantile IV

	Dependent Variable: Decile k of the TFPQ Distribution ($\rho_{mt}^{(k)}$)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
log(No. Estab./mi.^{2})	0.0540*	0.0551*	0.0480*	0.0522*	0.0433*	0.0365*	0.0438*	0.0520*	0.0559*
	(0.0132)	(0.0118)	(0.0092)	(0.0080)	(0.0086)	(0.0109)	(0.0117)	(0.0124)	(0.0138)
H_0: p Value	0.0000	0.0157	0.7378	0.6990	0.6891	0.7322	0.1296	0.0069	0.0000
log(No. Firms/mi.^{2})	0.0544*	0.0552*	0.0483*	0.0524*	0.0435*	0.0369*	0.0443*	0.0518*	0.0554*
	(0.0132)	(0.0117)	(0.0093)	(0.0080)	(0.0086)	(0.0109)	(0.0116)	(0.0124)	(0.0137)
H_0: p Value	0.0000	0.0130	0.5339	0.8455	0.9857	0.8941	0.2471	0.0097	0.0000
log(HHI No. Estab./mi.^{2})	0.0502*	0.0516*	0.0446*	0.0488*	0.0402*	0.0334*	0.0400*	0.0487*	0.0530*
	(0.0126)	(0.0112)	(0.0086)	(0.0076)	(0.0081)	(0.0102)	(0.0110)	(0.0116)	(0.0130)
H_0: p Value	0.0000	0.0098	0.2295	0.6351	0.9271	0.9840	0.5420	0.0236	0.0000
log(HHI No. Firms/mi.^{2})	0.0536*	0.0546*	0.0476*	0.0517*	0.0428*	0.0360*	0.0434*	0.0511*	0.0550*
	(0.0132)	(0.0118)	(0.0093)	(0.0081)	(0.0086)	(0.0108)	(0.0115)	(0.0122)	(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 (rounded)	900	900	900	900	900	900	900	900	900
Clusters (rounded)	300	300	300	300	300	300	300	300	300

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 5.2, for ready–mix concrete plants in CMF years 1982, 1987, and 1992. Year–specific constants are included but not reported and standard errors (in parentheses) 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. 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.

the bias it introduces. 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:

$$\rho_{mt}^{(k)} = \beta_t^{(k)} + \beta_c^{(k)} c_{mt} + g^{(k)}(n_{mt}) + \nu_{mt}, \qquad (13)$$

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. In Online Appendix Section A.4 I offer a simple Monte Carlo exercise that replicates the basic features of this exercise and shows first, that the bias is present, and second, how the bias correction restores the desired coverage of the grouped quantile IV estimator.

Recall from Section 5.2.1 the prediction of the selection effect hypothesis: $\beta_c^{(k)}$ should be positive and converge 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 $k = 1, \ldots, 9$ are presented in Table 6 and visually in Figure 2. 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 6 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.

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 5.1, 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, it tells us that establishments in the upper quantiles of the productivity distribution, those which 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.

6 Robustness and Other Considerations

In this section, I document several robustness checks. Returns to scale could explain much of the findings of Section 4.1, that competition is correlated with productivity, and so in Section 6.1 I develop a test to rule this mechanism out. Second, in 6.2, I consider a number of alternative specifications. Finally, I also offer, in Section 6.3, evidence from my data on potential within–firm mechanisms driving the result that the correlation between productivity and competition is primarily a treatment effect. However, this evidence is incomplete; it is meant primarily to suggest interesting avenues for future research.

6.1 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 $\hat{q}_{it}$ to be the constant returns to scale predicted output:

$$\hat{q}_{it} = \alpha_{Lt} l_{it} + \alpha_{K^E t} k_{it}^E + \alpha_{K^S t} k_{it}^S + \alpha_{Mt} m_{it} + \alpha_{Et} e_{it}. \tag{14}$$

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

$$q_{it} = (1+\gamma)\hat{q}_{it} + \tilde{\omega}_{it}, \tag{15}$$

where $\tilde{\omega}_{it}$ is the true productivity term. Rearranging equation (15) and taking expectations the relationship to productivity becomes clearer:

$$\mathbb{E}[q_{it} - \hat{q}_{it}] = \mathbb{E}[\tilde{\omega}_{it}] + \gamma\hat{q}_{it}. \tag{16}$$

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 endogenous, but the same demand shifters from Section 3.3 can be used as instruments for scale.[26] Model (1) of Table 7 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 the results for the overidentification test, which strongly reject the null of exogeneity. This is consistent with the maintained hypothesis of the paper—if the demand shifters affect 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).

6.2 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 8 for each of the four measures of competition; note that each coefficient estimate reported is from a distinct regression.

[26] Note that this test of CRS still maintains the other assumptions of Section 3.4, i.e. that inputs are flexible and homogeneous, and that input prices are homogeneous, ruling out, e.g., monopsony power.

Table 7: Competition, Productivity, and Returns to Scale

	Dependent Variable: TFPQ (ω_{it})				
	(1)	(2)	(3)	(4)	(5)
Scale Coeff. ($\hat{q}_{it}$)	0.0985*	-0.0273	-0.0139	-0.0392	-0.0131
	(0.0172)	(0.0311)	(0.0287)	(0.0353)	(0.0319)
log(No. Estab./mi.^{2})		0.0556*			
		(0.0121)			
log(No. Firms/mi.^{2})			0.0539*		
			(0.0122)		
log(HHI No. Estab./mi.^{2})				0.0615*	
				(0.0137)	
log(HHI No. Firms/mi.^{2})					0.0613*
					(0.0161)
Observations (rounded)	7400	7400	7400	7400	7400
Clusters (rounded)	300	300	300	300	300
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 6.1, for ready–mix concrete plants in CMF years 1982, 1987, and 1992. 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 (in parentheses) are clustered at the CEA level. Hansen J (over identification) test statistics are reported with associated p values. For reference, * signifies $p \leq 0.05$.

Table 8: Alternative Specifications

	Dependent Variable: TFPQ (ω_{it})					
	(1)	(2)	(3)	(4)	(5)	(6)
log(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)
log(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)
log(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)
log(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 (rounded)	5000	5900	7100	7400	4200	7400
Clusters (rounded)	300	300	300	300	300	300

Notes: Here I consider alternative specification results for the effect of competition on productivity residuals, for ready–mix concrete plants in CMF years 1982, 1987, and 1992. 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 (in parentheses) are clustered at the CEA level. For reference, * signifies $p \leq 0.05$.

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 from the within–market variation. The point estimate is similar to results from other specifications, but the there is insufficient variation to obtain a statistical signifigance. 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.

6.3 Mechanisms

If the relationship between competition and productivity is governed by a mechanism that behaves as a treatment effect, rather than the selection effect generated by industry dynamics models, where should we look to understand it? In this section I revisit the instrumental variables regression design from Section 4.1 and estimate

$$y_{it} = \beta_t + \beta c_{m(i)t} + \varepsilon_{it}, \tag{17}$$

where y_{it} stands in for a host of left-hand side variables that capture firm behavior and may help to explain how firms in more competitive markets reorganize to increase output. Table 9 presents results where y_{it} stands in for the ratio of ready–mix concrete revenues to total revenus, capturing specialization, and several factor input ratios: non–productive to productive labor hours, labor to capital, and equipment capital to structural capital.

The results paint a coherent picture: in more competitive markets, ready–mix concrete firms are more specialized, as we see from Models (1)–(4). They spend relatively more hours on management ("non–productive labor"), per Models (5)–(8), which suggests that much of the relative decline in labor expenditures, see Models (9)–(12) is coming from productive labor hours. Though statistically marginal, in Models (13)–(16) it also appears that there is greater expenditure on equipment rather than structural capital. This last result may be an artifact of the first, as ready–mix concrete trucks are equipment capital.

Combined, these results suggest that the answer lies in management rather than capacity utilization. If capacity utilization were the story, then we would expect to see the labor share of variable inputs such as labor share to be positively rather than negatively correlated with competition driven by demand shocks. This is confirmed by the last eight specifications. In Models (17)–(20) and (21)–(24), I check for correlations with the energy capital ratio, as both structural and equipment capital, and find no statistically significant result. Instead, the positive correlation with the ratio of non–productive to productive labor hours supports the hypothesis that better management is improving productivity. In Online Appendix

Table 9: Alternative Dependent Variables

	Dependent Variable:							
	RMC to Total Revenue Ratio				Non- to Productive Labor Ratio			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log(No. Estab./mi.^{2})	0.0142*				0.0195*			
	(0.0035)				(0.0058)			
log(No. Firms/mi.^{2})		0.0150*				0.0222*		
		(0.0038)				(0.0063)		
log(HHI No. Estab./mi.^{2})			0.0143*				0.0187*	
			(0.0036)				(0.0057)	
log(HHI No. Firms/mi.^{2})				0.0170*				0.0251*
				(0.0044)				(0.0070)
Observations (rounded)	7400	7400	7400	7400	7400	7400	7400	7400
Clusters (rounded)	300	300	300	300	300	300	300	300
First-Stage F	344.5	196.3	171.4	48.7	344.3	196.2	171.1	48.7
p Value	0	0	0	0	0	0	0	0
Hansen J Statistic	2.048	2.015	2.423	2.377	3.454	2.878	4.207	3.266
p Value	0.3591	0.3651	0.2977	0.3046	0.1778	0.2371	0.1220	0.1953
	Dependent Variable:							
	Labor to Capital Ratio				Equipment to Structural Capital Ratio			
	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
log(No. Estab./mi.^{2})	-0.0075*				0.5173*			
	(0.0019)				(0.2452)			
log(No. Firms/mi.^{2})		-0.0081*				0.4709		
		(0.0020)				(0.2672)		
log(HHI No. Estab./mi.^{2})			-0.0076*				0.5463*	
			(0.0020)				(0.2269)	
log(HHI No. Firms/mi.^{2})				-0.0095*				0.5688
				(0.0024)				(0.2965)
Observations (rounded)	7400	7400	7400	7400	7400	7400	7400	7400
Clusters (rounded)	300	300	300	300	300	300	300	300
First-Stage F	344.5	196.3	171.4	48.7	344.5	196.3	171.4	48.7
p Value	0	0	0	0	0	0	0	0
Hansen J Statistic	1.548	1.330	1.721	1.349	1.364	2.014	0.697	1.526
p Value	0.4610	0.5140	0.4229	0.5092	0.5056	0.3652	0.7057	0.4661
	Dependent Variable:							
	Energy to Structural Capital Ratio				Energy to Equipment Capital Ratio			
	(17)	(18)	(19)	(20)	(21)	(22)	(23)	(24)
log(No. Estab./mi.^{2})	0.0064				-0.0475			
	(0.0056)				(0.0425)			
log(No. Firms/mi.^{2})		0.0058				-0.0457		
		(0.0059)				(0.0500)		
log(HHI No. Estab./mi.^{2})			0.0081				-0.0484	
			(0.0059)				(0.0415)	
log(HHI No. Firms/mi.^{2})				0.0085				-0.0464
				(0.0072)				(0.0602)
Observations (rounded)	7400	7400	7400	7400	7400	7400	7400	7400
Clusters (rounded)	300	300	300	300	300	300	300	300
First-Stage F	344.5	196.3	171.4	48.70	344.5	196.3	171.4	48.70
p Value	0	0	0	0	0	0	0	0
Hansen J Statistic	4.057	4.162	3.668	3.919	3.612	4.045	3.515	4.287
p Value	0.1315	0.1248	0.1598	0.1409	0.1643	0.1323	0.1725	0.1172

Notes: Here I consider alternative specifications of the IV regressions with different left-hand side variables, as detailed in Section 6.3, for ready–mix concrete plants in CMF years 1982, 1987, and 1992. Year–specific constants are included but not reported and standard errors (in parentheses) 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 $p \leq 0.05$.

Section A.2.5, I include controls for establishment age and CEA area to consider alternative hypotheses: vintage capital and a mechanical effect of competition on productivity through geographic differentiation. The results, are inconsistent with either story, and so I do not believe they are first–order, but the data to not permit me to decisively rule such effects out. Additionally, in Online Appendix Section A.3, I introduce data from the Management and Organizational Practices Survey, but find no suggestive evidence on mechanisms from this.

This leaves the question open for future work, but not without guidance. Based on my findings in Section 6.1 above, the mechanism does not induce returns to scale, which is inconsistent with traditional stories based on endogenous sunk costs: advertisement and investment in reductions in variable cost (Sutton, 1991). Capacity utilization too, as I showed above, is inconsistent with the data. Finally, there is suggestive evidence that the effect is correlated with greater investment in non–productive worker hours, or management. One hypothesis that would be consistent with the above is that having competitors, and observing those competitors, generates information that allows managers to better disentangle signals and monitor workers. Therefore less competitive markets create room for managerial slack (Nalebuff and Stiglitz, 1983). What form does this managerial slack take? It may be that this takes the form of direct wastage: in the ready–mix concrete business, hardened concrete on the inside of a drum, which is costly to extract. Alternatively, and consistent with the relative reduction of productive labor hours, it may be that in more competitive markets, the reduction of managerial slack means cutting other slack, i.e. firing unproductive (but not non–productive) workers. Based on the results here, this organizational focus appears to be the right direction to push future research on the relationship between competition and productivity.

7 Discussion

The objective of this paper has been to test two competing hypotheses that could explain the positive correlation between productivity and competition. Motivating both of these hypotheses—the treatment effect and the selection effect of competition—are different theoretical models which raise different questions for future applied work. I have offered three pieces of evidence on the point: a direct test of the reallocation hypothesis in ready–mix concrete, a semi–parametric selection correction, and a grouped IV quantile approach. All of my evidence points to the conclusion that competition has a treatment effect on measured

Table 10: Decomposed Effect of a Standard Deviation Increase in Competition

Model	Output	Attributable to Covariance	Treatment	Selection
log(No. Estab./mi.^{2})	3.98%	0.25%	3.39%	0.34%
log(No. Firms/mi.^{2})	4.05%	0.23%	3.47%	0.35%
log(HHI No. Estab./mi.^{2})	4.13%	0.30%	3.80%	0.04%
log(HHI No. Firms/mi.^{2})	4.36%	0.26%	3.82%	0.29%

Notes: Here I present a decomposition of the effect of a one standard deviation increase in the competition index on physical output into its constituent parts: the covariance of output and marketshare (or "reallocation"), the treatment effect of competition on productivity, and the selection effect driven by greater exit of unproductive establishments. Estimates are based on a one standard deviation increase in 1987, per Table 1 as well as coefficient estimates from Tables 4 and 5.

productivity.

Combining the decomposition in Section 4.2 with the structural estimates in Section 5.1, I can model the effect of a one standard deviation change in competitiveness. In Table 10 I show how, across my different specifications, this increase affects productivity through different channels. For all specifications, output increases by approximately 4%. Less than a tenth of that comes through changes in the covariance between productivity and competition, i.e. the reallocation channel. And since differential selection is the extensive margin of reallocation, it too explains less than a tenth of the increase in output. The bulk of the change is coming through the treatment effect channel.

Understanding the relationship between competition and productivity is particularly important at a time when many are reconsidering some of the fundamental tenets of antitrust economics. Of late, there has been great public interest in understanding the apparent rise of market power across many industries (Eekhout et al., 2020) and the implications of this apparent rise for productivity in America (Shambaugh et al., 2018). Insofar as productive efficiencies come about through vigorous competition, this is a potentially powerful motivation for strong competition policy, rather than a counterweight as in Williamson (1968).

To summarize, the inventory of economic models offers two broad channels by which competitiveness might affect productivity: reallocation between firms and the exit of unproductive firms on the one hand, and increases in productivity stemming from changes within the firm on the other. The evidence presented here suggests that the latter is the productive direction for future work.

References

Ackerberg, D., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.

Adelman, M. A. (1969). Comment on the "h" concentration measure as a numbers-equivalent. *Review of Economics and Statistics*, 51(1):99–101.

Asker, J., Collard-Wexler, A., and De Loecker, J. (2019). (mis)allocation, market power, and global oil extraction. *American Economic Review*, 109(4):1568–1615.

Asplund, M. and Nocke, V. (2006). Firm turnover in imerfectly competitive markets. *Review of Economic Studies*, 73(2):295–327.

Atalay, E. (2014). Material prices and productivity. *Journal of the European Economic Association*, 12(3):575–611.

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., Draca, M., and Van Reenen, J. (2016). Trade induced technical change? the impact of chinese imports on innovation, it and productivity. *Review of Economic Studies*, 83(1):87–117.

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. (2016). Iv quantile regression for group-level treatments, with an application to the distributional effects of trade. *Econometrica*, 84(2):809–833.

Collard-Wexler, A. (2011). Productivity dispersion and plant selection in hte 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. (2015). Reallocation and technology: Evidence from the us steel industry. *American Economic Review*, 105(1):131–171.

De Loecker, J. and Goldberg, P. K. (2014). Firm performnace in a global market. *Annual Review of Economics*, 6:201–227.

Dertouzos, M. L., Lester, R., and Solow, R. M. (1989). *Made in America: Regaining the Productive Edge*. MIT Press, Cambridge, MA.

Doraszelski, U. and Jaumandreu, J. (2013). R&d and productivity: Estimating endogenous productivity. *Review of Economic Studies*, 80(4):1338–1383.

Doraszelski, U. and Satterthwaite, M. (2010). Computable markov-perfect industry dynamics. *The RAND Journal of Economics*, 41(2):215–243.

Dunne, T., Klimek, S., and Schmitz, J. (2010). Competition and productivity: Evidence from the post wwii us cement industry. Working Paper.

Dunne, T., Roberts, M. J., and Samuelson, L. (1988). Patterns of firm entry and exit in us manufacturing industries. *RAND Journal of Economics*, 19(4):495–515.

Eekhout, J., De Loecker, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2):561–644.

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 Krizan, C. J. (2001). *Aggregate Productivity Growth: Lessons from Microeconomic Evidence*, volume New Developments in Productivity Analysis, pages 303–372. University of Chicago Press.

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.

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. (2015). Dynamic olley-pakes productivity decomposition with entry and exit. *RAND Journal of Economics*, 46(2):362–375.

Nalebuff, B. and Stiglitz, J. (1983). Information, competition, and markets. *American Economic Review*, 73(2):278–283.

Neyman, J. and Scott, E. L. (1948). Consistent estimates based on partially consistent observations. *Econometrica*, 16(1):1–32.

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.

Shambaugh, J., Nunn, R., Breitwieser, A., and Liu, P. (2018). The state of compeittion and dynamism: Facts about concentration, start-ups, and related policies. Brookings Insittute: Hamilton Project Working Paper.

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 xistence 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.

Williamson, O. E. (1968). Economies as an antitrust defense: The welfare tradeoffs. *American Economic Review*, 58(1):18–36.