Firm Decentralization and Medium-Run Growth*

Daron Acemoglu†
John Van Reenen‡
Sean Wang§

February 2026

Abstract

The economic consequences of changing the distribution of power within firms is a fundamental question, yet little is known of the longer-run growth effects of giving more autonomy to local managers. We match direct measures on decentralization of decision rights within firms to output and inputs in two complementary datasets: U.S. Census Bureau and international company accounts across 18 countries. In both datasets we show that decentralization from headquarters to plant managers is robustly correlated with growth over 15 years. We introduce a new IV strategy exploiting the decentralization tendencies of previous parent companies (a "donor firm" method) to show that the positive relationship between decentralization and growth appears causal. We find similar results using an alternative IV based on industry decentralization trends in culturally similar countries. We highlight two mechanisms for the positive effect of decentralization on growth. First, our results are strongest in older firms still controlled by the founder, suggesting some founders may inefficiently hoard power. Second, more decentralized plants perform better when faced by exogenous increases in business opportunities (market size shocks), suggesting they have greater agility. Overall, our results suggest a tendency in firms towards excess centralization.

JEL No. O31, O32, O33, F23
Keywords: decentralization, growth, founder, agility

*Corresponding author: sean.y.wang@census.gov. We would like to thank Nick Bloom and Bob Gibbons for many stimulating discussions on these issues. This research was funded by the ESRC/UKRI through POID, IBM Global Universities Programs, Schmidt Futures, Sloan Foundation, and the Smith Richardson Foundation. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1745302 (Wang). Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure avoidance protection (Project 7512395: CBDRB-FY23-CES008-01, CBDRB-FY23-CES022-002).

†MIT and NBER
‡LSE and NBER
§U.S. Census Bureau

1 Introduction

Issues of power and authority are an intrinsic part of social, political, and economic life, and influential scholars from across social science disciplines have contributed to our understanding of the origins, characteristics, and potential consequences of these forces.[1] The management "guru's guru", Peter Drucker, is perhaps the most famous exponent of the benefits of decentralizing managerial power. The economy "has to be decentralized ... its organizations must be able to make fast decisions, be based on closeness to performance, closeness to the market, closeness to technology... which must be seen and utilized as opportunities for innovation"(Drucker, 1993).[2] Drucker's views were in sharp contrast to many traditional approaches which emphasized the importance of command and control. For example, Henri Fayol (1841-1925), is generally recognized as the founder of management science alongside his contemporary, Frederick Taylor. He saw the military as the ideal type of efficient organization, and his fourteen 'principles of management' included centralization, discipline, unity of command and direction and the subordination of individual interests (Fayol, 1933).[3]

Modern organizational economics points to the allocation of authority as a key part of the firm design (from Coase, 1937, Simon, 1951, through to Gibbons et al. (2012) and Bolton and Dewatripont, 2013). Despite this interest and a multitude of business case studies, formal empirical evidence of the consequences of delegating power *within* firms is very rare, especially over longer-run horizons. This is partly due to two major empirical challenges. First, the analysis requires both direct measures of decentralization and firm outcomes over long periods of time to abstract away from the influence of short-term shocks or business cycles. Second, since decentralization is a design choice, we need exogenous variation to isolate its causal effects. This is challenging as organizational features like decentralization tend to be persistent over time (e.g. Ichniowski et al. (1997); Gibbons and Roberts (2012); Gibbons and Henderson (2012)). The few recent studies with causal designs to study the impacts of decentralization have been in the context of the public sectors of developing countries, which have been more amenable to conducting experiments.[4]

We tackle these twin challenges of data and design by assembling two sets of longitudinal microdata on both decentralization and business outcomes. One uses administrative and survey data from the U.S. Census Bureau, and the other uses accounting and World Management Survey (WMS) data from 18 countries. We show across both datasets that establishments with a higher baseline level of decentralization from their headquarters subsequently experience higher levels growth (in sales, value-added, and employment). Since sales and value-added growth generally outpace employment growth, measures of labor productivity also tend to rise.

We introduce two complementary IV designs that exploit plausibly exogenous variation in the extent to which plants are given control rights from their headquarters. One design exploits the rich longitudinal firm ownership histories from the Census Bureau, and another exploits the cross-country comparability of the WMS. Across both settings and research designs, we find consistent

[1] For example, see the literatures following Marx (1867); Russell (1938); Parsons (1963); Dahl (1957); Weber (1978).

[2] See "Our Central Problem: Not Enough Decentralizing", The Drucker Institute https://drucker.institute/thedx/our-central-problem/ . Drucker himself attributed these insights to his favorite guru, Mary Parker Follett (1868-1933) "The ramifications of modern industry are too widespread, its organisation too complex, its problems too intricate for industry to be managed by commands from the top alone." (Robinson, 2005)

[3] More modern supporters of Fayol would include Jacques (1990), for example. See "In Praise of Hierarchy" https://hbr.org/1990/01/in-praise-of-hierarchy .

[4] We discuss Duflo et al. (2018); Bandiera et al. (2020); Balán et al. (2022) in the next subsection.

evidence that our results reflect the positive causal effects of decentralization on growth, as opposed to the differential selection of high-growth plants into adopting decentralization.

For U.S. data on decentralization and plant outcomes, we combine the Census Bureau's Management and Organizational Practices Survey (MOPS) with establishment-level data on shipments, value added employment and survival from the Census of Manufactures (CMF) and the Longitudinal Business Database (LBD). The MOPS data was collected as a supplement to the 2010 and 2015 Annual Survey of Manufactures (ASM). The survey asks respondents to directly report the degree of autonomy plant managers have from the central headquarters in six separate areas: capital investments, hiring of full-time employees, product introduction, marketing, pay, and product pricing. Our primary measure of decentralization is an index calculated from the plant manager's responses over these six decisions, which can be interpreted as the share of business decisions that a plant can make autonomously. We use the CMF as our primary source of plant-level data on the total value of shipments, value-added, and employment between 2002 and 2017. We focus on these as measures of plant size, and we additionally consider sales per worker and value-added per worker as measures of the plant's labor productivity. We augment the CMF with administrative data covering the near-universe of U.S. private, non-farm employers from the LBD as our primary source of data on firm characteristics and on establishment entry and exit. The LBD combines administrative payroll tax records with annual surveys of large firms to track establishment births, exits, job creation/destruction, and firm ownership structures from 1976 onwards (Chow et al., 2021).

To generate exogenous variation in plant decentralization, we introduce a "donor firm" methodology that exploits differences in the historical context in which an acquired plant operated prior to joining its current owner. To fix ideas, consider the diverging fates of two acquisitions by Microsoft. In 2014, it acquired the establishments that constitute the mobile phone division of the Finnish conglomerate Nokia. Despite its history of dominance, the division has been struggling to adapt to the smartphone revolution due in part to a "Soviet-style bureaucracy" in the company's design process (O'Brien, 2010). The acquisition had several elements that promised success — Nokia's CEO at the time was a former Microsoft executive, and the two organizations had a history of collaboration on their mobile products. Microsoft intended to use the acquisition to jumpstart the competitiveness of its mobile platform by tightly aligning hardware and software (as Apple has done), and it integrated several high-level Nokia executives into its organization, including putting the current Nokia CEO in charge of all Microsoft devices (Bort, 2013). However, the division's sales stagnated as the company struggled to align the mobile division's priorities with the rest of the firm, and it ultimately sold the division back to a startup run by former Nokia executives in 2016 for $350 million, a tiny fraction of the $9.4 billion it paid just two years earlier (Greene and Verbergt, 2016). In the same year, Microsoft made an even larger, $26 billion acquisition of the social networking startup LinkedIn without a substantive history of operating in the area. Contrary to the expectations of market skeptics, the firm recognized LinkedIn's "distinct brand, culture, and independence" and decided to retain its leadership team and operate it as a separate and autonomous entity that reports directly to the Microsoft CEO. In the years following the acquisition, LinkedIn has experienced steady growth and continued to "by and large ... operate as a separate entity from the rest of the company," even deferring planned integration with other Microsoft products to ensure they're "maintaining the trust of [their] membership" (Erlichman, 2019; Novet, 2019).

We use the consistent longitudinal coverage of both establishment and firm characteristics in

our Census data to systematically assess the validity of these examples in an instrumental variables (IV) framework. We assume (and provide supporting evidence) that (i) multi-establishment firms have propensities to decentralize across all their units (this "culture" may be the solution to long-standing firm-specific factors); (ii) when an establishment is bought by another firm, it tends to keep a portion of its previous control rights; (iii) conditional on the new firm's characteristics and the plant's economic environment, the original decentralization status of the "donor firm" is otherwise uncorrelated with any subsequent shocks to the plant (except through the plant's persistent degree of decentralization). We implement our IV strategy by identifying a set of plants that were acquired between 1982 and 2002 (up to twenty years before we start measuring plant outcomes) and using the characteristics of their donor firms in 1982 that predict firm decentralization (such as size) as an instrument for the eventual decentralization propensity of the acquired plants. This strategy is related to the long-run institutional persistence literature, which uses long-dated characteristics of a geographical area (at a critical historical conjunction) as plausibly exogenous determinants of key institutional features (e.g., Acemoglu et al. (2001, 2005); Dell (2010)). Once these institutions bed down, they often persist even when the initial characteristics change or are no longer relevant to outcomes.

Our results from the U.S. data show that plant decentralization is robustly correlated with 15-year-long growth rates and that the effects seem to be driven by the causal effect of persistent decentralization on labor productivity. For example, in our baseline OLS regressions, we find that a one standard deviation increase in decentralization (equivalent to getting full autonomy for an additional 1.2 out of 6 business decisions) is associated with a 0.4 percentage point (p.p.) increase in annual sales and value-added growth, and a 0.3 p.p. increase in jobs growth.

Our baseline IV results are larger for outputs, but smaller for employment, suggesting that the channel may operate through labor productivity growth. We confirm that the effects of decentralization on the sales or value-added per worker at surviving plants are especially large in our IV regressions. These effects also remain similar after adding an array of detailed controls, including structured management (as in Bloom et al. (2012, 2019)) and the detailed characteristics of the parent firm. Finally, as a placebo exercise, we find that our instrument has no effect on the growth rate of single-unit firms — a setting where decentralization from headquarters to plant managers would not be operating, but other potential concerns with the exogeneity or excludability of our "donor" firm strategy would have remained relevant.

To assess the degree to which our results generalize to other countries and settings, we also create another, internationally comparable dataset based on the World Management Survey. The WMS is a survey of plant managers that has asked questions on decentralization from the 2006 wave onwards, nearly identical to MOPS. We match the WMS to firm-level accounting data from BVD ORBIS and focus on the firm-level growth in sales and employment in 6-year increments as our principal outcomes of interest.

We exploit the international dimension of our linked WMS data to construct a complementary IV strategy to our Census data "donor" firm design. Specifically, we follow the literature that considers culture as an important driver of decentralization decisions (e.g. Hofstede (1984)), and we find that firms in English-speaking countries (Australia, Canada, New Zealand, and the U.K.) are more likely than firms in other countries to follow the decentralization propensities of comparable firms in the U.S., despite the fact that production technologies may be similar across countries within

any given industry. We use the differential willingness of English-speaking countries to align with the decentralization practices of their American peers (firms in the same industry) as a source of non-technological variation in decentralization, and we consider the corresponding difference in firm outcomes across the industry-country pairs.

Despite the differences in the underlying source of variation, our results from the WMS are consistent with the positive causal effect of decentralization on firm growth. Our baseline OLS estimates indicate a one standard deviation increase in decentralization is associated with a 0.6 annual p.p increase in sales growth and a 0.4 p.p increase in employment growth. When we focus on the U.S.-induced industry-level correlation in decentralization in other English-speaking countries, we again find corroborating evidence on the causal effects of decentralization.

Finally, we highlight two related questions. Why, if decentralization is beneficial are so many firms centralized? And what are the potential mechanisms behind these positive, long-run effects of decentralization? On the first question, we find that the effects of decentralization are especially strong in firms that are over a decade old, but are still controlled by the founder. As a result, the negative effects of *centralization* may reflect in part an inefficient organization structure that stems from the desire to "hoard power" by firm owners. On the second issue of mechanisms, we find that decentralized plants are more agile — growing faster in industries that experienced sustained exogenous growth in product demand over the 15-year time period. This greater agility suggests that local information may be a persistent advantage that plant managers have over company headquarters in both good and bad times.

So, taken together, the data suggests that in the long-run decentralization on average does foster growth, broadly in line with the Drucker hypothesis.

After a brief literature review, the rest of the paper is organized as follows. Section 2 describes our data, Section 3 the empirical strategy, Section 4 presents the main MOPS results, Section 5 examines WMS, Section 6 mechanisms and Section 7 concludes. Online appendices have more data details (A) and further results (C).

1.1 Related Literature

Our paper is most closely related to the small but growing literature on the empirical effects of decentralization on organizational performance. Kala (2019) uses a policy change to study the effects of decentralizing control from the Indian federal government to managers of state-owned enterprises. Decentralization depends on (partly nonlinear) rules relating to the prior profitability and net assets position of these enterprises. She finds that decentralization increases firm value-added up to ten years later. There were also distributional impacts increasing markups but reducing worker amenities (such as housing). Bandiera et al. (2020) uses a randomized control trial in Pakistan to examine the impact of increasing autonomy to government procurement officers (600 officers in 26 districts of the Punjab). Autonomy enables the government to secure lower procurement prices without lowering product quality. Balán et al. (2022) also have a randomized control trial, this time decentralizing the collection of property taxes from the central state agency to local village chiefs in the Congo. They find that decentralization raised tax collection (but also bribes) as the chiefs had better local information on which households were more likely to pay their taxes. Our paper is consistent with

these three papers in finding a positive causal effect of decentralization on performance.[5] But it differs as we focus on the private sector in mainly developed countries, rather than decentralizing control rights from government entities in emerging economies. To our knowledge, we are the first paper systematically documenting the longer run growth effects of decentralization within *private* firms. This is an interesting context, as there are many reasons why public sector entities could have inefficient organizational architectures, especially in countries with weak state capacity.[6]

Our paper is also related to the literature that focuses specifically on the performance of private firms under decentralization. Aghion et al. (2021) examine the short-run growth responses of decentralized firms to the negative economic shock of the Global Financial Crisis. Their data does not allow them to look at longer run like ours, which is why they focus on business cycle variation. Bradley et al. (2011) report a positive relationship between firm independence (which they interpret as a proxy for greater autonomy in resource allocation decisions) and firm survival during downturns using Swedish data. Neither paper has designs to capture exogenous differences in decentralization.

We also contribute to the existing literature on the determinants of decentralization (for example, see the survey in Aghion et al. (2014)). On the theory side, our paper relates to the literature on decentralization within the firm (see Gibbons and Roberts (2012) or Garicano and Rayo (2016) for recent surveys) and incomplete contracts (see Gibbons and Roberts (2012)). In particular, Lange (1936); Hayek (1945); Hart and Moore (2005); Acemoglu et al. (2007) analyzes the optimal allocation of authority and how it depends on the trade-off between the nature of the agent's local information and the potential for conflicts between the principal's and agent's incentives. Empirically, Rajan and Wulf (2006); Blundell et al. (2016) document a movement towards decentralized firms, and Caroli and Van Reenen (2001); Bresnahan et al. (2002); Guadalupe and Wulf (2010); Katayama et al. (2018); McElheran (2014); Dessein et al. (2019) highlight factors that influence the extent of decentralization within firms. Our "donor firm" IV strategy additionally emphasizes the role of path dependency as a determinant of decentralization, which supports the literature on organizational persistence (see Feldman et al. (2019) for a survey), while our international IV highlights the role of culture in addition to technology (e.g., Hofstede (1984)).

2 Data

We detail all datasets in Appendix A, but summarize some main features here.

2.1 U.S. Establishment-level Data

Management and Organizational Practices Surveys (MOPS) We draw our direct measures of decentralization and management practices from the Management and Organizational Practices Surveys (MOPS), the world's largest survey on organizational practices. It was conducted by the U.S. Census Bureau as a supplement to the Annual Survey of Manufactures (ASM) and responses were legally mandated, ensuring a very high sample response rate. We use the first and second waves

[5] It is also consistent with the message of Duflo et al. (2018). The authors conducted an experiment in India, doubling environmental inspections in the treatment group, but made them random, so removing discretion from local inspectors. There was little effect which they interpret as being due to the removal of autonomy from the inspectors who were using their local information to more effectively target plants.

[6] More broadly, our paper is also related to the large literature on autonomy and school effectiveness. See, e.g., Galiani et al. (2008); Clark (2009); Hanushek et al. (2013); Abdulkadiroğlu et al. (2016).

referring to reference years 2010 and 2015 (with recall questions back to 2005). The design was based on the World Management Survey (WMS) that we discuss below and was generally answered by a plant manager (for more background on the MOPS, see Buffington et al. (2016) and Appendix A).

One block of the survey contained questions on decentralization regarding the degree of autonomy the plant manager had from the central headquarters. The dimensions of decentralization related to six areas: (a) capital investments, (b) hiring of full-time employees, (c) product introduction, (d) marketing, (e) pay, and (f) product pricing. For example, the capital question reads: "In 2005 and 2010, what was the dollar amount that could be used to purchase a fixed/capital asset for this establishment without prior authorization from headquarters?" There were five choices for the question on autonomy in capital investments, starting with "Under $1,000" (coded as 1) up until "$1 million or more" (coded as 5). The other questions had three possible responses. For example, the labor hiring question was: "In 2005 and 2010, where were decisions on hiring permanent full-time employees made? "Only at this establishment" (coded as 3); "Both at this establishment and at headquarters" (coded as 2); "only at headquarters" (coded as 1).

Our baseline approach is to follow the extant literature and normalize the scale of all six decentralization questions to be between 0 and 1, and calculate the plant's decentralization index as the simple arithmetic average of the plant manager's autonomy scores over these six decisions. This approach gives the decentralization index an intuitive interpretation as approximately the share of the six business decisions where the plant has maximal autonomy. We examined alternative ways to build the decentralization index and found our results to be robust across different approaches.[7]

Each survey asked the respondent to report the plant's degree of decentralization separately for the reference year and five years earlier. As a result, we observe a plant's decentralization status up to four times - in 2010 and 2015 (from the contemporaneous questions in the 2010 and 2015 surveys), as well as in 2005 and 2010 (from the 5-year recall questions in the 2010 and 2015 surveys). The 2010 recall responses are useful as an additional source of data for 2010, and the 2005 recall responses are especially valuable because they measure decentralization at a time close to the beginning of our plant outcomes period that otherwise would not be covered by any existing organizational surveys. Bloom et al. (2019) has shown that these recall responses tend to be high quality whenever the tenure of the manager responding to the survey started at least one year before the recall period, so we correspondingly only use recall responses from managers with at least seven years of tenure at the establishment.

Other Census Data Given that the sampled plants for the MOPS are drawn from the universe of manufacturing plants, we focus on measuring long-run outcomes at the plant level using the 2002 and 2017 Censuses of Manufactures (CMF). The Economic Censuses are conducted every five years, and 2017 is the most recent year of data available. Defining the start of the sample period in 2002 ensures that we are able to maximize the time horizon for measuring plant outcomes while still observing a plant's decentralization status close to the start of our sample period. The CMF provides plant-level data on its total value of shipments (which we use as a proxy for sales), value-added, and employment. We focus on these three measures as our focal measures of plant size, and we additionally consider shipments per worker and value-added per worker as measures of the plant's labor productivity.

[7]For example, our empirical results are similar if we z-score and aggregate over each of the six responses, or if we discretize the overall decentralization measure to a binary above/below median indicator.

Although the CMF provides important plant-level detail on sales and value-added for the near-universe of manufacturing plants, it is missing detail for non-manufacturing plants and during non-Census years. Consequently, we use administrative data covering the universe of US employers from the Longitudinal Business Database (LBD) as our primary source of data on firm characteristics, ownership, and establishment entry and exit. The LBD combines administrative payroll tax records with annual surveys of large firms to track establishment births, exits, job creation/destruction, and firm ownership structures from 1976 to the present(Chow et al., 2021). We exploit this longitudinal coverage of both establishment and firm characteristics to identify mergers or acquisitions and to construct instruments for the historic size of "donor" firms. We also use the LBD as our primary measure of establishment exit.

For constructing our measure of export-based shocks to market size we use trade data. Establishment revenues by product class are available from the Longitudinal Firm Trade Transactions (LFTTD) data which contain the universe of export shipments at the firm-product level. For the growth in a foreign product by country export market we use the UN COMTRADE data which has exports and imports by year at the HS six-digit by country level.

2.2 International micro-data from WMS and BVD

The MOPS data is the largest single-country database containing direct measures on decentralization that currently exists. But being from a single country raises questions about the generalizability of the results to other nations. To investigate this issue we use the World Management Survey (WMS) which contains information on decentralization and other features of the firm. In short, the WMS is a telephone-based survey of plant managers that has been ongoing since 2004. From 2006 onward, respondents were asked questions on decentralization. The wording of the MOPs questions was based on these WMS questions, so they are essentially identical. The two key differences are that first, the WMS only has the first four of the six MOPS questions (relating to capital, labor, product introduction, and marketing). Second, the WMS method was more open as the interviewers were able to probe the respondents for practical examples to more accurately score the managers.

We used three waves of WMS in 2006, 2009, and 2013 across 18 WMS countries. WMS asked many other questions regarding firm, worker, and manager characteristics (e.g. management practices, industry, education, ownership, etc.). To obtain outcome data, we match the WMS into firm-level accounting data from BVD ORBIS in order to obtain sales, employment, age, etc. Using this outcome data we estimate empirical models analogous to the ones in MOPS.

2.3 Descriptive Statistics

Some descriptive statistics on MOPS are contained in Table 1 for the baseline OLS sample of about 7,300 establishments. The mean of the plant decentralization index is 0.371 with a standard deviation of 0.195. It can be seen that the firms are relatively large compared to the whole LBD, with an average plant employment of over 100 and an average firm employment of over 1800. The typical plant accounts for just under a quarter (23%) of the shipments in the firm. The middle part of the table shows changes between 2002 and 2017. Unsurprisingly since this is the manufacturing sector, the average plant was shedding about 2.7% of its workforce per year. Nevertheless, output has not fallen so quickly, with only a 0.2% annual fall in shipments, implying a substantial increase in labor

productivity.

In Figure 1 we show bivariate binscatters of annualized plant growth over 15 years (2002-2017) and decentralization. There is a strong upwards gradient for all three measures - sales, value added and employment. The slope of the gradient appears similar for the output measures, but more shallow for the input jobs measure, suggesting a positive association of labor productivity with decentralization.

We show the relationship between (i) firm sales and employment growth (from BVD) and WMS decentralization in Figure 2 and (ii) plants employment growth and decentralization (both from WMS) in Figure C1. As with Figure 1 from MOPS we see a strong upwards sloping relationship in the raw data. We now move to some more formal econometric investigation of these intriguing statistical relationships.

3 Empirical Strategy

3.1 The relationship between decentralization and establishment growth

We start by estimating the relationship between establishment growth and decentralization using the following specification:

$$\Delta Y_{i,2002-17} = \mu \text{DEC}_{i,2005} + \gamma_n + \Theta X_i + \Delta \epsilon_i, \tag{1}$$

where the dependent variable $\Delta Y_{i,2002-17}$ is the annualized growth in plant i's outcome Y_i between 2002 and 2017, and the explanatory variable of interest is $\text{DEC}_{i,2005}$, the plant's i's decentralization index in 2005 (as measured by MOPS). For outcomes that measure plant size (i.e., total shipments, value-added, and employment), we use the Davis et al. (1998) transformation (henceforth, "DHS") that measures relative changes both at the intensive and extensive margin.[8] For outcomes that measure labor productivity (i.e., shipments per worker or value-added per worker), we restrict our sample to plants that had positive employment in 2017, and we consider the log-difference in outcomes between 2002 and 2017. Since outcomes always enter as a long difference between 2002 and 2017, we implicitly net out any persistent differences in each outcome at the plant level. In addition, we allow for arbitrary differences in industry growth rates over our sample period by including four-digit NAICS industry fixed effects (γ_n). Depending on the specification, we also allow for additional control variables, X_i, such as establishment baseline log(employment) and log(payroll) (measured in 2002), establishment structured management practices (as measured in 2005), and other characteristics (X_i) that we discuss below. We cluster all standard errors, (ϵ_i) at the firm level to allow for arbitrary correlations in outcomes between establishments in the same firm.

An obvious concern in interpreting μ as the effect of decentralization is that there could be omitted factors correlated with growth and DEC. For example, Caliendo et al. (2020) find that firms add organization layers in response to positive demand shocks. To mitigate this concern, we have used the plant's lagged decentralization index (from 2005, the first year decentralization data are available), estimated in long-differences and allowed for industry trends. We assess the degree to

[8] Specifically, we define $\Delta Y_{i,2002-17} = \frac{Y_{i,2017}-Y_{i,2002}}{(Y_{i,2017}+Y_{i,2002})/2} \times \frac{100}{2017-2002}$, and we assign exiting plants to $Y_{i,2017} = 0$. Plant-level value-added is negative in a couple of observations, so we censor these at zero.

which this concern may be driving our OLS estimates by systematically allowing for plant growth rates to depend flexibly on other lagged observables. The most general specifications include plant baseline size (log employment, l, and log payroll, w, in 2002), plant practices (structured management scores from 2005), firm characteristics in 2002 (log employment, log payroll, log number of plants in the firm and log age) and the plant's share of firm employment in 2002.

3.2 Exogenous variation in decentralization

Although adding plant, firm, and industry controls can help rule out some potential confounding mechanisms, decentralization may still be correlated with unobserved plant-level variables that affect growth. Consequently, we describe how we exploit differences between plants' historic owners as a source of variation in a plant's current propensity to be decentralized. As noted in the previous section, the Census Bureau maintains longitudinally consistent identifiers at both the plant-level and the firm-level in the LBD from 1976 to the present day. By following a plant's firm-level identifiers over time, we can identify establishments that have changed owners through a M&A event as well as the characteristics of their previous owners, who we describe as "donor firms" (for more details on how we identify mergers, see Appendix A.2). Our key assumption is that a plant's degree of autonomy is likely to be persistent over time, whereas direct economic relationships between the plant and the donor firm cease following a change in ownership. There is a wealth of evidence from case studies and also a smaller econometric literature suggesting strong organizational persistence (e.g. Gibbons and Henderson (2012)).

To make this approach concrete, define $\bar{D}_j$ as a firm's *permanent* propensity to decentralize control to its plants. Then defining $J(i,t)$ as the firm that owns plant i in period t, we assume:

$$E\left[\Delta\epsilon_{i,2002-17}|\bar{D}_{J(i,1982)}, l_{i,1982}, w_{i,1982}, \gamma_n\right] = 0 \tag{2}$$

In other words, for plants that have changed ownership between 1982 and 2002, the decentralization propensity of the plant's donor firm in 1982 is uncorrelated with the unobserved factors that affect the plant's growth rate between 2002–2017. It is plausible that the characteristics of a parent firm two decades after spinning off an establishment do not have a direct effect on the growth prospects of this establishment (conditional on plant controls) except through the legacy of the organizational form that the firm "donated" to it's previously owned plant.

If we were able to observe each donor firm's decentralization propensity, we would be able to test for the causal effect of decentralization by using the donor firm's decentralization propensity as an instrument. However, the set of donor firms where we directly observe the decentralization of at least one establishment is small, both because donor firms from 1982 are unlikely to survive at least through to 2010 (by definition they were the subject of at least a partial acquisition) and because MOPS itself is a partial sampling of plants in the manufacturing sector. So instead we use observable correlates of the firms' propensity to decentralize.[9]. Our primary instrument is the firm's total log(employment), but we show results using a variety of other 1982 observables such as the log(firm payroll), log(number of plants in the firm), etc.[10]

[9] i.e.

$$\bar{D}_{J(i,1982)} = f[Z_{j',1982}, \Lambda(l_{i,1982}, w_{i,1982}), \psi_n]$$

[10] We also considered machine learning approaches to build a predictor of firm decentralization trained on MOPS

Our final IV estimation equations are then:

$$\text{DEC}_i = \pi Z_{j',1982} + \Lambda\left(l_{i,1982}, w_{i,1982}\right) + \psi_n + \nu_i \tag{3}$$

$$\Delta Y_{i,2002-17} = \beta \text{DEC}_i + \Theta\left(l_{i,1982}, w_{i,1982}\right) + \gamma_n + \Delta\epsilon_{i,2002-17} \tag{4}$$

where $Z_{j',1982}$ is the log employment of donor firm j' in 1982 and DEC_i is the average decentralization index of plant i across all waves of the MOPS. In our preferred specification, we separately control for the establishment's employment and payroll in 1982 to ensure that we are isolating the variation from each donor firm's other plants for identification, and in some specifications, we also additionally include the plant's structured management score and the plant's recipient firm characteristics in 2002 to assess the degree to which our IV estimates may be driven by other channels unrelated to decentralization.

What are the threats to our identification strategy embodied in Equation (2)? The core assumption is that the characteristics of the donor firm in 1982 that we use to proxy for decentralization propensity are uncorrelated with the establishment's future growth shock between 2017 and 2002. This assumption could be violated if there was a feature of the donor firm that made it more likely that its spun-off plants would grow faster than average 20 years later (or similarly if it was sold to a high-growth firm where all establishments grew in later periods). For example, perhaps larger firms were more likely to sell (unobservably) high-potential growth establishments. We view this as *ex-ante* unlikely because this effect should be reflected in the plant's growth performance 1982-2002 which we capture by the observables in Equation (4) such as 2002 and 1982 plant characteristics and 2002 firm characteristics. However, a placebo test of this concern is whether we observe faster growth in a spun-off establishment where decentralization should be irrelevant. We consider establishments that become single-unit firms as a placebo. Since our MOPS decentralization measure refers to decentralization from the headquarters to the local plant manager, the decentralization channel is irrelevant for single-unit firms, whereas other potential channels would apply for both single-unit and multi-unit firms.

A second concern is that although Equation (2) holds, the causal mechanism is not through decentralization but rather through some other inherited organizational feature from the donor firm. Because we have rich organizational information from MOPS we can go further and directly investigate this alternative channel by repeating our baseline estimation strategy, but using other organizational measures than decentralization. In particular, we may be concerned that the structured management practices (that Bloom et al (2019) show to be correlated with short-run growth) are also inherited from donor firms and confound the decentralization effect. Consequently, we directly control for a plant's structured management score to assess whether our results are driven by the persistence of management practices. We find that our decentralization term dominates with a little-changed coefficient on decentralization. Our view is that plant autonomy, once granted, is particularly hard to rescind, so the long-run donor firm design we have developed here is well suited to examining the impact of decentralization.

waves. In the event, this did not buy us much more, as a single variable such as size tends to be a very strong predictor of decentralization as discussed below.

4 Main MOPS Results

4.1 Decentralized and Growth: OLS Results

Our first result is that a plant's baseline degree of decentralization is robustly correlated with its subsequent, long-run growth. Table 2 shows estimates of Equation (1) using the growth in the establishment's total value of shipments in Panel A, the growth in value added (shipments less materials) in Panel B, and employment in Panel C, all taken from the Census of Manufacturers. Column (1) regresses annualized growth over the fifteen-year period of 2002–2017 on decentralization in 2005 controlling for for industry dummies. The later columns cumulatively add more controls. It is immediately clear across all three outcomes and all four specifications, the coefficient on decentralization is positive and statistically significant.

In column (2) of Table 2, we control for the plant's log employment and log wagebill. In terms of magnitudes, the coefficient of 2.256 for the relationship between baseline decentralization and sales growth in our preferred specification is non-trivial. From Table 1, we can see that the standard deviation of our decentralization index is 0.195. Hence, a one standard deviation increase in decentralization is associated with a 0.440 (= 2.256 x 0.195) percentage point increase in sales growth per year. The relationship between decentralization and value-added growth is nearly identical, suggesting a one standard deviation rise in decentralization is associated with a 0.433 (= 2.223 x 0.195) percentage point increase in value-added growth. The correlation with employment growth is somewhat smaller in magnitude: about 0.311 (= 1.594 x 0.195) percentage points. This suggests a net positive association of decentralization with labor productivity, a finding we will assess more directly in the next table. In column (3) we also include the plant's management score and in column (4) we additionally include controls for the current parent firm's characteristics (in 2002). The coefficient on decentralization falls only slightly moving from the basic specification to the final column with a rich set of controls (e.g. from 2.322 to 2.191 in Panel A).

Our outcomes in Table 2 are in "DHS" form so combine both the intensive and extensive margins so a natural question is whether our growth results are driven by plant exit or by growth among the survivors. Appendix Table C4 explicitly puts plant exit as the outcome in Equation (1) to consider the extensive margin. Either in our baseline specification or one with the most general set of controls, we find that decentralized establishments are more likely to survive between 2002 and 2017 than their centralized counterparts. A one standard deviation increase in decentralization is associated with a fall in exit of 1.38 percentage points.[11]

In Table 3 we look more directly at the growth of plant productivity by presenting estimates of the growth in shipments per worker (Panel A) and value-added per worker (Panel B) for the sample of plants that have positive sales, value-added, and employment in 2017.[12] In column (2), our baseline specification, the coefficient (standard error) on the relationship between shipments per worker and decentralization is 0.457 (0.245). This implies that a one standard deviation increase in decentralization is associated with about a 0.089 (= 0.457 x 0.195) percentage point increase in annual labor productivity growth. The estimates on value-added per worker (in Panel B) are similar

[11] Since the DHS measure treats plant exit as equivalent to a -200% drop in plant outcomes, this extensive margin effect is equivalent to a $1.38 \times \frac{2}{2017-2002} = 0.184$ annualized size effect. In other words, approximately 42% of the relationship between decentralization and sales growth is due to changes in the extensive margin, and three-fifths is due to changes in the intensive margin.

[12] Although the exit channel is irrelevant for this sample, the relationship between size growth and decentralization is qualitatively similar to the estimates in Table 2.

in magnitude but more noisy, being only (weakly) significant in the final column.

There are two reasons why the coefficients in Table 3 are smaller than in Table 2. First, the coefficients on decentralization in the output equation are larger than those in the labor input equation. For example, the difference in column (2) of Table 2 between sales in Panel A and employment in Panel C is 0.662 =(2.256 - 1.594). This is still larger than the 0.457 coefficient in Panel A column (2) for a second reason. In contrast to the results for plant size where exiters really do have zero sales and employment, productivity can only be defined for survivors, so we are forced to drop exiting plants (which is why we have 6,100 plants in Table 3 instead of 7,300 plants in Table 2). Given our empirical finding of a positive relationship between survival and decentralization discussed above (and documented in Appendix Table C4), we *underestimate* the relationship between productivity and decentralization in Table 3 since the least productive establishments are most prone to exit.

4.2 Decentralization and growth: IV Results

The previous subsection provided robust evidence that establishments that were initially decentralized grew faster than their centralized counterparts over the 2002-2017 period. We now document that exogenous variation in a plant's decentralization, based on the decentralization of its previous owner in 1982 (and before) also appears to generate significantly faster long-run growth.

4.2.1 Core IV results

Table 4 shows the First Stage of our IV regressing current decentralization on our main proxy for the donor firm's propensity to decentralize - its (log) employment in 1982. The order of the columns mirrors the earlier tables where we build up from the most basic specification in column (1) to the most general regression in the final column with the richest set of controls. The instrument is highly significant in all columns and implies that larger firms are more likely to be centralized. One reason for this could be that coordinating actions across plants is more important in large corporations (for example, in pricing decisions). The firm size effect is not negligible. Using column (2), a one standard deviation increase in firm size (1.7 from Table 1) reduces the decentralization index by about 9% (=(1.7x0.01)/.195)) of standard deviation. Note that the addition of the plant's 1982 establishment size in column (2) of Table 4 means that the coefficient on firm size is identified solely from the variation in size of the donor firm's *other* establishments, and is, therefore, an important check on our conditional exogeneity assumption.

Table 5 presents our key results using our three major growth outcomes of sales, value-added and employment. There is a clear positive and significant causal effect of decentralization on sales in all columns of Panel A. The F-statistics at the base of the columns range between 42 and 55 and confirm that our instrument has considerable power. The estimates are again similar for value added in Panel B, which suggests that our results are not driven by increases in materials.

Interestingly, none of the impacts on employment are significant in Panel C of Table 5, except for the simplest specification of column (1). Indeed, in the final column, the point estimate is negative (albeit with a large standard error). This suggests that the positive impacts of decentralization on employment growth we saw in OLS may be spurious. The results are contrary to theories of "empire-building" where plant managers prioritize plant employment size over revenue growth. Rather, they

suggest that plant managers hold superior local information that allows them to operate more efficiently (Aghion and Tirole, 1997).

To look at the productivity impact more directly, Table 6 contains the IV productivity regression.[13] We find significant and large effects of decentralization. Using our baseline specification in column (2), we find that a one standard deviation increase in decentralization leads to a 3.16 percentage point increase in sales per worker and a 2.65 percentage point increase in value-added per worker.

The IV results are a lot larger than our OLS estimates. Focusing on the sales growth effects in the final columns, the IV results are about seven times as big. Below, we will investigate the validity of our instruments further, but in terms of interpretation, we think it is unlikely that the downward OLS bias can be fully explained by attenuation bias or omitted variables. More likely, the reason for the larger IV coefficients is due to the establishments that provide the identifying variation for the instrument (a LATE interpretation). Implicitly, the plants driving the IV estimates are those who have stuck with the organizational pattern of their original donor firms despite changes in ownership. In other words, plants that are able to most successfully resist changes in their organizational structure across ownership changes may also be the ones that benefit the most from decentralization. For example, consider a decentralized plant being acquired by a firm with a high propensity to centralize. Forcing such a firm to centralize may have strong negative effects on its performance if there is a general tendency towards over-centralization at the firm level, especially if acquired establishments tend to be in "peripheral" or new business areas in which the parent firm has weaker knowledge or expertise (Andrade et al., 2001; Renneboog and Vansteenkiste, 2019). Section 6 will present some evidence suggesting that this is indeed the case.

4.2.2 Placebo Test of IV

The nature of our instrument and the result of our specification checks support the interpretation that the variation behind the instrument is unrelated to unobserved shocks to plant outcomes between 2002 and 2017. However, the primary threat to a causal interpretation is that historic differences in firm ownership may be affecting subsequent growth through channels other than decentralization. Although our control for structured management practices in column (3) of Table 6 tests for this explanation, it is possible that there are other transmission mechanisms.

In Table 7, we present a direct test of the exclusion restriction through a placebo check. If persistent decentralization is the key mediator between the historic firm and subsequent plant growth, then the historic decentralization propensity of the donor firm should have no effect on plants where the decentralization channel is irrelevant. Since decentralization is defined as the delegation of control from firm headquarters to other plants, one observable placebo group should be the establishments that serve as the headquarters to its firm - i.e., single-unit firms.[14] We estimate the reduced form analog of our IV specification on all primary outcomes, and we find that the effects of historic firm size are generally small in magnitude and statistically insignificant in all cases. Moreover, given that the relationship between original firm size and subsequent decentralization is negative, the placebo estimates for value-added and value-added per worker are of the wrong sign (positive).

[13] The relationship between size growth and decentralization for this sample of survivors is again qualitatively similar to Table 5.

[14] For more details on how we construct the placebo sample, see Appendix A.2.

4.2.3 Other robustness tests of IV

We conduct a number of robustness tests on these results. Appendix Table C1 reports the first stage estimates for the survivors in the IV productivity sample. Reassuringly, the first-stage estimates are similar to the first-stage estimates in the overall sample, so our productivity estimates in Table 6 are not driven by potential differences in the persistence of decentralization for these plants.

Appendix Table C2 expands the set of historic 1982 firm characteristics to include historic firm wagebill and establishment count as additional instruments in our baseline IV specification. Note that by including the additional characteristics as logs, this specification implicitly allows for variation in average pay or establishment size to also explain subsequent decentralization and firm outcomes. The results corroborate the main results: we identify positive and significant decentralization effects on sales, value-added, and labor productivity, but zero effect on employment growth. Moreover, we cannot reject the null in the over-identification test that our IV estimates are the same across these instruments, despite the distinct conceptual channels they may capture. This provides further reassurance that the effects of these instruments are operating through decentralization rather than through an alternate channel.

Finally, although we consider plant-level outcomes between 2002–2017 and allow for exit in the size regressions, there is a concern that plants must survive at least until the 2010 MOPS to be included in our sample. If the plants that exit between 2002 and the 2010 MOPS have different responses to decentralization than the plants that survive until the 2010 MOPS, then this selection introduces bias in our IV estimates even if the instrument was as good as randomly assigned and operated solely through the decentralization channel. To test for the role of this selection channel, we estimate the *reduced form* equation implied by our IV system of equations (3) and (4). Since the reduced form equation does not require us to observe decentralization, we can estimate the regression on all acquired plants that fit our sample selection criteria without requiring that they are included in the 2010 MOPS. Appendix Table C3 reports reduced form estimates of the relationship between historic firm size and subsequent plant outcomes, including exit. Given that the relationship between historic firm size and decentralization is negative, the reduced form results are consistent with our IV results, but slightly noisier — plants that are expected to be more decentralized have faster output and labor productivity growth and are less likely to exit.

4.3 Reconciling OLS and IV estimates

Taken together, our OLS estimates in Section 4.2 and our IV estimates in Section 4.3 each presents robust evidence for a positive causal effect of decentralization on plant output growth. But there is uncertainly over the *magnitude* of the effects. For example, our preferred OLS estimate from Table 2 column (2) is that a 0.1 increase in the plant's decentralization index (equivalent to an extra 10% of decisions being delegated from headquarters to the plant) increases the plant's annualized sales growth by 0.23 p.p., while our preferred IV estimate from Table 5 column (2) is that a 0.1 increase in the plant's decentralization index increases the plant's annualized sales growth by 1.57 p.p. - a seven-fold difference (the results on value added suggest a five-fold difference).

In general, OLS and IV estimates may differ due to measurement error in the endogenous variable, omitted variables bias, or heterogeneous treatment effects (DiTraglia and Garcia-Jimeno, 2021; Ishimaru, 2024). Our preferred explanation for the difference between the OLS and IV results is

measurement error in our decentralization measure. In particular, we argue that since the components of the MOPS questionnaire are discrete and bounded between 0 and 1, measurement error tends to have a "mean-reverting" structure that biases OLS estimates downwards and IV estimates upwards.[15] As a result, the true effect tends to be bounded between the OLS and IV estimates. Without high quality validation data or additional instruments that are orthogonal to this measurement error, it is difficult to characterize the exact degree and nature of the measurement error, which is required to point-identify the true decentralization effect (Abowd and Stinson, 2013; Schennach, 2016; Bingley and Martinello, 2017). Instead, we can refine our bounds by conservatively assuming a wide range of empirically relevant models for measurement error, and under each assumed model, calculating the true coefficient that would be consistent with our estimated effects. For more details of our approach, including expressions for the expected bias in OLS and IV estimates under measurement error, see Appendix B.

Table 8 presents our adjusted estimates under three different assumed models for measurement error. The first model assumes that the measurement error for the treatment variable is equivalent to the error structure from a binary variable with a constant i.i.d. probability that the variable is mismeasured. The second and third model assumes in turn that the treatment is a simple arithmetic index over either three or six underlying binary variables, each with their own i.i.d. probability of mismeasurement. Under each assumed model, the bias and variance of the measurement error can by determined by the sample mean for decentralization and the empirical OLS and IV estimates. These statistics, in turn, are sufficient to characterize and adjust for the implied bias in our estimates. The ratio between the OLS and IV estimates differs slightly when we consider either sales growth or value-added growth, so we repeat the exercise for each of the two outcomes.[16] Recall that for sales growth we obtained an OLS effect of 2.3 and IV of 15.7. Our bias corrected effects have a tighter range between 5.7 and 7, closer to OLS than IV. Similarly, our bias corrected decentralization effects on value added growth range from 4.6 to 5.9 compared to 2.2 for OLS and 12.0 for IV.

We considered two other explanations of the differences in our OLS vs IV estimates for output growth, but found them less persuasive: omitted variables and heterogeneous treatment effects. First, the degree of omitted variables bias necessary to rationalize the difference between the two sets of results would not only need to be large, but it should also shrink as we progressively add controls that are correlated with the treatment and the outcome (Altonji et al., 2005; Oster, 2019). Instead, we find that our baseline results remain stable after controlling for either management or a flexible set of firm characteristics. In Panels A and B of Tables and 5, the coefficient on decentralization do not materially change as we add controls. The exception to this is employment growth, where the coefficient does fall across columns for IV in Table 5 and even flips sign in the final column (although is statistically insignificant). This suggests the IV helps deal with some of the omitted variable bias for jobs.

We also do not find clear evidence that heterogeneous treatment effects can explain the difference between our estimates. For example, we non-parametrically assess the potential scope for any

[15] To see this, consider the binary case. Measurement error means that true decentralization will be classified as full centralization (and vice versa). This type of measurement error leads a valid IV first stage to be "too low" by compressing support of the predicted value of the endogenous variable while leaving the reduced form regression unaffected. The resulting IV coefficient is then too large. The Appendix goes through this simple binary case and then a generalization to our context in detail.

[16] We do not use the estimates for employment in this exercise, since the difference between the two sets of estimates is less consistent and more likely to reflect omitted variables bias in the OLS. We discuss this point more extensively at the end of this section.

heterogeneous treatment effects by estimating flexible quantile regressions to assess the relationship between decentralization and the full distribution of establishment outcomes. If the effect of decentralization is highly heterogeneous, we would expect to see a greater dispersion of outcomes at more decentralized plants. We instead find results to the contrary - the distribution of sales growth at more decentralized plants is more compressed.[17]

In summary, we believe that the discrepancy between the OLS and IV estimates on output growth is due to non-classical measurement error, biasing OLS downwards and IV upwards. Our preferred treatment effect estimates imply a causal impact of decentralization on sales growth of between 6 and 7 and an effect on value-added growth of between 5 and 6. A 0.1 point increase in decentralization increases sales growth by between 0.6 and 0.7 p.p. per annum. This is closer to the OLS estimate of 0.2 than the IV estimate of 1.2. Of course, while there is still uncertainty over the exact magnitude of the decentralization effects, all these bias-adjusted suggest that decentralization has a meaningfully positive effect on output growth.

5 Analysis using the World Management Survey

As discussed above, we now turn to the WMS data which asked four identical questions to the MOPS survey over a similar time period, but covers 18 countries and used a different collection methodology. Our baseline estimates regress the annualized six-year growth rates of sales and employment on the initial level of decentralization. For example, we regress the growth of sales 2006-12 on decentralization in 2006 for the firms surveyed in 2006 (and similarly the 2009-15 growth on 2009 decentralization and 2013-2019 growth on 2013 decentralization).[18]

5.1 Basic WMS Results

Table 9 presents the WMS results in an analogous way to those from MOPS with the sales growth regressions in Panel A and the jobs growth regressions in Panel B. We see a very clear picture of a strong and positive relationship between growth and initial decentralization. Column (1) controls only for NAICS5 trends (SIC3 dummies), year dummies, and country dummies. Column (2) adds in initial employment.[19] Since decentralization in the WMS is parameterized as a z-score, we can interpret the coefficients directly as the average change for a 1 s.d. increase in the decentralization score. The coefficients in Column (2) imply that a one standard deviation increase in the decentralization index is associated with a 0.621 percentage point increase in annual sales growth and a 0.399 percentage point increase in annual employment growth. These are a bit larger than the MOPS results where the OLS marginal effects were 0.440 and 0.311 for sales and jobs respectively, but are in the same ballpark. Column (3) includes initial management scores constructed as in Bloom and Van Reenen (2007) and Column (4) includes a wide range of other firm and plant controls.[20] Although

[17]Note that the sample in our IV regressions comprise of establishments that have changed firm ownership between 1982 and 2002, whereas the sample in our OLS regressions include all plants for which we have a valid measure of decentralization from 2005. Although in principle, these sample differences could contribute to the difference between our estimates, we also find that the IV estimates are larger than comparable OLS regressions that are estimated on exactly the same sample.

[18]Since we match WMS to accounting data, it is harder to implement the very long 15 year differenced specifications we can perform on MOPS, as sales and employment are not mandatory accounting items for SMEs in most countries.

[19]The wagebill is missing from accounts data for too many firms to include it as a control

[20]These additional controls include union density, whether the firm is a foreign multinational, the fraction of employees with a college degree, firm size, average hours worked, whether the CEO is onsite, whether the firm is

these controls are jointly (and sometimes individually) significant, the coefficients and significance of decentralization are largely unchanged from the earlier columns, similarly to MOPS. The coefficient on decentralization in the productivity regressions in Panel C of 9 are positive, but only (weakly) significant in the last two columns.

Overall, we find the WMS results in Table 9 are remarkably consistent with those of US MOPS. Given the standard errors, we cannot reject surprisingly similar magnitudes, despite the great differences in the samples and survey methods.

5.2 Alternative IV strategy for WMS

We do not have the same rich data on ownership changes in the WMS as we have from the US Census LBD, so we cannot implement the donor firm IV. However, we consider the idea of using decentralization in the industry of the technology leader (the US) as an IV for decentralization in other countries. The idea is that the exogenous element of decentralization that is driven by fundamental technological causes, rather than country-industry-specific performance, is reflected in the decentralization status of US industries.

Existing work suggests that culture is an important driver of decentralization decisions (e.g. Hofstede, 2001). We calculate the average decentralization in each US SIC three-digit industry (NAICS-5 digit) using only the US firms. Empirically, we find that the industry-level correlation of US decentralization was much stronger with other English-speaking countries (Australia, Canada, New Zealand, and the UK) than other nations. The binscatter in Figure C2 shows a strong positive correlation between industry decentralization in these "Anglo" countries, whereas the relationship was generally weaker for other countries (suggesting underlying cultural similarities). Hence, we focus on this sub-sample of the data. We then use this as a shifter of decentralization in the four Anglo-Saxon countries in the same industries. Column (1) of Table C5 estimates a similar specification to column (3) of Table 9 for the subset of non-US Anglo-Saxon countries in the WMS. Since the instrument is at the SIC three-digit industry level, we can longer include three-digit dummies, so we include two-digit dummies and identify from the three-digit variation within a two-digit industry. All standard errors are clustered at the three-digit level.

Although we lose three-quarters of all firms in Table C5 compared to Table 9 due to focusing on the "Anglo" nations, and are using more conservative standard errors, we still observe a positive and significant relationship between sales growth and initial decentralization in column (1). The coefficients are slightly larger than those of Table 9 (0.997 vs. 0.633). Column (2) presents the First Stage. Consistent with the industry-level bivariate scatterplot of Figure C2, there is a positive and significant correlation between US industry decentralization and that of Anglo-Saxon firms. The F-statistic of 7 is not strong, however.[21] Column (3) presents the IV results, showing a large positive coefficient of decentralization on growth that is significant at the 10% level. These results are similar to the findings in MOPS, where the IV coefficient is larger than the OLS one.

We repeat these specifications using employment growth as an outcome in Panel B of Appendix Table C5. The IV estimate for the effects of decentralization on firm employment is larger than the effects on sales, although given the size of the standard errors we cannot reject the equality of the

publicly owned, a multi-plant dummy, average plant size and dummies for ownership type (government, dispersed, family and other private).

[21] Note that the same specification in column (2) on the non-Anglo-Saxon firms has essentially zero correlation: a coefficient (standard error) of 0.001(0.041).

coefficients. We use labor productivity as an outcome in Panel C. Here, the OLS results in column (1) are positive, but small and insignificant. The coefficient on the IV results is negative, but also completely insignificant. Given the size of the sample and the lower power of the IV, we conclude that although there is evidence for a positive and causal impact of decentralization on size in the WMS, we cannot say anything very clear about productivity.

5.3 Plant-level regressions for employment growth

The WMS analysis is based on firm-level regressions from BVD as sales information is unavailable at the establishment level in company accounts. However, the WMS did ask for plant employment, so we can consider this as an outcome. Note that this can only be done on observations that are in the WMS panel in order to have plant-level employment (in Table 9 we exploit firm-level accounting data from other sources generally available in each year, even if the WMS survey was not performed on that firm in that year). Since the WMS panel is very unbalanced, we focus on the 2006 decentralization variable and follow the evolution of employment through to the next wave of the WMS. Here, we can track over a thousand plants until the next wave of the WMS, which is by far the largest sample for a long-difference regression. Figure C1 has the binscatter between plant employment growth and initial decentralization, showing a positive relationship.

5.4 Summary on WMS results

The international results from the WMS broadly confirm what we have found in US MOPS. There is a robust and positive correlation between decentralization and sales growth, which is smaller in OLS than IV.

6 Digging deeper: Reasons for over-centralization and mechanisms behind the decentralization effect on growth

We have documented evidence that decentralized establishments grow systematically faster than their more centralized counterparts in two complementary datasets both in OLS and using alternative IV designs. This raises two connected questions. If decentralization benefits growth, why do more firms remain centralized? And second, what is the mechanisms why decentralization aids growth?

To investigate this, we first investigate the hypothesis that there may be an inefficiently high degree of centralization by examining whether founders may be reluctant to cede control. Then, we examine whether decentralized firms are more agile in responding to new market opportunities, which would be the case if local managers were better able to exploit their local information.

6.1 Why are so many firms so centralized?

If decentralization on average creates faster growth, why is it the case that so many firms remain centralized? Growth is not the same as profitability, of course, and it could be the case that costs rise as fast as output. We tended to find positive effects, however, and could not detect much effects on much higher costs. For example, employment rises less quickly than sales and we do not find that average wages rise in response to decentralization.

One possibility is that some owner-CEOs may inefficiently hoard power at the center. A class of firms where this might be the case is firms controlled by the founder, even when the firm has sufficiently matured that we would expect the founder to bring in more professional management. There are numerous examples of this in the economics and finance literature, history and popular culture.[22] For example, (Johnson et al., 1985) examined event studies around the death of a CEO of publicly listed US firms. They found that cumulative abnormal returns fell in general after a CEO death, but actually *rose* if the CEO was also a founder.

MOPS does not have founder information, but WMS does. Hence we re-estimate Table 9 but allow an interaction between decentralization and a dummy for whether a firm was old (above median age) and still controlled by the founder in Table 10. We also include linear terms in decentralization and the "old founder" dummy. The results clearly show that the coefficient on the decentralization interaction is significant and positive in the sales and employment growth equations. Indeed, in Panel B the linear decentralization coefficient (reflecting the impact of decentralization on non-old founder growth) although positive, is statistically insignificantly different from zero.

The magnitudes are economically non-trivial. Whereas a one standard deviation increase in decentralization in column (2) of Table 10 is predicted to increase growth by 0.4 pp a year for most firms, it increases growth by 2.7 pp for old founders. This seems consistent with the hypothesis that there may be systematic power-hoarding by many firms, and this might explain the positive impact of decentralization on growth.

6.2 Decentralization and Market Size opportunities

We next turn to another leading hypothesis on the benefits of autonomy, namely that decentralized organizations may respond with greater agility to new market opportunities. We explore the mechanism of adaptability by testing whether decentralized plants respond more sensitively to opportunity shocks. To define "opportunity shocks", we begin by following Mayer et al. (2020) and construct a measure of exogenous increases to market size driven by shocks to export demand at the six-digit NAICS level as:

$$\Delta OPP_{n,2002-17} = (X/Y)_{n,2002} * \sum_{i,c,c \neq US} [(X_{njc,2002}/X_{n,2002}) * \Delta(M_{jc,2002-17})] \tag{5}$$

where $X_{njc,2002}$ denotes exports from industry n of product j to country c in 2002, $X_{n,2002}$ denotes total exports from industry n in 2002, and $\Delta(M_{jc,2002-17})$ denotes the change in imports of product j by country c from all countries (other than the US) between 2002 and 2017 (using the Davis-Haltiwanger-Schuh transformation). This industry-level export measure is then interacted with $(X/Y)_{n,2002}$, the ratio of exports to shipments in industry n in 2002. The advantage of this measure (as opposed to a firm-level measure) is that it can be defined for non-exporting firms as well as exporters.

The idea behind this design is that a firm in an industry which is heavily exporting in 2002 to a market (product-country pair) where demand grows strongly over the next 15 years (e.g. selling computers in China), will enjoy a large exogenous increase to its demand. Unfortunately, we do not have the sufficiently disaggregated export data to implement this on WMS, so we focus on MOPS

[22] For example, (Jayaraman et al., 2000; Wasserman, 2003; Lee et al., 2017). An example in popular culture would be Logan Roy from the HBO hit series 'Succession'.

data for this exercise.

The opportunity shock is in long-differences (ΔOPP_{it}) and our preferred method is to use Census years to maximize the number of observations (so those which end in a 2 or 7). The firm export data comes from LFFTD. The country-product trade data for exports is from UN COMTRADE (available 1962-2018) and the industry-level export and shipments data is aggregated from the Census of Manufactures.

We then generalize Equation (1) to estimate:

$$\begin{aligned}\Delta lnY_{i,2002-17} =& \beta \left(\text{DEC}_i \times \Delta\text{OPP}_{i,2002-17}\right) + \mu \text{DEC}_i + \alpha \Delta\text{OPP}_{i,2002-17} \\ &+ \Theta \left[l_{i,2002}, w_{i,2002}, (l_{i,2002}, w_{i,2002}) \times \Delta\text{OPP}_{i,2002-17}\right] + \gamma_n + \Delta\epsilon_{i,2002-17},\end{aligned} \tag{6}$$

where $\Delta\text{OPP}_{i,2002-17}$ is our constructed exogenous trade shock. Note that since γ_n, our baseline industry fixed effects, are at the more aggregated NAICS4 level, while our industry shocks are defined at the NAICS6 level, the linear industry shock is separately identified from the within NAICS4 variation in export demand.

Given the exogeneity of the demand shock, the main coefficient of interest, β, is identified under a similar assumption to our assumption for Equation (1). Namely, we assume that conditional on plant size, the plant's degree of decentralization is uncorrelated to the plant's unobserved sensitivity to external demand shocks. Ideally, we would relax this assumption using a generalization of the IV approach from Equations (3) and (4), but precisely estimating β is a demanding exercise, and our variation from the instrument is not sufficiently strong to detect meaningful differences in the responses to shocks. To mitigate concerns with omitted variables bias, we focus on outcomes that are consistent between the OLS and the IV regressions — sales growth and value-added growth.

Table 11 presents results from expanding our OLS specification to allow for differential effects of decentralization by the industry's exposure to market size shocks (as in Equation (5)). The controls are consistent with those in Table 2 with the simplest specifications in Column (1) and the most demanding specifications in Column (4). Across all specifications, we find that decentralized establishments respond more to these market opportunity shocks than centralized firms, as indicated by the positive and (weakly) significant interaction term between decentralization and the shock for output growth. this is true whether we use sales growth in Panel A or value added in Panel B.

The magnitudes are not trivial. For a firm with mean decentralization (0.37), according to column (2) of 11 the marginal effect of a demand shock for a plant with the mean level of decentralization is a 1.11 (= (0.371 x 4.102) - 0.409) percentage point increase in annual sales growth. By contrast, for a plant with a one standard deviation (0.195) higher level of decentralization, the same demand shock generates a 1.91 (= (0.566 x 4.102) - 0.409) percentage point increase in sales growth. So, a one standard deviation increase in a plant's decentralization index increases its responsiveness to demand shocks by an extra 72% (= (1.91 - 1.11)/1.11) relative to the average plant in our sample.

The results in this section suggest that an important mechanism for the faster growth of decentralzied firms is that react with more agility to market opportunities. And one reason why more firms remain too centralized is that older owner-managers are reluctant to "let go of the reins" as they exercise their preferences for power (e.g. Fehr et al, 2013).

7 Conclusions

We present what we believe is the first analysis of the effects of decentralization on medium-run micro growth (over one to two decades). It is also the first causal design in the private sector and in a set of developed countries. We exploit panel data on (i) US establishments (MOPS) and (ii) cross-country firm-level data (WMS). In both datasets, we find compelling evidence that more autonomous establishments grow faster. We develop a new "donor firm" IV methodology exploiting exogenous variation in plant decentralization based on the organizational persistence of plants spun off from firms of different organizational firms. This uses decentralization-related characteristics of the firms who originally owned a plant, many decades after the plant switched owners.

In summary, our key findings are that more decentralized plants (whose local managers have more autonomy from headquarters over capital, labor, innovation, product, pricing, and pay) firms grow faster than their centralized counterparts. This correlation appears to be causal when using quasi-exogenous variation in plant decentralization using our "donor firm" IV strategy. Moreover, the growth-decentralization relationship is also apparent in cross-country firm-level panel data from the WMS both in OLS and using an alternative "cultural" IV. Finally, we show that one mechanism underlying the positive impact of decentralization on growth is that local managers appear to respond with more agility to export market size shocks, consistent with more efficient use of local information.

Although we believe the evidence for these benefits of decentralization is compelling, the question remains of why more firms do not grant their local managers more autonomy. We suggest that this is because there is a tendency for CEOs to hoard power and present suggestive evidence showing our results are much stronger in older firms that remain under the control of the founder. This is an area where more evidence is needed.

References

Abdulkadiroğlu, Atila, Joshua D. Angrist, Peter D. Hull, and Parag A. Pathak (2016) "Charters without Lotteries: Testing Takeovers in New Orleans and Boston," *American Economic Review*, 106 (7), 1878–1920.

Abowd, John M. and Martha H. Stinson (2013) "Estimating Measurement Error in Annual Job Earnings: A Comparison of Survey and Administrative Data," *The Review of Economics and Statistics*, 95 (5), 1451–1467.

Acemoglu, Daron, Philippe Aghion, Claire Lelarge, John Van Reenen, and Fabrizio Zilibotti (2007) "Technology, Information, and the Decentralization of the Firm," *The Quarterly Journal of Economics*, 122 (4), 1759–1799.

Acemoglu, Daron, Simon Johnson, and James Robinson (2005) "The Rise of Europe: Atlantic Trade, Institutional Change, and Economic Growth," *American Economic Review*, 95 (3), 546–579.

Acemoglu, Daron, Simon Johnson, and James A. Robinson (2001) "The Colonial Origins of Comparative Development: An Empirical Investigation," *American Economic Review*, 91 (5), 1369–1401.

Aghion, Philippe, Nicholas Bloom, Brian Lucking, Raffaella Sadun, and John Van Reenen (2021)

"Turbulence, Firm Decentralization, and Growth in Bad Times," *American Economic Journal: Applied Economics*, 13 (1), 133–169.

Aghion, Philippe and Jean Tirole (1997) "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105 (1), 1–29.

Aghion, Phillipe, Nicholas Bloom, and John Van Reenen (2014) "Incomplete Contracts and the Internal Organization of Firms," *The Journal of Law, Economics, and Organization*, 30 (suppl_1), i37–i63.

Altonji, JosephÂ G., ToddÂ E. Elder, and ChristopherÂ R. Taber (2005) "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools," *Journal of Political Economy*, 113 (1), 151–184.

Andrade, Gregor, Mark Mitchell, and Erik Stafford (2001) "New Evidence and Perspectives on Mergers," *Journal of Economic Perspectives*, 15 (2), 103–120.

Atalay, Enghin, Ali Hortaçsu, Mary Jialin Li, and Chad Syverson (2019) "How Wide Is the Firm Border?" *The Quarterly Journal of Economics*, 134 (4), 1845–1882.

Balán, Pablo, Augustin Bergeron, Gabriel Tourek, and Jonathan L. Weigel (2022) "Local Elites as State Capacity: How City Chiefs Use Local Information to Increase Tax Compliance in the Democratic Republic of the Congo," *American Economic Review*, 112 (3), 762–797.

Bandiera, Oriana, Michael Carlos Best, Adnan Qadir Khan, and Andrea Prat (2020) "The Allocation of Authority in Organizations: A Field Experiment with Bureaucrats," Working Paper 26733, National Bureau of Economic Research.

Bingley, Paul and Alessandro Martinello (2017) "Measurement Error in Income and Schooling and the Bias of Linear Estimators," *Journal of Labor Economics*, 35 (4), 1117–1148.

Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John Van Reenen (2019) "What Drives Differences in Management Practices?" *American Economic Review*, 109 (5), 1648–1683.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen (2012) "The Organization of Firms Across Countries," *The Quarterly Journal of Economics*, 127 (4), 1663–1705.

Blundell, Richard, David A. Green, and Wenchao Jin (2016) "The UK wage premium puzzle: How did a large increase in university graduates leave the education premium unchanged?" Working Paper W16/01, IFS Working Papers.

Bort, Julie (2013) "Steve Ballmer Tells Employees: 'Stephen Elop Will Be Coming back to Microsoft'," September.

Bradley, Steven W., Howard Aldrich, Dean A. Shepherd, and Johan Wiklund (2011) "Resources, environmental change, and survival: asymmetric paths of young independent and subsidiary organizations," *Strategic Management Journal*, 32 (5), 486–509.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt (2002) "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence," *The Quarterly Journal of Economics*, 117 (1), 339–376.

Buffington, Catherine, Lucia Foster, Ron Jarmin, and Scott Ohlmacher (2016) "The Management and Organizational Practices Survey (MOPS): An Overview*," working Paper, U.S. Census Bureau, Center for Economic Studies.

Caliendo, Lorenzo, Giordano Mion, Luca David Opromolla, and Esteban Rossi-Hansberg (2020) "Productivity and Organization in Portuguese Firms," *Journal of Political Economy*, 128 (11), 4211–4257.

Caroli, Eve and John Van Reenen (2001) "Skill-Biased Organizational Change? Evidence from A Panel of British and French Establishments," *The Quarterly Journal of Economics*, 116 (4), 1449–1492.

Chow, Melissa C., Teresa C. Fort, Christopher Goetz, Nathan Goldschlag, James Lawrence, Elisabeth Ruth Perlman, Martha Stinson, and T. Kirk White (2021) "Redesigning the Longitudinal Business Database,"Technical Report w28839, National Bureau of Economic Research.

Clark, Damon (2009) "The Performance and Competitive Effects of School Autonomy," *Journal of Political Economy*, 117 (4), 745–783.

Dahl, Robert A. (1957) "The concept of power," *Behavioral Science*, 2 (3), 201–215.

Davis, Steven J., John C. Haltiwanger, and Scott Schuh (1998) *Job Creation and Destruction*, 1: The MIT Press.

Dell, Melissa (2010) "The Persistent Effects of Peru's Mining Mita," *Econometrica*, 78 (6), 1863–1903.

Dessein, Wouter, Desmond (Ho-Fu) Lo, and Chieko Minami (2019) "Coordination and Organization Design: Theory and Micro-evidence," CEPR Discussion Papers 13938, C.E.P.R. Discussion Papers.

DiTraglia, Francis J. and Camilo Garcia-Jimeno (2021) "A Framework for Eliciting, Incorporating, and Disciplining Identification Beliefs in Linear Models," *Journal of Business & Economic Statistics*, 39 (4), 1038–1053.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan (2018) "The Value of Regulatory Discretion: Estimates from Environmental Inspections in India," *Econometrica*, 86 (6), 2123–2160.

Erlichman, Jon (2019) "Three years after Microsoft acquisition, LinkedIn keeps quietly climbing - BNN Bloomberg," October.

Feldman, Maryann P., Serden Ozcan, and Toke Reichstein (2019) "Falling Not Far from the Tree: Entrepreneurs and Organizational Heritage," *Organization Science*, 30 (2).

Galiani, Sebastian, Paul Gertler, and Ernesto Schargrodsky (2008) "School decentralization: Helping the good get better, but leaving the poor behind," *Journal of Public Economics*, 92 (10), 2106–2120.

Garicano, Luis and Luis Rayo (2016) "Why Organizations Fail: Models and Cases," *Journal of Economic Literature*, 54 (1), 137–192.

Gibbons, Robert and Rebecca Henderson (2012) "Relational Contracts and Organizational Capabilities," *Organization Science*, 23 (5), 1350–1364.

Gibbons, Robert, Niko Matouschek, and John Roberts (2012) "Decisions in Organizations," in Gibbons, Robert and John Roberts eds. *The Handbook of Organizational Economics*: Princeton University Press.

Gibbons, Robert and John Roberts (2012) "The Handbook of Organizational Economics," economics Books, Princeton University Press.

Greene, Jay and Matthias Verbergt (2016) "Microsoft Takes Another Step Out of the Cellphone Business, as Nokia Plots a Return," *Wall Street Journal*.

Guadalupe, Maria and Julie Wulf (2010) "The Flattening Firm and Product Market Competition: The Effect of Trade Liberalization on Corporate Hierarchies," *American Economic Journal: Applied Economics*, 2 (4), 105–127.

Hanushek, Eric A., Susanne Link, and Ludger Woessmann (2013) "Does school autonomy make sense everywhere? Panel estimates from PISA," *Journal of Development Economics*, 104, 212–232.

Hart, Oliver and John Moore (2005) "On the Design of Hierarchies: Coordination versus Specialization," *Journal of Political Economy*, 113 (4), 675–702.

Hayek, F. A. (1945) "The Use of Knowledge in Society," *The American Economic Review*, 35 (4), 519–530.

Hofstede, Geert (1984) *Culture's Consequences: International Differences in Work-Related Values*: SAGE.

Ichniowski, Casey, Kathryn Shaw, and Giovanna Prennushi (1997) "The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines," *The American Economic Review*, 87 (3), 291–313.

Ishimaru, Shoya (2024) "Empirical Decomposition of the IV-OLS Gap with Heterogeneous and Nonlinear Effects," *The Review of Economics and Statistics*, 106 (2), 505–520.

Jayaraman, Narayanan, Ajay Khorana, Edward Nelling, and Jeffrey Covin (2000) "CEO Founder Status and Firm Financial Performance," *Strategic Management Journal*, 21 (12), 1215–1224.

Johnson, Bruce, Robert Magee, Nandu Nagarajan, and Harry Newman (1985) "An analysis of the stock price reaction to sudden executive deaths: Implications for the managerial labor market," *Journal of Accounting and Economics*, 7 (1), 151–174.

Kala, Namrata (2019) "The Impacts of Managerial Autonomy on Firm Outcomes," Working Paper 26304, National Bureau of Economic Research.

Katayama, Hajime, Kieron J. Meagher, and Andrew Wait (2018) "Authority and communication in firms," *Journal of Economic Behavior & Organization*, 155, 315–348.

Lange, Oskar (1936) "On the Economic Theory of Socialism: Part One," *The Review of Economic Studies*, 4 (1), 53–71.

Lee, Joon Mahn, Byoung-Hyoun Hwang, and Hailiang Chen (2017) "Are founder CEOs more overconfident than professional CEOs? Evidence from S&P 1500 companies," *Strategic Management Journal*, 38 (3), 751–769.

Marx, Karl (1867) *Capital: A Critique of Political Economy (I): The Process of Capitalist Production*, 1: McMaster University Archive for the History of Economic Thought.

Mayer, Thierry, Marc J. Melitz, and Gianmarco I.P. Ottaviano (2020) "Product Mix and Firm Productivity Responses to Trade Competition," *The Review of Economics and Statistics*, 1–59.

McElheran, Kristina (2014) "Delegation in Multi-Establishment Firms: Evidence from I.T. Purchasing," *Journal of Economics & Management Strategy*, 23 (2), 225–258.

Novet, Jordan (2019) "Microsoft paid $26 billion for LinkedIn, then mostly left it alone—and CEO Jeff Weiner is good with that," December.

O'Brien, Kevin J. (2010) "Nokia's New Chief Faces Culture of Complacency," *The New York Times*.

Oster, Emily (2019) "Unobservable Selection and Coefficient Stability: Theory and Evidence," *Journal of Business & Economic Statistics*, 37 (2), 187–204.

Parsons, Talcott (1963) "On the Concept of Political Power," *Proceedings of the American Philosophical Society*, 107 (3), 232–262.

Rajan, Raghuram G and Julie Wulf (2006) "The Flattening Firm: Evidence from Panel Data on the Changing Nature of Corporate Hierarchies," *The Review of Economics and Statistics*, 88 (4), 759–773.

Renneboog, Luc and Cara Vansteenkiste (2019) "Failure and success in mergers and acquisitions," *Journal of Corporate Finance*, 58, 650–699.

Russell, Bertrand (1938) *Power: A New Social Analysis*: W.W.Norton.

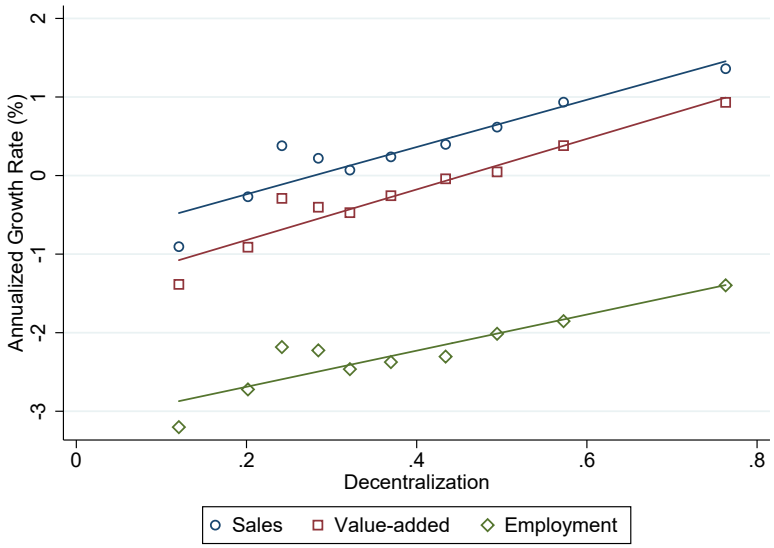
Schennach, Susanne M. (2016) "Recent Advances in the Measurement Error Literature," *Annual Review of Economics*, 8, 341–377.

Wasserman, Noam (2003) "Founder-CEO Succession and the Paradox of Entrepreneurial Success," *Organization Science*, 14 (2), 149–172.

Weber, Max (1978) *Economy and Society: An Outline of Interpretive Sociology*: University of California Press.

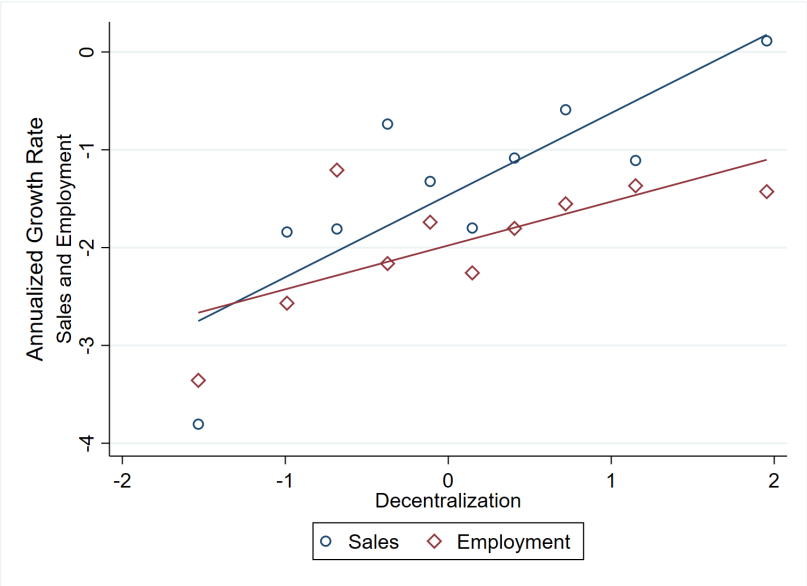
Tables and Figures

Figure 1: US plant growth is increasing in average decentralization (MOPS)



Note: Binscatter of average annualized plant level sales, value-added, and employment growth between 2002 and 2017 (measured by DHS) as a function of the plant's average decentralization status from all waves of the MOPS. Decentralization is a continuous index for the plant's average degree of decentralization across all six questions and all waves of the MOPS.

Figure 2: Firm growth is increasing in initial decentralization (WMS)



Note: Binscatter of average annualized firm-level sales and employment growth from BvD (in percentage points) as a function of initial average decentralization from WMS. WMS data pooled across all available years.

Table 1: Summary statistics

(a) MOPS samples

Variable:	Mean	SD
Decentralization index (2005)	0.371	[0.195]
Baseline characteristics:		
Log shipments	10.10	[1.029]
Log value-added	9.300	[1.063]
Log employment	4.601	[1.005]
Log firm employment	7.492	[1.678]
Share of firm shipments	0.230	[0.307]
Changes (2002-17):		
Shipment growth	-0.245	[6.669]
Value-added growth	-0.733	[7.127]
Employment growth	-2.657	[5.509]
Exit (LBD)	0.130	[0.336]
Industry demand shock	0.192	[0.228]
Additional IV stats:		
Decentralization index (any)	0.374	[0.166]
1982 log employment	4.659	[1.123]
1982 log firm employment	7.474	[2.417]

(b) WMS samples

Variable:	Mean	SD
Decentralization (z-score)	0	[1]
Baseline characteristics:		
Log sales	11.28	[1.38]
Log employment (plant)	5.19	[0.89]
Log employment (firm)	5.79	[1.08]
Old Founder	0.11	[0.31]
Changes (6 years):		
Sales growth	-1.34	[11.60]
Firm employment growth	-1.93	[10.80]

Note: For panel A, establishment-level summary statistics for our baseline OLS sample (approximately 7,300 establishments). In the last three rows of panel A, we report additional summary statistics for plants in our IV sample (approximately 4,400 establishments). "Decentralization index (2005)" includes only valid responses from the 2005 recall questions of the MOPS, while "decentralization index (any)" includes responses from all waves of the MOPS. All baseline characteristics are measured in 2002. The industry demand shock is the industry-level export shock, as described in Equation 5. For panel B, establishment and firm-level summary statistics for our WMS OLS sample (approximately 8,900 establishments).

Table 2: Decentralization predicts subsequent plant size growth

	Long Differences (DHS), 2002-17			
	(1)	(2)	(3)	(4)
	A: Sales			
Decentralization	2.322***	2.256***	2.282***	2.191***
	(0.438)	(0.439)	(0.437)	(0.442)
	B: Value-added			
Decentralization	2.281***	2.223***	2.249***	2.158***
	(0.474)	(0.475)	(0.472)	(0.476)
	C: Employment			
Decentralization	1.669***	1.594***	1.620***	1.471***
	(0.368)	(0.367)	(0.366)	(0.369)
# Est.	7300	7300	7300	7300
# Firms	3400	3400	3400	3400
Controls:				
Baseline est. size		Y	Y	Y
Management			Y	Y
Baseline firm chars.				Y

Note: OLS coefficients with standard errors in parentheses (clustered at the firm level). Decentralization is a continuous index for the plant's average degree of decentralization across all six questions in MOPS in 2005. The sample is all plants with a valid decentralization score in 2005 and positive shipments, value-added, and employment in 2002. The dependent variable is the total annualized difference in total plant shipments (Panel A), value-added (Panel B), or employment (Panel C) between 2002 and 2017, as measured by DHS (Davis-Haltiwanger-Schuh) growth rates ($\Delta Y_{i,2002-17} = \frac{Y_{i,2017}-Y_{i,2002}}{(Y_{i,2017}+Y_{i,2002})/2} \times \frac{100}{2017-2002}$). "Baseline establishment size" controls allow for differences by the plant's log employment and log wagebill from the 2002 CMF. "Management" controls allow for differences by the plant's structured management score from the 2005 MOPS. "Baseline firm characteristics" controls allow for differences by the plant's 2002 ownership firm's log(employment), log(wagebill), log(establishment count in firm), and log(firm age) from the 2002 LBD, as well as the plant's share of total firm output from the 2002 CMF. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

Table 3: Decentralization predicts subsequent plant productivity growth

	Long Differences (Log), 2002-17			
	(1)	(2)	(3)	(4)
	A: Sales per worker			
Decentralization	0.413*	0.457*	0.453*	0.538**
	(0.249)	(0.245)	(0.245)	(0.248)
	B: Value-added per worker			
Decentralization	0.419	0.489	0.487	0.572*
	(0.321)	(0.320)	(0.320)	(0.322)
# Est.	6100	6100	6100	6100
# Firms	3000	3000	3000	3000
Controls:				
Baseline est. size		Y	Y	Y
Management			Y	Y
Baseline firm chars.				Y

Note: OLS coefficients with standard errors in parentheses (clustered at the firm level). Decentralization is a continuous index for the plant's average degree of decentralization across all six questions in MOPS in 2005. The sample is all plants with a valid decentralization score in 2005 and positive shipments, value-added, and employment in both 2002 and 2017. The dependent variable is the total annualized difference in total plant shipments per worker (Panel A) or value-added per worker (Panel B) between 2002 and 2017, as measured by logs. "Baseline establishment size" controls allow for differences by the plant's log employment and log wagebill from the 2002 CMF. "Management" controls allow for differences by the plant's structured management score from the 2005 MOPS. "Baseline firm characteristics" controls allow for differences by the plant's 2002 ownership firm's log(employment), log(wagebill), log(establishment count in firm), and log(firm age) from the 2002 LBD, as well as the plant's share of total firm output from the 2002 CMF. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

Table 4: Pre-merger firm characteristics predict subsequent decentralization

Dep. Variable	Average Decentralization			
	(1)	(2)	(3)	(4)
Orig. firm size	-0.00795***	-0.00988***	-0.00998***	-0.00938***
	(0.00122)	(0.00135)	(0.00135)	(0.00141)
# Est.	4400	4400	4400	4400
# Firms	1600	1600	1600	1600
Controls:				
Orig. est. size		Y	Y	Y
Management			Y	Y
Baseline firm chars.				Y

Note: OLS coefficients with standard errors in parentheses (clustered at the firm level). Original firm size is the log employment of the plant's owner firm in 1982. The sample is all plants that experienced a merger or acquisition event between 1982 and 2002, have a valid decentralization score in any wave of the MOPS, and positive shipments, value-added, and employment in 2002. The dependent variable is the plant's average decentralization score from all waves of the MOPS. "Original establishment size" controls allow for differences by the plant's log(employment) and log(wagebill) from the 1982 LBD. "Management" controls allow for differences by the plant's structured management score (constructed as in Bloom et al. (2019) from 16 question) questions from all waves of the MOPS. "Baseline firm characteristics" controls allow for differences by the plant's ownership 2002 firm's log(employment), log(wagebill), log(establishment count in firm), and log(firm age) from the 2002 LBD, as well as the plant's share of total firm output from the 2002 CMF. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

Table 5: Causal impacts of decentralization on plant size growth

	Long Differences (DHS), 2002-17			
	(1)	(2)	(3)	(4)
	A: Sales			
Decentralization	23.89***	15.70***	16.95***	13.67**
	(6.130)	(5.304)	(5.268)	(5.496)
	B: Value-added			
Decentralization	24.50***	12.09**	13.30**	9.401
	(6.747)	(5.746)	(5.704)	(6.024)
	C: Employment			
Decentralization	11.67***	2.013	3.162	-0.154
	(4.495)	(4.092)	(4.030)	(4.387)
# Est.	4400	4400	4400	4400
# Firms	1600	1600	1600	1600
First stage F	42.4	53.9	54.6	44.2
Controls:				
Orig. est. size		Y	Y	Y
Management			Y	Y
Baseline firm chars.				Y

Note: Two-stage least-squares coefficients with standard errors in parentheses (clustered at the firm level). We instrument for the plant's average decentralization status from all waves of the MOPS using the log employment of the plant's owner firm in 1982. The sample is all plants that experienced a merger or acquisition event between 1982 and 2002, have a valid decentralization score in any wave of the MOPS, and positive shipments, value-added, and employment in 2002. The dependent variable is the total annualized difference in total plant shipments (Panel A), value-added (Panel B), or employment (Panel C) between 2002 and 2017, as measured by DHS (Davis-Haltiwanger-Schuh) growth rates ($\Delta Y_{i,2002-17} = \frac{Y_{i,2017}-Y_{i,2002}}{(Y_{i,2017}+Y_{i,2002})/2} \times \frac{100}{2017-2002}$). "Original establishment size" controls allow for differences by the plant's log(employment) and log(wagebill) from the 1982 LBD. "Management" controls allow for differences by the plant's structured management score (constructed as in Bloom et al. (2019) from 16 question) questions from all waves of the MOPS. "Baseline firm characteristics" controls allow for differences by the plant's ownership 2002 firm's log(employment), log(wagebill), log(establishment count in firm), and log(firm age) from the 2002 LBD, as well as the plant's share of total firm output from the 2002 CMF. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

Table 6: Causal impacts of decentralization on plant productivity growth

	Long Differences (Log), 2002-17			
	(1)	(2)	(3)	(4)
	A: Sales per worker			
Decentralization	13.72***	16.23***	16.23***	16.61***
	(3.772)	(3.780)	(3.762)	(4.113)
	B: Value-added per worker			
Decentralization	15.32***	13.61***	13.63***	13.78***
	(5.074)	(4.755)	(4.720)	(5.117)
# Est.	3700	3700	3700	3700
# Firms	1400	1400	1400	1400
First stage F	35.9	48	48.4	38.4
Controls:				
Orig. est. size		Y	Y	Y
Management			Y	Y
Baseline firm chars.				Y

Note: Two-stage least-squares coefficients with standard errors in parentheses (clustered at the firm level). We instrument for the plant's average decentralization status from all waves of the MOPS using the log employment of the plant's owner firm in 1982. The sample is all plants that experienced a merger or acquisition event between 1982 and 2002, have a valid decentralization score in any wave of the MOPS, and positive shipments, value-added, and employment in both 2002 and 2017. The dependent variable is the total annualized difference in total plant shipments per worker (Panel A) or value-added per worker (Panel B) between 2002 and 2017, as measured by logs. "Original establishment size" controls allow for differences by the plant's log(employment) and log(wagebill) from the 1982 LBD. "Management" controls allow for differences by the plant's structured management score (constructed as in Bloom et al. (2019) from 16 question) questions from all waves of the MOPS. "Baseline firm characteristics" controls allow for differences by the plant's ownership 2002 firm's log(employment), log(wagebill), log(establishment count in firm), and log(firm age) from the 2002 LBD, as well as the plant's share of total firm output from the 2002 CMF. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

Table 7: IV placebo: historic donor firm size does not predict future growth for single-unit establishments

	Long Differences, 2002-17				
	(1)	(2)	(3)	(4)	(5)
Orig. firm size	-0.0672	0.00945	-0.0595	-0.0114	0.151
	(0.116)	(0.119)	(0.0969)	(0.0773)	(0.0930)
# Est.	9000	9000	9000	3600	3600
# Firms	9000	9000	9000	3600	3600
Outcome	Sales	VA	Emp	Sales/Emp	VA/Emp
Controls:					
Orig. est. size	Y	Y	Y	Y	Y

Note: OLS coefficients with standard errors in parentheses (clustered at the firm level). The sample is all plants that changed firm ownership between 1982 and 2002 and are single-unit firms in 2002. The definition of the instrument and all outcomes are the same as in Tables 5 and 6. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

Table 8: Measurement error adjusted decentralization effects

	Long Differences (DHS), 2002-17	
	Sales	Value-added
Baseline estimates:		
$\hat{\beta}_{\text{OLS}}$	2.256	2.223
$\hat{\beta}_{\text{IV}}$	15.70	12.09
$\hat{\beta}_{\text{IV}}/\hat{\beta}_{\text{OLS}}$	6.959	5.439
Adjusted estimates:		
Binary treatment	7.008	5.886
Index (3 components)	6.201	5.165
Index (6 components)	5.568	4.619

Note: Point estimates for the effect of decentralization on plant growth, adjusted for bias from (non-classical) measurement error. Each row assumes a different underlying structure for measurement error for decentralization. The first model assumes that the measurement error for the treatment variable is equivalent to the error structure from a binary variable with a constant i.i.d. probability that the variable is mismeasured. The second and third model assumes that the treatment is a simple mean over either three or six underlying binary variables each with their own i.i.d. probability of mismeasurement. For each outcome, the parameters of the assumed error structure is calibrated using the observed sample mean of decentralization (from Table 1) and the ratio of the corresponding baseline IV and OLS estimates (from column 2 of Tables 5 and 2, respectively) according to Equation 12.

Table 9: Decentralized firms have higher growth rates: Cross country WMS data

	6-year Long Differences (Log)			
	(1)	(2)	(3)	(4)
	A. Sales			
Decentralization	0.633***	0.621***	0.657***	0.705***
	(0.198)	(0.199)	(0.204)	(0.203)
Observations	9,155	9,155	9,155	9,155
	B. Employment			
Decentralization	0.421**	0.399*	0.377*	0.424*
	(0.213)	(0.216)	(0.224)	(0.227)
Observations	9,224	9,224	9,224	9,224
	C. Sales per worker			
Decentralization	0.204	0.216	0.292*	0.300*
	(0.175)	(0.175)	(0.176)	(0.178)
Observations	9,058	9,058	9,058	9,058
Controls:				
Initial size		Y	Y	Y
Management control			Y	Y
Extended controls				Y

Note: OLS regressions based on WMS plant-level data from 18 countries. Decentralization (z-score) is measured in various years between 2006 and 2009 and outcomes are subsequent growth in log(sales) in Panel A, log(employment) in panel B, and log(sales/employment) in Panel C. These are annualized change in log values over the subsequent six years. Standard errors are clustered by firm. Initial size controls are base year log(plant employment) and log(firm employment). Management controls are the z-score of management across 18 practices. Extended controls include basic controls plus whether the establishment belongs to a domestic multinational, whether the establishment belongs to a foreign multinational, the fraction of unionized employees, and the fraction of workers with a college degree. All columns control for three-digit industry dummies, country dummies, and time dummies. Up to 2,203 firms underlying these regressions. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

Table 10: The growth-decentralization relationship is strongest for old founders

	6-year Long Differences (Log)			
	(1)	(2)	(3)	(4)
	A. Sales (8,887 obs)			
Old Founder × Decentralization	2.253***	2.297***	2.312***	2.279***
	(0.830)	(0.826)	(0.825)	(0.809)
Decentralization	0.445**	0.429**	0.463**	0.511**
	(0.193)	(0.194)	(0.198)	(0.198)
Old Founder	0.695	0.750	0.690	0.229
	(0.787)	(0.788)	(0.788)	(0.792)
	B. Employment (8,977 obs)			
Old Founder × Decentralization	1.242**	1.271**	1.263**	1.238**
	(0.607)	(0.610)	(0.611)	(0.599)
Decentralization	0.328	0.302	0.281	0.322
	(0.224)	(0.228)	(0.234)	(0.237)
Old Founder	1.278**	1.298**	1.333**	0.914
	(0.626)	(0.630)	(0.638)	(0.658)
	C. Sales per worker (8,820 obs)			
Old Founder × Decentralization	1.074	1.100	1.128	1.134
	(0.731)	(0.714)	(0.713)	(0.708)
Decentralization	0.102	0.116	0.183	0.196
	(0.171)	(0.171)	(0.173)	(0.174)
Old Founder	-0.155	-0.070	-0.186	-0.258
	(0.655)	(0.651)	(0.651)	(0.658)
Controls:				
Initial size		Y	Y	Y
Management control			Y	Y
Extended controls				Y

Note: OLS regressions based on WMS plant-level data from 18 countries. Decentralization (z-score) is measured in various years between 2006 and 2009 and outcomes are subsequent growth in log(sales) in Panel A, log(employment) in panel B, and log(sales/employment) in Panel C. These are annualized change in log values over the subsequent six years. "Old Founders" are firms over a decade old who are controlled and owned by the original founder. Standard errors are clustered by firm. Initial size controls are base year log(plant employment) and log(firm employment). Management controls are the z-score of management across 18 practices. Extended controls include basic controls plus whether the establishment belongs to a domestic multinational, whether the establishment belongs to a foreign multinational, the fraction of unionized employees, and the fraction of workers with a college degree. All columns control for three-digit industry dummies, country dummies, and time dummies. Up to 2,203 firms underlying these regressions. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

Table 11: Decentralized plants are more responsive to demand shocks

	Long Differences (DHS), 2002-17			
	(1)	(2)	(3)	(4)
	A: Sales			
Shock × decen.	4.427***	4.315***	4.021**	3.678**
	(1.593)	(1.588)	(1.573)	(1.548)
Shock	-0.841	-0.691	-0.653	-0.556
	(0.840)	(0.834)	(0.822)	(0.812)
	B: Value-added			
Shock × decen.	5.385***	5.288***	5.007***	4.636***
	(1.716)	(1.710)	(1.694)	(1.672)
Shock	-1.405	-1.252	-1.217	-1.114
	(0.896)	(0.889)	(0.878)	(0.870)
# Est.	17500	17500	17500	17500
# Firms	6700	6700	6700	6700
Controls:				
Baseline est. size		Y	Y	Y
Management			Y	Y
Baseline firm chars.				Y

Note: OLS coefficients with standard errors in parentheses (clustered at the firm level). Decentralization is the plant's average decentralization status from all waves of the MOPS. Shock is a continuous measure of the exogenous export demand shock faced by the plant's NAICS6 industry, as defined in equation 5. The sample is all plants with a valid decentralization score in any wave of the MOPS and positive shipments, value-added, and employment in 2002. "Baseline establishment size" controls allow for differences by the plant's log employment and log wagebill from the 2002 CMF. "Management" controls allow for differences by the plant's structured management score from the 2005 MOPS. "Baseline firm characteristics" controls allow for differences by the plant's 2002 ownership firm's log(employment), log(wagebill), log(establishment count in firm), and log(firm age) from the 2002 LBD, as well as the plant's share of total firm output from the 2002 CMF. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

A Data Appendix

A.1 Management and Organizational Practices (MOPS) Survey Design

The Census Bureau, the National Science Foundation, and the MIT Initiative on the Digital Economy, the Sloan Foundation and the Kauffman Foundation jointly funded the Management and Organizational Practices Survey (MOPS). It was fielded in 2011 and 2016 as a supplement to the 2010 and 2015 Annual Survey of Manufacturers (ASM), with response mandated by law. The design was based on the World Management Survey (WMS) and was mailed to the establishment plant manager (see Bloom et al. (2019) or Brynjolfsson and McElheran 2016). It was adapted from the WMS format of open-ended telephone interviews to the closed format of the US survey through two years of development and cognitive testing by the Census Bureau. Paper copies of the survey instrument along with instructions for responding online were mailed to the physical addresses of all establishments in the ASM sample, attention of the plant manager. The typical respondent for the MOPS 2015 was a plant manager). Most respondents (58.4% in 2010 and 80% in 2015) completed the survey electronically, with the remainder completing the survey by paper. Non-respondents were mailed a follow-up letter after six weeks if no response had been received. A second follow-up letter was mailed if no response had been received after 12 weeks. The first follow-up letter included a copy of the MOPS instrument.

The sampling frames for the 2010 and 2015 MOPS were the 2010 and 2015 ASM respectively, which were around 50,000 plants in each wave. We had a response rate of 78% in 2010 and 74% in 2015. Bloom et al. (2019) report the results for linear probability models for each step in the sampling process. The establishments that were mailed and responded to the MOPS survey are somewhat larger and more productive compared to those that did not respond, but these differences are quantitatively small. Importantly, the decentralization questions are only asked for plants who said the headquarters was not on the same site. For most of our analysis, we further restrict the sample to establishments with at least four non-missing responses to the six decentralization questions.

Decentralization Measures The "Organization Module" of MOPS is only for plants where headquarters is off site - plants with headquarters on site (question 17) are told to skip this section. The survey contained six questions on decentralization with four of these covering the same domain as the WMS – plant manager autonomy over (a) capital investments (q23), (b) hiring of full time employees (q18), (c) product introduction ("innovation") (q20) and (d) advertising (q22) – with two additional question on (e) pay increases of at least 10% (q19), and (f) product pricing decisions (q21).

For each question, respondents were asked to choose among three options capturing where the specific decisions were made "only at this establishment", "only at headquarters", or "both at this establishment and at headquarters". There were five choices for the question on autonomy in capital investments, starting with "Under $1,000" up until "$1 million or more". We consider several ways of constructing a single index of decentralization as well as disaggregating question by question. For example, our main measure reflects the fraction of decisions decentralized, and is constructed by coding responses to each question to evenly spaced values between 0 and 1 (with 0 for the least decentralized response and 1 for the most decentralized response) and then averaging this score

across the six categories.

Other Measures in MOPS The survey also included measures of structured management practices; of decentralization from the plant manager to workers and the use of data driven decision-making (DDD). It also had some background questions on the establishment (such as the proportion of employees with college degrees) and respondent.

Performance Measures from other datasources In addition to our management data, we also use data from other Census and non-Census data sets to create our measures of performance (size, productivity, growth, and survival). We use establishment-level data on sales and labor inputs from the ASM to create measures of labor productivity. We also use measures of material inputs to create value added.

We construct measures of establishment survival and growth using the Longitudinal Business Database (LBD). In a handful of cases, a plant with valid decentralization data is missing one or more performance measures or additional MOPS controls. To fully utilize the decentralization data while ensuring a consistent sample, we impute the missing performance measures by using the industry by year-recall group mean, if available, or just by the year-recall group mean. Whenever we include a potentially imputed measure in the regression, we allow the imputed observations for that measure to have a separate intercept by including the corresponding indicator for whether the measure was imputed.

A.2 Identifying mergers and acquisitions

For our IV estimation strategy, we identify plant acquisitions in the Longitudinal Business Database (LBD) by isolating cases where a plant leaves its previous firm to join another, ongoing firm. The LBD tracks the parent *firm* of each establishment in each year through its firm identifiers. Unlike establishment identifiers, firm identifiers in the LBD are not necessarily longitudinally consistent. As a result, while changes in a plant's firm ownership should be accompanied by changes in its firm identifier, an establishment's firm identifier may also change due to other reasons that are unrelated to a change in ownership, like reorganizations into different legal forms or idiosyncratic changes in EIN filing behavior Chow et al. (2021). To eliminate these alternate scenarios that would also result in a change to an establishment's firm identifiers, we define an *acquisition* as the case when a plant changes firm identifiers to a firm that has had strictly positive employment and payroll during at least one prior year.[23] Correspondingly, our IV sample is the set of all plants in the 2002 Census of Manufacturers that:

1. Switched firm identifiers between 1982 and 2002
2. Experienced an acquisition between 1982 and 2002
3. Is in a multi-unit firm in 2002

The multi-unit firm restriction is not strictly necessary to identify acquisitions, but it helps ensure that the plants in our IV sample are exposed to the decentralization channel and distinct from our placebo sample.

[23]This is similar to the approach from Atalay et al. (2019) to identify mergers using the LBD.

Figure A1 provides an example of an acquired plant that would be included in our IV sample. Circles are establishments — the red (dark) circle is the focal establishment, and grey (light) circles are other establishments that may be in the same firm. Dashed lines indicate establishments and firms whose existence is not required for the definition of our sample, but is also not ruled out by our restrictions. In this example, the focal establishment switched ownership from Firm 1 to Firm 2 between t and t-1. In addition, Firm 2 also existed prior to the ownership change, and is a multi-unit firm. On the other hand, we do not require Firm 1 to survive following the acquisition event, nor do we require Firm 1 to be a multi-unit firm prior to the acquisition.[24]

In contrast, we define our placebo sample as all plants that may have changed ownership but where the decentralization channel would *not* apply. To construct this sample, we include all plants that changed firm identifiers between 1982 and 2002, including both acquired plants and plants that may have undergone a reorganization or an idiosyncratic change in EIN filing behavior. Crucially, we require that the plant is a single-unit firm in 2002 — so the plant and the headquarters coincide — and that the plant never reported a valid decentralization response between 2002 and 2017. These plants are effectively the complement of our IV sample, but we additionally relax our acquisitions-only restriction to ensure that we are including ownership changes and spin-offs that might be a substantial portion of this sample. Figure A2 provides an example of a spun-off plant that would be in our placebo sample. In this example, the focal establishment also switched ownership from Firm 1 to Firm 2, but we do not impose any other restrictions besides that Firm 2 is a single-unit firm following the ownership change.

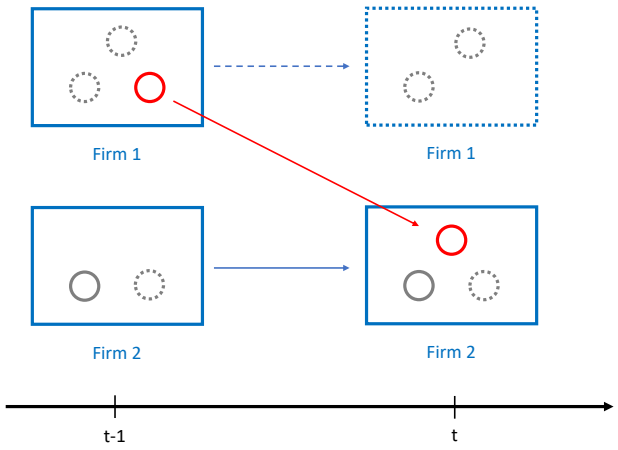


Figure A1: Acquired plant example

[24]Note, however, that by controlling for original establishment size in most of our IV regressions (including our preferred estimates), we are effectively isolating the identifying variation from multi-unit origin firms.

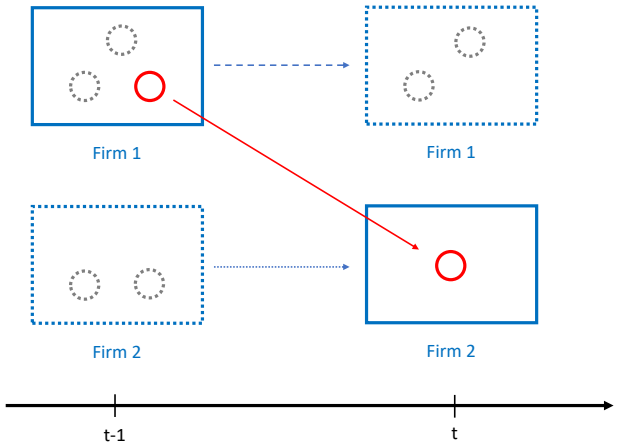


Figure A2: Placebo plant example

A.3 Exports

We use disaggregate establishment revenues by product class to the Longitudinal Firm Trade Transactions (LFTTD) data which contain the universe of export shipments at the firm-product level. To construct our measure, we first match the product categories from LFTTD (ten-digit Harmonized System categories, or HS10) to the 7-digit NAICS product classes contained in the ASM using the Pierce and Schott (2009) concordance. Next, we aggregate exports to the 7-digit NAICS level and calculate the change in exports in each product by year.

A.4 World Management Survey (WMS) International Data

Firm-level Accounting Databases

The WMS sampling frame was based on the Bureau van Dijk (BVD) ORBIS which is composed of the Amadeus dataset for Europe; Icarus for the United States, Oriana for Japan and China. These databases all provide sufficient information on companies to conduct a stratified telephone survey (company name, address, and a size indicator). These databases also typically have accounting information on employment, sales and assets. Apart from size, we did not insist on having accounting information to form the sampling population, however. Amadeus is constructed from a range of sources, primarily the National registries of companies (such as Companies House in the United Kingdom). Icarus is constructed from the Dun & Bradstreet database, which is a private database of over 5 million U.S. trading locations built up from credit records, business telephone directories, and direct research. Oriana is constructed from the Teikoku Database in Japan.The full WMS consists of 34 countries but because we need decentralization data in we are restricted to 18 countries.

The Organizational Survey

In every country the sampling frame for the organization survey was all firms with a manufacturing primary industry code with between 50 and 5,000 employees on average over the most recent three years of data. Interviewers were each given a randomly selected list of firms from the sampling frame. More details are available in Bloom, Sadun and Van Reenen (2012) where we compare the sampling frame with Census demographic data from each country and show that the sample is broadly representative of medium sized manufacturing firms. We also analyzed sample selection - the response rate was 45% and respondents appear random with respect to company performance, although larger firms where slightly more likely to respond. We collected a detailed set of information on the interview process itself (number and type of prior contacts before obtaining the interviews, duration, local time-of-day, date and day-of-the-week), on the manager (gender, seniority, nationality, company and job tenure, internal and external employment experience, and location), and on the interviewer (we can include individual "analyst" fixed effects, time-of-day, and subjective reliability score). We used a subset of these "noise controls" (see text) to help reduce residual variation in some specifications.

In analyzing organizational surveys across countries we also have to be extremely careful to ensure comparability of responses. One step was the team all operated from two large survey rooms in the London School of Economics. Every interviewer also had the same initial three days of interview training, which provided three "calibration" exercises, where the group would all score a role-played interview and then discuss scoring together of each question. This continued throughout the survey, with one calibration exercise every Friday afternoon as part of the weekly group training sessions. Finally, the analysts interviewed firms in multiple countries since they all spoke their native language plus English, so interviewers were able to interview firms from their own country plus the UK and US, enabling us to remove interviewer fixed effects.

The construction of the degree of decentralization measures (from Central Headquarters to Plant Manager) is discussed in the main text. The questions are addressed to the plant manager. We only keep observations where at least two of the four decentralization questions were answered (and we include a control for the number of non-missing questions in the set of noise controls). We drop observations where the plant manager is also the CEO (5% of firms). In cases were the Central Headquarters is on the same site as the plant we interviewed we add a dummy variable to indicate this (one of the noise controls) to reflect potentially greater monitoring. We use the data from the 2006 wave in all cases except when we analyze changes in decentralization as an outcome where we exploit the fact that we ran another wave in 2009 and 2010 for a sub-sample of firms.

As a check of potential survey bias and measurement error we performed repeat interviews on 72 firms in 2006, contacting different managers in different plants at the same firm, using different interviewers. To the extent that our organizational measure is truly picking up company-wide practices these two scores should be correlated, while to the extent the measure is driven by noise the measures should be independent. The correlation of the first interview against the second interviews was 0.513 (p-value of 0.000), with no obvious (or statistically significant) relationship between the degree of measurement error and the decentralization score. That is to say, firms that reported very low or high decentralization scores in one plant appeared to be genuinely very centralized or decentralized in their other plants, rather than extreme draws of sampling measurement error.

Firm-level variables

Our firm accounting data on sales, employment, capital (fixed assets), profits and intermediate inputs came from BVD ORBIS. Whether the variable is reported depends on the accounting standards in different countries. Sales are deflated by a three digit industry producer price index. BVD has extensive information on ownership structure, so we can use this to identify whether the firm was part of a multinational enterprise. We also asked specific questions on the multinational status of the firm (whether it owned plants aboard and the country where the parent company is headquartered) to be able to distinguish domestic multinationals from foreign multinationals.

We collected many other variables through our survey including information on plant size, skills, organization, etc. as described in the main text. We also collected management practices data in the survey. These were scored following the methodology of Bloom and Van Reenen (2007), with practices grouped into four areas: operations (three practices), monitoring (five practices), targets (five practices), and incentives (five practices). The shop-floor operations section focuses on the introduction of lean manufacturing techniques, the documentation of processes improvements, and the rationale behind introductions of improvements. The monitoring section focuses on the tracking of performance of individuals, reviewing performance, and consequence management. The targets section examines the type of targets, the realism of the targets, the transparency of targets, and the range and interconnection of targets. Finally, the incentives section includes promotion criteria, pay and bonuses, and fixing or firing bad performers, where best practice is deemed the approach that gives strong rewards for those with both ability and effort. Our management measure uses the unweighted average of the z-scores of all 18 dimensions.

B OLS and IV bias under mean-reverting measurement error

We consider the impact of non-classical measurement error on our OLS and IV estimates of decentralization on growth. To frame intuition and make our main point, we first consider the case of a binary endogenous variable. We then generalize to an environment close to our empirical exercise which averages across multiple discrete variables. We provide both analytical and simulation-based results.

B.1 Binary treatment

Consider the following simplified model:

$$y_i = \beta D_i + \theta W_i + \epsilon_i,$$

where $D_i \in \{0,1\}$ is the true binary variable of interest and $\hat{D}_i = \pi(1-D_i) + (1-\pi)D_i = D_i + \pi(1-2D_i)$ is the observed variable that is "wrong" with probability π. W_i is a confounding variable that may be correlated with D_i. Under this setup,

$$\begin{aligned} E\left[\hat{D}|D=0\right] &= \pi > 0 \\ E\left[\hat{D}|D=1\right] &= 1-\pi < 1 \end{aligned}$$

so clearly $E\left[\hat{D}|D\right] \neq D$, and measurement error cannot be classical.

How might this affect our results? Consider as a simple baseline example, a case where

$$D_i = 1(Z_i + W_i + \eta_i) > 0 \quad (7)$$

$$y_i = D_i - W_i + \epsilon_i \quad (8)$$

So Z_i is a valid instrument for binary D_i, and without controlling for W_i, there would also be negative omitted variables bias. We simulate the empirical model in Equations 7 and 8 by drawing 10,000 sets of $W_i, \epsilon_i, Z_i, \eta_i$ from i.i.d. standard normal distributions. As a starting point, we assume that all random variables are drawn from a standard normal distribution, and we consider the impact of measurement error on both the OLS and IV regressions (with and without controlling for the confounder W_i).

In Panel A of Table B1, we assume that there is a potentially omitted variable W, but that decentralization is measured without error (i.e. $\pi = 0$). The long OLS in column (3), the short IV in column (3), and the long IV in column (4) are both able to recover the true parameter value of 1. On the other hand, the short OLS regression in column (1) is biased downwards substantially due to the omitted variable.

On the other hand, in Panel B of Table B1, we assume that we only observe $\hat{D}_i$ that is correct 75% of the time, and incorrect other 25% of the time (i.e. $\pi = 0.25$). Even when we are able to control for the confounding variable W_i, the OLS regression is attenuated by almost 60% in column (2). On the other hand, the IV estimate is around 90% *too* large in column (3), and controlling for the confounding variable in column (4) does not change our results. We can unpack the IV results

Table B1: OLS and IV estimates under simulated measurement error

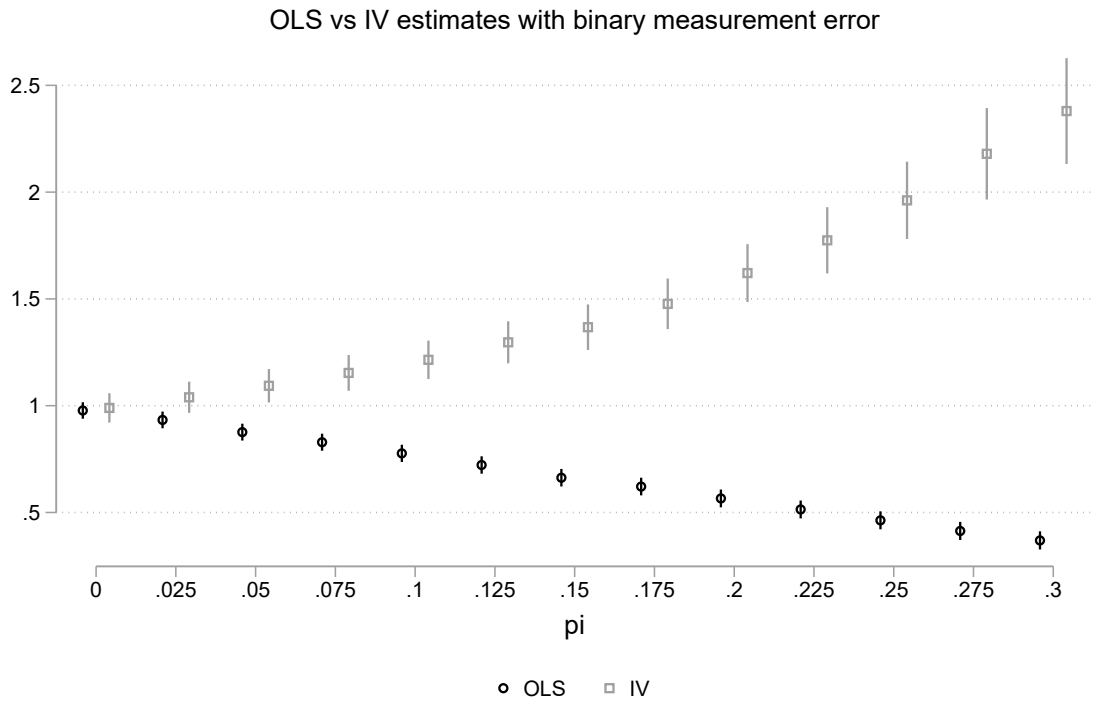
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: True model ($\pi = 0$):					
d	0.0713	0.978	1.010	0.988		
	(0.0264)	(0.0219)	(0.0602)	(0.0423)		
w		-0.991		-0.993		
		(0.0109)		(0.0137)		
z					0.236	0.233
					(0.0135)	(0.00357)
	Panel B: Noisy model ($\pi = 0.25$):					
d	0.0285	0.408	1.936	1.893		
	(0.0264)	(0.0220)	(0.134)	(0.105)		
w		-0.814		-0.988		
		(0.0108)		(0.0177)		
z					0.236	0.122
					(0.0135)	(0.00470)
N	10000	10000	10000	10000	10000	
Model:	OLS (short)	OLS (long)	IV (short)	IV (long)	Reduced form	First stage

Note: Estimates from simulated model of Equations 7 and 8. $W_i, \epsilon_i, Z_i, \eta_i$ are all from i.i.d. standard normal distributions. The dependent variable in columns 1-4 is the model outcome y_i, and the dependent variable in column 5 is the (potentially noisily observed) treatment D_i. Heteroskedasticity-robust standard errors in parentheses.

by by comparing the reduced form and first stage estimates for the two sets of regressions in columns (5) and (6). As expected, the reduced form is identical across Panels A and B, since it does not depend on the endogenous regressor D_i. On the other hand, since measurement error compresses the range of $\hat{D}_i$, it also leads the noisy first stage to be too small in Panel B, which in turn leads the IV estimate to be too large. As a result, the IV estimate is around 4.6 times larger than the OLS, even when we are controlling for the omitted variable W_i.

How do these estimates depend on the degree of measurement error π? In our simple example, we can get a sense of this by abstracting from the omitted variable W_i, and calculating the OLS and IV estimates for different hypothetical levels of π in Figure [1]. We see that even mild degrees of measurement error can inflate the ratio between the IV and OLS estimates. For example, $\pi = .1$ corresponds to a ratio of 1.56. This ratio increases quickly with π as the IV estimates become larger (and noisier) and the OLS estimates become more attenuated.

Figure B3: OLS vs IV estimates by π



B.2 General model

In general, for a noisily measured treatment

$$\hat{D}_i = D_i + \alpha + \rho\left(D_i - \bar{D}\right) + \epsilon_i, \tag{9}$$

where ρ captures the bias in the measurement error and ϵ captures its variance, the probability limit of the OLS and IV estimates are

$$E\left[\hat{\beta}_{OLS}\right] = \beta\frac{(1+\rho)\,\sigma_D^2}{(1+\rho)^2\,\sigma_D^2+\sigma_\epsilon^2} = \underbrace{\frac{\beta}{1+\rho}}_{\text{bias}}\underbrace{\left(1+\frac{\sigma_\epsilon^2}{(1+\rho)^2\,\sigma_D^2}\right)^{-1}}_{\text{attenuation}} \tag{10}$$

$$E\left[\hat{\beta}_{IV}\right] = \frac{\beta}{1+\rho}. \tag{11}$$

It then follows that

$$\frac{E\left[\hat{\beta}_{IV}\right]}{E\left[\hat{\beta}_{OLS}\right]} = 1+\frac{\sigma_\epsilon^2}{(1+\rho)^2\,\sigma_D^2}. \tag{12}$$

This implies that for a given model of measurement error in $\hat{D}_i$ and a given set of OLS and IV estimates, it would be sufficient to only know ρ to back out the true effect β from the biased estimate $\hat{\beta}_{IV}$. However, it is not possible to separately identify ρ from $\sigma_\epsilon^2/\sigma_D^2$ simply by comparing $\hat{\beta}_{IV}$ and $\hat{\beta}_{OLS}$, unless we impose additional restrictions on $\hat{D}_i$ such that ρ and $\sigma_\epsilon^2/\sigma_D^2$ are determined by a

single parameter (as in the binary case).

We can examine several hypothetical models for measurement error in our decentralization measure, and evaluate what each model implies for the bias in $\hat{\beta}_{IV}$ and the true value of $\hat{\beta}$. The key to each model is that there exists a single parameter θ which determines ρ and $\sigma_\epsilon^2/\sigma_D^2$, and which in turn, allows us to define a function $G(\theta) = \frac{\sigma_\epsilon^2}{\sigma_D^2}/(1+\rho)^2$. We can then, in theory, combine our OLS and IV estimates in a GMM framework to back out (θ, β) — the key sufficient statistic for the non-classical measurement error (under our assumed model) and the true coefficient β.

For each model of measurement error in this section, we assume that the true decentralization component $D_{ij} = 1(Z_i + \eta_i) > c$, where Z_i, η_i are i.i.d. standard normal random variables, and c is determined by the assumed model and measurement error parameter π such that the *noisy* measure of decentralization matches the empirical sample mean. For each model, we numerically approximate its conditional mean and variance by drawing 1,000,000 observations. We calculate the analytical bias in OLS and IV using the numerical estimates for ρ and $\sigma_\epsilon^2/\sigma_D^2$ and e

B.2.1 Example 1: index over binary variables with i.i.d. measurement error

As a natural extension of our motivating example, we can assume that each decentralization measure is an average of six binary variables D_{ij}, each of which has a probability π that it is mismeasured. So

$$\begin{aligned} D_i &= \sum_{j=1}^{6} D_{ij}/6 \\ \hat{D}_{ij} &= \pi(1 - D_{ij}) + (1-\pi) D_{ij}. \end{aligned}$$

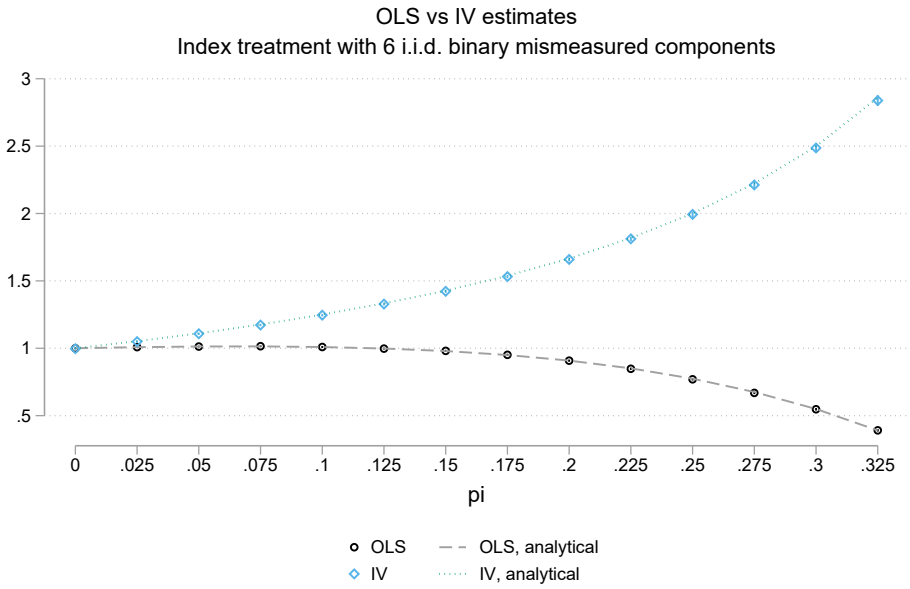
Since both expressions are linear, it then follows that

$$E\left[\hat{D}_i | D_i\right] = \pi + (1 - 2\pi) D_i.$$

This implies that the bias between $\hat{\beta}_{IV}$ and β is actually identical between the multi-response index case and the simple binary case, since we can iterate over expectations. However, the variance of the measurement error would be smaller in the index case than in the single binary variable case, so the attenuation of $\hat{\beta}_{OLS}$ would also be smaller.

In practice, with six independently mismeasured components, the bias component and the attenuation component nearly offset each other. This appears to be a genuine coincidence rather than some special property of binary measurement error — as we fix π and increase the number of components from 1 to 6, the bias term remains constant while the degree of attenuation decreases. However, the rate of attenuation changes in such a way such that the range of values π for which the overall degree of bias for $\hat{\beta}_{OLS}$ is small actually expands substantially.

Figure B4: OLS vs IV estimates by π



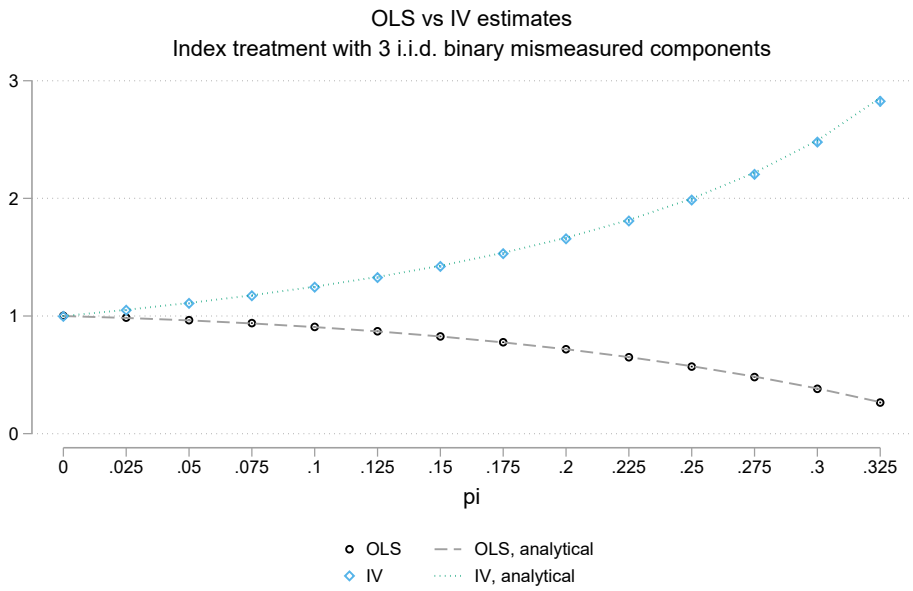
B.2.2 Example 2: clustered binary variables

To approximate the fact that the individual decentralization measures tend to fall within one of three "clusters," we can also explicitly assume that rather than averaging over six independent measures, we are actually aggregating over three factors: capital, labor, and product, with each factor potentially appearing in multiple questions. If we abstract away from the different weights that each factor might receive, then the treatment of this model more or less follows from the 6-measure i.i.d. example. So,

$$\begin{aligned} D_i &= \sum_{j=1}^{3} D_{ij}/3 \\ \hat{D}_{ij} &= \pi\left(1 - D_{ij}\right) + \left(1 - \pi\right) D_{ij}. \end{aligned}$$

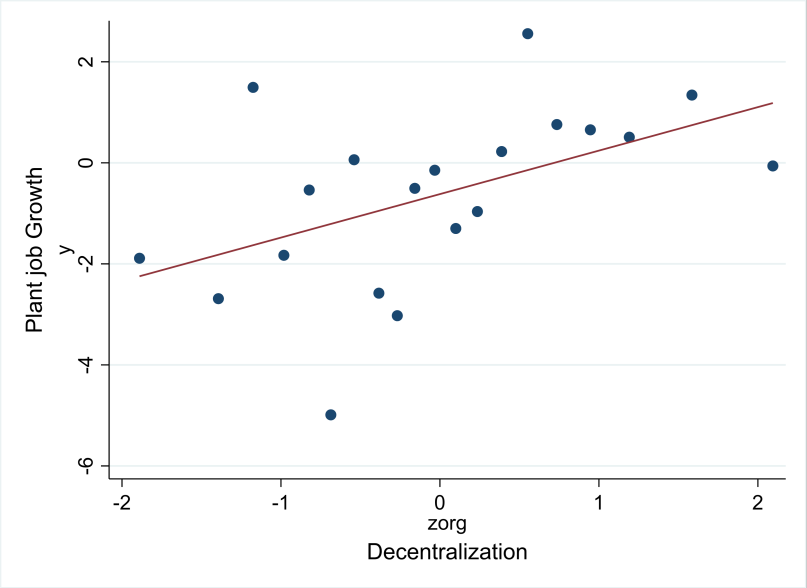
Notice, however, that in this case the attenuation factor is more powerful, which leads to relatively more bias for $\hat{\beta}_{OLS}$.

Figure B5: OLS vs IV estimates by π



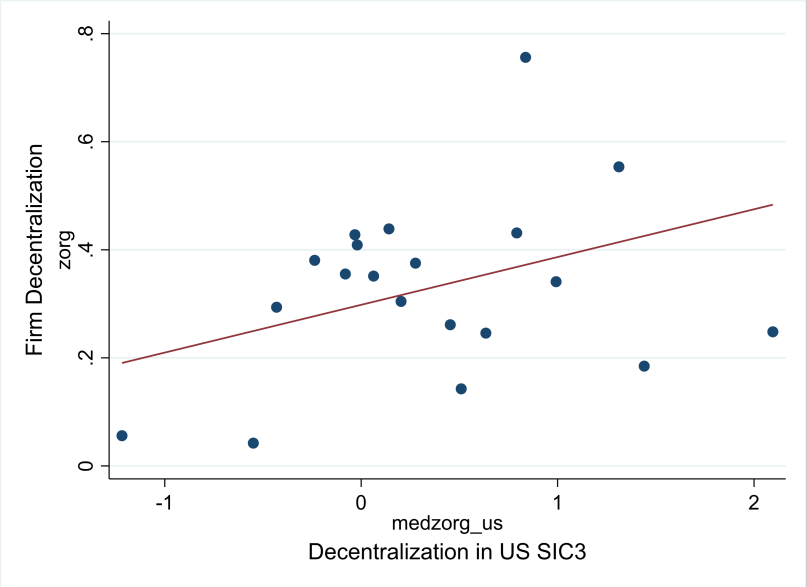
C Additional Tables and Figures

Figure C1: WMS plant employment growth is increasing in initial decentralization



Note: Binscatter of average annualized plant-level employment growth from WMS (in percentage points) as a function of initial average decentralization. WMS data pooled across all available years.

Figure C2: WMS industry-level decentralization is correlated between the US and other Anglo-Saxon countries



Note: Binscatter of average U.S. plant decentralization by industry as a function of average plant decentralization by industry in other Anglo-Saxon countries (Australia, Canada, New Zealand, and the UK). WMS data pooled across all available years.

Table C1: IV First Stage for productivity sub-sample

Dep. Variable	Average Decentralization			
	(1)	(2)	(3)	(4)
Orig. firm size	-0.00778***	-0.00982***	-0.00989***	-0.00925***
	(0.00130)	(0.00142)	(0.00142)	(0.00149)
# Est.	3700	3700	3700	3700
# Firms	1400	1400	1400	1400
Controls:				
Orig. est. size		Y	Y	Y
Management			Y	Y
Baseline firm chars.				Y

Note: OLS coefficients with standard errors in parentheses (clustered at the firm level). Original firm size is the log employment of the plant's owner firm in 1982. The sample is identical to the sample in Table 6. The dependent variable is the plant's average decentralization score from all waves of the MOPS. "Original establishment size" controls allow for differences by the plant's log(employment) and log(wagebill) from the 1982 LBD. "Management" controls allow for differences by the plant's structured management score (constructed as in Bloom et al. (2019) from 16 question) questions from all waves of the MOPS. "Baseline firm characteristics" controls allow for differences by the plant's ownership 2002 firm's log(employment), log(wagebill), log(establishment count in firm), and log(firm age) from the 2002 LBD, as well as the plant's share of total firm output from the 2002 CMF. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

Table C2: IV Results using additional characteristics of donor firms (payroll and number of plants)

	Long Differences, 2002-17				
	(1)	(2)	(3)	(4)	(5)
Decentralization	15.59***	12.01**	2.075	15.97***	13.51***
	(5.308)	(5.743)	(4.095)	(3.736)	(4.693)
# Est.	4400	4400	4400	3700	3700
# Firms	1600	1600	1600	1400	1400
First stage F	18.1	18.1	18.1	16.4	16.4
Dep variable	Sales	VA	Emp	Sales/Emp	VA/Emp
J-test p	0.662	0.459	0.729	0.567	0.812
Controls:					
Orig. est. size	Y	Y	Y	Y	Y

Note: Two-stage least-squares coefficients with standard errors in parentheses (clustered at the firm level). We instrument for the plant's average decentralization status from all waves of the MOPS using the log(employment) of the plant's owner firm ("donor firm") in 1982 as in main IV table, but now also add 1982 donor firm log(payroll) and log (number of establishments). The sample in the first three columns matches the sample in Table 5, and the sample in the last two columns matches the sample in Table 6. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

Table C3: IV reduced form

	Long Differences, 2002-17					
	(1)	(2)	(3)	(4)	(5)	(6)
Orig. firm size	-0.0837**	-0.0609	-0.0191	-0.122***	-0.0849***	0.00559**
	(0.0384)	(0.0392)	(0.0319)	(0.0244)	(0.0320)	(0.00237)
# Est.	12500	12500	12500	6200	6200	12500
# Firms	4300	4300	4300	2600	2600	4300
Outcome	Sales	VA	Emp	Sales/Emp	VA/Emp	Exit
Controls:						
Orig. est. size	Y	Y	Y	Y	Y	Y

Note: OLS coefficients with standard errors in parentheses (clustered at the firm level). Original firm size is the log employment of the plant's owner firm in 1982. Note that this sample is not conditional on observing decentralization, so we can expand our IV samples to include all plants that are not covered by any wave of the MOPS. These are the estimated reduced forms implied by our IV specification. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

Table C4: Decentralized plants are less likely to exit

Dep. Variable	Plant Exit by 2017					
	(1)	(2)	(3)	(4)	(5)	(6)
Decentralization	-0.0708***	-0.0616***	-0.345	-0.243	-0.0362	-0.0300
	(0.0225)	(0.0226)	(0.245)	(0.257)	(0.0350)	(0.0349)
Export shock					0.0296	0.0270
					(0.0572)	(0.0566)
Shock × decen.					-0.125	-0.116
					(0.107)	(0.106)
# Est.	7300	7300	4400	4400	7300	7300
# Firms	3400	3400	1600	1600	3400	3400
First stage F			53.9	44.2		
Controls:						
Baseline est. size	Y	Y	Y	Y	Y	Y
Management		Y		Y		Y
Baseline firm chars.		Y		Y		Y

Note: OLS coefficients with standard errors in parentheses (clustered at the firm level). The dependent variable is a binary indicator for whether the plant had positive payroll in 2017 (as measured by the LBD). The sample and controls in each pair of columns correspond to Tables 2, 5, and 11, respectively. Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.

Table C5: Instrumenting firm decentralization with US decentralization in the same industry (WMS)

Method	(1) OLS	(2) First Stage	(3) 2SLS
		A. Sales	
Decentralization	0.997**		3.865*
	(0.384)		(2.037)
US industry Decentralization		0.198** (0.075)	
Observations	1,576	1,576	1,576
		B. Employment	
Decentralization	0.743**		5.408**
	(0.313)		(2.264)
US industry Decentralization		0.186** (0.071)	
Observations	1,591	1,591	1,591
		C. Sales per worker	
Decentralization	0.058		-0.834
	(0.317)		(1.934)
US industry Decentralization		0.190** (0.074)	
Observations	1,560	1,560	1,560

Note: Regressions based on WMS firm-level data (unbalanced panel) from "Anglo-Saxon" countries (Australia, Canada, UK, and New Zealand). Decentralization (z-score) is measured in 2006 and 2009 and outcomes are subsequent growth in sales, employment, or sales/employment (annualized change in log values over subsequent 6 years). Standard errors are clustered by three-digit industry. Decentralization in the US is the median level of decentralization in the US in the same three-digit industry the firm operates in. Controls are year dummies, country dummies, 2-digit industry dummies, and (in the same year as decentralization) management score, union density, the fraction of employees with a college degree, whether the establishment belongs to a domestic multinational, whether the establishment belongs to a foreign multinational, log(firm size) and log(plant employment). Stars indicate the level of statistical significance: * 10%, ** 5%, and *** 1%.