Data in Action: Data-Driven Decision-Making in U.S. Manufacturing[*]

PRELIMINARY– DO NOT CITE OR DISTRIBUTE

Erik Brynjolfsson
MIT Sloan School of Management
100 Main Street, E62-414
Cambridge, MA 02142
erikb@mit.edu

Kristina McElheran
University of Toronto
105 St. George Street
#7033
Toronto, ON M5S 3E6
kristina.mcelheran@rotman.utoronto.ca

This Draft: February 22, 2015

Abstract

We investigate the adoption and performance effects of data-driven decision-making (DDD) in U.S. firms. Our study relies on a novel survey of managerial practices conducted by the U.S. Census Bureau combined with other detailed data on information technology (IT) investment and firm performance in the manufacturing sector. Using a differences-in-differences approach, we find that plants adopting DDD exhibit an average increase of 3% in value-added between 2005 and 2010. For the average plant in our sample, this increase is on par with the productivity benefit from investing an additional $5 million in IT capital, or $60 thousand per employee over the five-year period. Because IT *use* is managerially and statistically distinct from IT *investment*, we take this as evidence that it can be just as important for firm performance. Moreover, DDD displays important complementarities with other organizational choices such as employee education and the allocation of decision rights. Furthermore, the variation in the timing and distribution of DDD effects suggests there are complex and costly organizational adjustments as firms adapt to new technological possibilities.

[*] Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. All errors are our own. We think the MIT Initiative on the Digital Economy for generous funding for this research.

1. INTRODUCTION

Economists and management scholars have long been interested in uncovering the drivers of persistent performance differences among firms that appear otherwise similar in terms of size, industry focus, and investments in labor, capital, and other inputs (e.g., Jensen and McGuckin 1996, Syverson 2011, Henderson and Gibbons forthcoming). A significant body of prior evidence points to information technology (IT) as a potential differentiator among firms (e.g., Brynjolfsson and Hitt 1995, Dunne et al. 2004, Syverson 2011). Yet more-recent work on the role of IT in firm productivity emphasizes important differences among organizations in how they apply generic IT capital in the production process and the extent to which IT is actually integrated into essential firm activities (Bresnahan and Greenstein 1996, Aral and Weill 2007, McElheran forthcoming).

Given the significant increase in data storage and processing capabilities and recent widespread attention among practitioners to "big data" (Gartner, 2014; Brynjolfsson, Hitt and Kim 2011), this paper focuses on one potentially important application of IT: collecting and using data to guide managerial decision-making in firms. Recent work has focused on management practices as a potentially critical component of heterogeneous firm performance (Bloom et al. 2014). Data -driven decision-making, in particular, is one practice that has been associated with superior firm performance in a modest sample of large public firms (Brynjolfsson et al. 2011).

Our paper uses data recently collected by the U.S. Census Bureau to examine these phenomena in more detail, in a representative sample of firms, and with unprecedented visibility into organizational design choices and complementary activities at the firm. It extends our current understanding by providing the first evidence that applying IT to the particular practice of data-driven decision-making is correlated with better firm performance, even controlling for both generic IT investment and adoption of more structured management practices. In robust cross-sectional analyses as well as our preferred differences-in-differences specification, we find that increased levels of data-driven decision-making are associated with increases in value-added of over 3%, as well as higher levels of total factor productivity

(TFP) and other performance measures (*pending disclosure*). For the typical plant in our sample, the magnitude of this effect is comparable to investing an additional $5 million in IT capital, or $60 thousand per employee over the five-year period. This effect is robust to controlling for IT investment in a number of ways, suggesting that the *use* of the IT is just as important as the level of investment. However, this productivity differential is not uniform across firms: establishments that belong to multi-unit firms, establishments that are larger or have greater capital intensity, those that are older, and those with a more complex operating environment show a lower correlation of productivity with DDD, perhaps because they have been less able to make complementarity organizational changes or because they require more time to do so. This pattern suggests a complex and costly organizational adjustment on multiple margins as firms adapt to new technological possibilities. Consistent with this interpretation, DDD displays interesting complementarities with other organizational choices such as the level of IT investment, the education of the plant's workforce, and the allocation of decision-making within the firm.

In Section 2 we describe the motivation for this study and link it to prior literature. In Section 3 we introduce the data, our analysis sample, and how we measure data-driven decision-making. Section 4 explores which firms adopt DDD. Section 5 examines the relationship between DDD and firm performance. Section 6 presents our conclusions.

2. CONCEPTUAL MOTIVATION & PRIOR LITERATURE

Because organizations can be thought of as "information processors" (Galbraith, 1974), it would be surprising if large changes in the costs of data collection, storage and analysis made possible by computers did not lead to large changes in other processes and systems in organizations. However, changing human systems, many of which are not well-understood, is costly and time-consuming. Thus, the reducing the cost of digital data can be expected to have uneven effects on firms as the complementary changes are invented.

In particular, digital data collection codifies information, which makes it more explicit and less tacit. When firms invest in both the IT infrastructure to collect data and the managerial decision-making process required to select which data to collect (and take a stand on why it is important), firms typically go through a lengthy process of "learning what they know" by consulting with employees from throughout the firm and discovering what sorts of knowledge reside in scattered locations throughout the firm. Importantly, this process is useful for capturing tacit knowledge that employees gain through less formalized channels, and codifies and centralizes it in accessible databases. If firms, like people, "can know more than they can tell" (Polanyi 1964), the steps necessary to implement data-driven decision-making help increase access to more objective information.

The economic value of this information is zero unless it leads to a change in action. Yet, it's not always easy to act on this new codified and accessible information. In complex modern enterprises, this typically requires not only changes in decision rights and processes, but also complementary changes on multiple dimensions within the organization (performance pay, promotion, hiring, measurement, etc.) and even culture. This is inevitably costly and time-consuming, and often subject to mistakes. On the other hand, once implemented, these changes can be difficult for competitors or quickly imitate, providing a margin of competitive advantage, and measurable differences in performance.

3. DATA

This study is one of the first to take advantage of new data on management practices collected by the U.S. Census Bureau in 2011 and released for non-internal use in 2014 (for more details see Bloom et al., 2014). The Management and Organizational Practices Survey (MOPS) was included as a supplement to the Annual Survey of Manufactures (ASM), which targets roughly 50,000 of the over 300,000 establishments in the U.S. manufacturing sector. It is generated from five-year rotating sample frames that begin in years ending with "4" or "9" and end in years ending with "8" or "3." The ASM sample is constructed with the intention of generating representative annual coverage of the manufacturing sector and, as a result, somewhat over-weights large plants.[1]

[1] Plants with over 1,000 employees are sampled with certainty; the likelihood of sampling is lower but increasing with size for all plants below this threshold. We use the ASM sampling weights where appropriate to provide statistics that are relatively more representative of the entire population of U.S. manufacturing establishments. The certainty sample accounted for approximately 67% of the total value shipped in the entire U.S. manufacturing sector

The survey was sent by mail and electronically to the ASM respondent for each establishment, which was typically the accounting, human-resource, or general plant manager. Non-respondents were given up to three follow-up telephone calls if no response had been received within three months. The ultimate response rate was 78% (Bloom et al., 2014), which is extremely high for surveys of this nature.

The survey comprises 36 multiple-choice questions split into three sections. The first section, labeled "management practices" and based on work by Bloom and Van Reenen (2007) has 16 questions focused primarily on monitoring, communication, and incentives at the plant. Examples include the collection and communication of key performance indicators (e.g., production targets, costs, quality, inventory, absenteeism, and on-time deliveries), the speed with which under-performing employees are reassigned or dismissed, and the basis for promotion and performance bonuses. The second section, labeled "organization," has 13 questions focused on decision-making within the firm. Inspired by work such as Bresnahan, Brynjolfsson and Hitt (2002) and Bloom, Sadun and Van Reenen (2012), it has six questions about different types of decisions ranging from human resources management to product strategy and advertising, and whether they are made at the local plant, at headquarters, or both. In addition, based on prior work by Brynjolfsson, Hitt and Kim (2011), there are two questions on the availability and use of data to support decision-making at the establishment. The final section of the survey captures "background" characteristics of the respondent and establishment, primarily the respondent's position and tenure at the plant as well as the number and education of the plant's employees and their union status. The full questionnaire is available at http://bhs.econ.census.gov/bhs/mops/form.html. Bloom et al. (2014) explores the distribution and performance implications of the practices covered in the first section. While we leverage many different aspects of the survey in this paper, we take as our point of departure the data-related questions from the second section, as well as a couple of questions directly related to data collection and use from the first section (see below).

---

in 2007 (see http://www.census.gov/manufacturing/asm/how_the_data_are_collected/index.html).

Sample Boundaries. Our analysis requires that we restrict our attention to establishments that have positive value added, positive employment, and positive imputed capital in the ASM.[2] In order to keep our sample stable across specifications,[3] we further restrict our analysis to records with complete responses to the data-driven decision-making questions, headquarters status, and a critical mass of the management practices questions.[4]

Panel Data Sample. Although the survey only took place once, in conjunction with the 2010 ASM, all of the questions asked respondents to report what they recall the state of practice and other features of the establishment to have been in 2005. The ASM itself provides data on various plant characteristics for both 2005 and 2010. Thus, we use the difference between what is reported for 2005 and for 2010 to employ a differences-in-differences research design. Of course, the validity of this approach depends in part on the quality of recall. As one gauge of quality, we explored similarities between the 2005 recall questions regarding plant-level employment and actual IRS records for that year and find the differences to be negligible (*pending disclosure*). We include a measure of the discrepancy between the 2005 MOPS and 2005 ASM employment numbers as a "noise" control in many specifications. We also control for the tenure of the respondent at the plant, although the impact of these additional controls is trivial (see Table 6).

Linking the 2010 MOPS sample to the 2005 ASM reduces the size of our analysis sample from roughly 34,000 observations in the 2010 complete-information cross section to a balanced panel of

[2] This is to make the standard productivity calculations possible and to exclude low-quality records that may introduce systematic biases to the estimation. To meet these requirements automatically requires a successful match between the MOPS and the ASM and that the establishment be flagged as worthy of tabulation in the national statistics. Another technical condition for the panel analysis (and to get controls such as age) requires the establishment to have a valid linkage to the Longitudinal Business Database (LBD).

[3] A stable analysis sample is essential to avoiding inadvertent disclosure of proprietary firm information and is requested by the U.S. Census Bureau when using the non-public data.

[4] Specifically, we require that firms respond to questions 1,2,5,6,8,9,11,13,14,15,&16. Omitted from this required list are questions that we worry have a high potential to be confounded with firm performance, such as how difficult it was to meet performance targets, or what percentage of employees received performance bonuses (which would depend on sufficiently good performance to justify them). We explore but do not require responses to questions 3 & 4 on the frequency of with which key performance indicators are reviewed. We exclude these questions from our core analysis because they are highly correlated with our other measures, have a relatively higher percentage of missing data, and offer no additional insights while restricting the sample size.

roughly 18,000 observations.[5] The loss of observations comes primarily from the change in ASM sample frame in 2008 and from the requirement that plants be at least five years old; a small number also lack complete information for 2005 and get dropped as a result. Due to these restrictions, our balanced sample is slightly biased towards plants that are larger and slightly more productive compared to the entire ASM mail-out sample. Comparisons between the 2010 balanced sample, the 2010 complete-information cross section, and the 2010 ASM for key variables are provided in the Appendix (see Table A1). Where appropriate, we leverage the ASM sampling weights to generate statistics that are relatively more representative of the entire population of manufacturing establishments. That said, we interpret our core results as being principally informative about the population of larger manufacturing establishments in the U.S.

Subject to these limitations, examining these practices in a panel setting is a significant advance, as it expands the scope for addressing unobserved heterogeneity among establishments that could bias estimates of the return to adopting DDD. Our results suggest that this is a valid concern (see Section 5). Finally, our sample is representative of the diversity of activities comprising U.S. manufacturing, covering all industries from food to furniture manufacturing (86 industries at the 4-digit NAICS level of aggregation).

Data-Driven Decision-Making (DDD). The key questions about how firms use data to support managerial decision-making are questions 27 and 28 of the MOPS. Respondents are asked to choose a value on a 5-point Likert scale according to "what best describes the **availability of data** to support decision-making at this establishment", and "what best describes the **use of data** to support decision-making at this establishment." Empirically, they are highly correlated and we combine information from both to reduce measurement error and to help separate out firms that are closer to the frontier of practice.

In addition, question 2 of the MOPS asks about the number of key performance indicators (KPIs) tracked at the establishment. Respondents are primed with the following example: "Metrics concerning production, cost, waste, quality, inventory, energy, absenteeism and deliveries on time." Our prior, which

[5] Exact records counts are suppressed in the interest of disclosure avoidance.

has been corroborated by qualitative interviews independent of the Census Bureau data collection process, was that the number of identified and tracked performance measures is an essential measure of the breadth and/or intensity of data gathering at the establishment. Thus, we combine the core DDD questions with this objective measure of whether plants collect data on 10 or more KPIs.

Furthermore, having appropriate targets against which to compare real-time or historical data ought to facilitate decision-making on whether the production system is performing appropriately, identifying where it is not (and how badly), and formulating appropriate actions. Again, this assumption was corroborated in independent interviews with plant managers. Question 6 of the MOPS asks about the presence and time frame of production targets (short-term, long-term, or combined). We take the combined approach to target setting as a more nuanced managerial understanding of the dimensions of performance that need to be monitored and controlled.

Thus, for our core investigation, we create a combined indicator for being at the frontier of data availability (question 27) and use of data in decision-making (question 28), extensive use of key performance indicators (question 2), and employing the most nuanced approach to target-setting (question 6). We call this "data-driven decision-making" (DDD) throughout the rest of this paper. A summary of the adoption of these different combinations of data-related managerial practices is provided in Table 2a. The first two columns show how adoption changed over time in the balanced panel. The third shows results for the 2010 complete-information cross section, which has more small and young establishments and slightly less DDD adoption, on average.

The first two rows of Table 2a show that, by 2010, the perceived availability and use of data in U.S. manufacturing was quite widespread. Using the ASM sampling weights to extrapolate a bit better to the population of manufacturing plants, nearly 60% of establishments report being relatively intensive **both** in collection **and** use of data (third row of Table 2a). However, when this is combined with the arguably more objective measure of how many key performance indicators are tracked at the plant, only 34% are also tracking 10 or more dimensions of performance by 2010. Only 25% report that they furthermore are comparing this information to short-term and long-term targets. These latter conditions help identify the population of establishments that are objectively closer to the frontier of data-related management practices – an important distinction given the popularity of all things "data" over the past decade and more. A rapid growth of DDD is apparent in the data, only 9% of plants reported intensive DDD according to our definition in 2005 (although 43% nevertheless report that they had either a "great deal of data" or "all the data we need" to support decision-making).[6] It is worth noting that plants that are part of multi-unit firms appear more likely to adopt DDD in both periods (see Table 2b).

Relying on this combination of practices to identify DDD is empirically justified by a polychoric principal factor analysis (see Table 3). Applying this technique, appropriate for factor analysis of discrete variables, to these four dimensions of practice reports a single factor with an eigenvalue of 2.28 accounting for 57% of the variance in the balanced sample in 2010. An oblique promax rotation confirms a single factor; similar results also obtain for principal-component factor analysis (available upon request). Nevertheless, we further explore the robustness of our findings to a wide range of different definitions of DDD in Table 8 (see Section 5)

Management Practices. A key concern for identifying the relationship between DDD and higher productivity is the possibility that DDD may merely proxy for "better management" at the establishment. Bloom et al. (2014) show robust positive correlations between the adoption of more "structured" management and performance measures identical to the ones we study here. If DDD and structured management are correlated, then measures of DDD adoption in isolation could simply pick up the effect of these other management practices. To address this concern, we construct indices of structured management that are similar to those used by Bloom et al. (2013) but that *omit the data-related measures discussed above*. Also in contrast to this other work, we further eliminate measures that we believe may be confounded with performance, such as the likelihood of reassigning or dismissing underperforming

[6] Adding restrictions to how DDD is defined in terms of the frequency of data review (questions 3 and 4) or how KPIs are displayed throughout the firm (question 5) further reduces the number of firms at the frontier – however, we found that this started to lead to problems of small cells (and hence risked violating disclosure avoidance guidelines) and no significantly different estimation results.

workers quickly (which may tend to happen more when the firm, in general, is performing badly) and whether performance bonuses were paid in the prior year (again, which might depend on whether performance was sufficiently good to warrant bonus pay in the first place). We create a Z-score indexing the remaining non-DDD, non-performance questions from the first section of the survey as our main proxy for structured management.[7] Based on explorations of the broader survey instrument using discrete principal factor analysis – and in an attempt to use as much of the information we have to hand – we separately construct and include an index for how quickly an underperforming non-manager or manager is dismissed or re-assigned (questions 15 and 16).[8]

Complementary Practices and Investments. If firms will make adjustments on multiple fronts to take advantage of new technology, this raises the interesting question of what, specifically, they are changing at the same time. An obvious candidate is the level of IT investment. We calculate IT capital stock using the ASM and CMF questions on hardware expenditure dating back to 2002 and software expenditure questions dating back to 2006. Specifically, we use the Bureau of Economic Analysis (BEA) deflators and a perpetual inventory approach, combining hardware and software investment, imputing values for years in which they are missing[9], and depreciating at the rate of 35% per year.

Human capital is another obvious potential complement to DDD. Questions 34 and 35 ask about the percentage of managers and non-managers, respectively, with bachelor's degrees. We use a combined measure to a) control for labor quality in our performance regressions and b) explore whether DDD is complementary to having a more educated workforce.

---

[7] This consists of normalizing the response of all of the following questions to the 0-1 interval and summing them to create the composite management score: 1, 5, 8, 9, 11, 13, & 14. These questions cover how the firm reacted to an exception in its production process, whether and where display boards showed output and other key performance indicators, who was aware of production targets at the plant, what the basis (if any) was for performance bonuses for managers and non-managers, and the basis for promotion of managers and non-managers (performance and ability versus other factors such as tenure or family connections).

[8] These questions are negatively correlated with our productivity measure and often show up as a separate factor, justifying a separate treatment in our analysis.

[9] For plant-years where IT expenditure information is missing, we impute the missing values using the average of the IT investment from the closest before and after years that have non-missing values. For instance, if IT investment in 2008 is missing, we impute it using the average IT investment for the plant in 2007 and 2009 or using the 2007 and 2010 values if 2009 is missing. Similar logic is applied to missing values from other years. Our core results are robust to excluding observations with missing IT data.

Finally, we explore whether the relationship of DDD to productivity depends on the organization of decision-making within the firm. To do this, we construct an index of delegation from the 6 questions on where decision-making for different types of activities lies within the firm. This measure is only valid for non-headquarters firms and is therefore used only on the subsample of observations that belong to multi-establishment firms (roughly 14,000 plants in the balanced panel).

Controls. Our main specifications rely on value-added as the dependent variable and plant-level fixed effects to control for time-invariant unobserved heterogeneity. For both cross-sectional and panel analyses, we include a large number of time-varying controls collected from the ASM files. These include depreciated capital stock (calculated following the perpetual inventory method), labor measured in terms of the number of employees, and energy inputs.[10] This approach is useful for measuring a plant's productivity, because it estimates how much output the plant creates controlling for how much it spends on primary inputs, similar to a standard total factor productivity (TFP) approach. We also conduct a standard TFP analysis, were we estimate a log-linear production function with the aforementioned inputs but with industry-specific factor shares, taking the residual as a measure of plant productivity. We further explore alternative measures of plant performance, including the total value of exports and total "markup", which is calculated as the operating profit (value added minus wages and salaries) divided by total sales at the plant (*pending disclosure*).

We use other observable characteristics of the plant to control for potential unobserved heterogeneity and to explore interesting variation in the data. An important feature of manufacturing plants is their "multi-unit status", which is an indicator of whether a plant belongs to a firm with more than one manufacturing establishment. We also use information on how many other plants belong to the same parent firm, which is calculated using the Longitudinal Business Database (LBD). We also observe and use information on whether or not the plant is a headquarters (in addition to being a production unit – non-production administrative offices are collected in a separate survey), how many products it produces

[10] The energy consumption measure is calculated by combining expenditures on electricity and fuels, logging the value, and then winsorizing it to reduce the impact of outliers and help with disclosure avoidance. We log the capital and labor measures, as well, to address the highly skewed nature of their distributions.

at the 7-digit NAICS industry code level, its age (also from the LBD), and whether or not it accepts orders electronically (e-commerce). Weighted means and standard deviations for all of these variables are reported for the balanced sample in Table 1; pairwise correlations can be found in the Appendix (Table A2).

In order to control for possible heterogeneity in the quality of response to the survey that might vary systematically in ways we can address, we include "noise controls" in many specifications. These include: (1) measures for the distance between ASM and MOPS reported employment for 2005 and 2010; (2) online filing indicator; (3) date of filing and date; (4) day of week; (5) tenure of the respondent; (6) seniority of the respondent.

Industry Statistics. In all of our cross-sectional analyses, we control for industry using either 4- or 6-digit NAICS indicators. In terms of understanding the phenomenon, it is separately interesting to observe variation by industry in DDD adoption, a topic which we explore in depth, next.

4. ADOPTION OF DATA-DRIVEN DECISION-MAKING

While our core research question concerns the performance implications of data-driven decision-making, it is both interesting and useful to step back and first look at which firms and industries are more intensively using DDD over the five-year period we observe. Adoption of DDD varies considerably by certain characteristics of firms and industries. Table 4 provides some more descriptive statistics on DDD adoption, breaking it down by both year and 3-digit NAICS industry for the balanced sample. Top industries for DDD adoption by 2010 include: Transportation (NAICS 336), Food (NAICS 311), Paper (NAICS 322) and Chemicals (NAICS 325) manufacturing. Casual empiricism points to similarities among these industries related to the capital intensity of production and the tendency to work on continuous or near-continuous flows of product. Their greater use of DDD could be due to greater prevalence of operating environments where more standardized processes, greater monitoring, and responsiveness to exceptions in the production process are essential (Hayes and Wheelright 1979). We explore capital intensity, in particular, in more detail below.

In contrast, laggard industries include Apparel and Leather (NAICS 315 and 316 combined[11]), Furniture (NAICS 337), and Printing (NAICS 323). While there is significant heterogeneity within industries, these latter production environments more often involve discrete manufacturing techniques and tend to be more labor-intensive. The most significant change in the prevalence of DDD practices is in the Beverage and Tobacco manufacturing industries, which range from wineries to cigarette manufacturers, and Transportation, which includes aerospace and automobile manufacturing.

To more systematically explore the potential drivers of DDD adoption, we estimate a standard probit model of adoption (David 1969) on the 2010 complete-information cross section, with our preferred definition of DDD as the binary dependent variable. Table 5 reports the average marginal effects calculated at the sample mean of the other covariates. We control for logged employment in all columns and the industry variation described above is controlled for with 86 4-digit NAICS controls in all specifications. Thus, the results should be interpreted as within-industry relationships.

To begin, greater investment in IT at the establishment is correlated with the presence of DDD (see columns 1-5, first row). This correlation is consistent with complementarities between IT and DDD (Athey and Stern 1998, Brynjolfsson and Milgrom 2013), whereby the firms that have more IT reap a greater reward from DDD and vice versa -- a relationship we explore further in Section 5.[12] Having an educated workforce and structured management at the plant are highly predictive of DDD adoption, again consistent with potential complementarities (see all columns of Table 5). Consistent with Table 2b, Column 3 reports that belonging to a multi-unit firm is correlated with a greater likelihood of DDD adoption. Also, being in the top quartile of the capital stock distribution is associated with a greater likelihood of DDD, consistent with an operating environment that is more automated and more dependent on data collection and interpretation. All of these relationships are statistically significant, most at the 1%

[11] We combine these industries in our descriptive statistics to help prevent inadvertent disclosure of firm participation in the survey.

[12] While a formal test for complementarities is beyond the main focus of this current version of the paper, by revealed preference, if two firm choices are more productive in combination than they are separately, they ought to be observed to occur together more frequently in the population than pure chance would predict (Brynjolfsson and Milgrom 2013).

level, and for the most part empirically large. These patterns are important to keep in mind when we interpret the relationship between DDD and firm performance in the next section.

Columns 4 and 5 restrict the sample to only the multi-unit establishments. Delegated decision-making at the plant has a negative correlation with DDD, though the effect is too noisy to be statistically significant in most specifications. Plants that are also the headquarters for larger firms are less likely to adopt DDD, as well (column 4), although this effect becomes noisy when we control for the number of establishments at the parent firm (column 5). Plants belonging to firms with a larger number of establishments (i.e., large firms in terms of sites, not just employment) are significantly more likely to have DDD by 2010.

Because the bulk of our analyses are run on the balanced sample, column 6 shows the equivalent specification for this more restrictive sample. The results remain largely unchanged, although the magnitude on the IT coefficient is reduced to the point that it is not statistically different from zero. This could be related to the fact that the younger and smaller establishments that tend benefit most from the marginal investment in IT (Tambe and Hitt 2012) are excluded from this sample.

5. DATA-DRIVEN DECISION-MAKING AND FIRM PERFORMANCE

In order to investigate the relationship between DDD and performance, we take a conventional approach to modeling the plant production function. Assume that the establishment production function is as given in equation (1):

$$Y_{it} = A_{it} K_{it}^{\alpha} L_{it}^{\beta} E_{it}^{\gamma} IT_{it}^{\lambda} e^{\mu SM_{it}} e^{\eta X_{it}} e^{\delta DDD_{it}} \tag{1}$$

where Y_{it} is real value added (output - materials), A_{it} is productivity, K_{it} denotes the establishment's capital stock at the beginning of the period, L_{it} is the labor force, E_{it} is the establishments consumption of energy inputs, IT_{it} is the establishment's IT capital stock (hardware and software) at the beginning of the period, SM_{it} is a measure of structured management at the establishment, X_{it} is a vector of additional factors like industry and education, and DDD_{it} is our measure of data-driven decision-making.[13]

Taking logs provides a tractable form to take to the data:

$$Log(Y_{it}) = a\log(K_{it}) + \beta\log(L_{it}) + \gamma\log(E_{it}) + \lambda\log(IT_{it}) + \mu SM_{it} + \eta X_{it} + \delta(DDD_{it}) + f_i + \varepsilon_{it} \quad (2)$$

where the productivity term has been decomposed into a set of industry (or establishment) fixed effects f_i and an added stochastic term, ε_{it}. Because we often have multiple establishments per firm, we cluster our standard errors at the firm level.

Although our core findings rest on panel data techniques, in Table 6 we first explore the conditional correlation between DDD in the 2010 cross section and value-added, controlling for capital stock, employment, energy expenditure, and fuel costs. As discussed in Section 3, this is similar in spirit to a TFP analysis, except that it demands a bit less of the data. Most importantly, this approach has the advantage of including the entire sample for which we have information in 2010 (particularly smaller and younger plants). Recall that this sample consists of roughly 34,000 plants across the U.S. Manufacturing sector, making it by far the largest and most diverse data set yet available for this type of study. However, it poses some challenges for identification. To address this as much as possible, we include a rich set of industry controls at the 6-digit NAICS level and other establishment controls such as multi-unit status.

Column 1 establishes the baseline correlation of value-added with DDD, reporting a large coefficient of 0.120, which would be equivalent to 12.7% greater value added for firms that follow this practice at a relatively intensive level. To separate out the effect of data-driven management practices from generic investment in IT, columns 2 and 3 separately control for IT capital stock. Column 2 uses the logged value of deflated IT capital accumulated at the plant since Census started tracking hardware investment in 2002 and software investment in 2006. This is potentially useful in light of the prior literature emphasizing the relationship between generic IT investment and productivity (e.g., Brynjolfsson and Hitt 1995; Dunne, Foster et al. 2004) and the logical relationship between IT investment and data

[13]We put the management score and x_{it} controls to the exponential for the convenience of including them in levels rather than logs.

collection and use. Perhaps surprisingly, the coefficient on DDD remains quite robust to this inclusion, suggesting that IT investment and use have separate and important relationships to firm performance.

Dividing the IT variable into “High IT” (top quartile) and “Medium IT” (50-75th percentile) in column 3 to address its highly skewed distribution does not change the point estimate. However, it illustrates that firms at the top of the IT distribution enjoy significantly higher productivity – the important variation is not taking place at the mean of the sample.

Column 4 addresses the possibility that firms relying on data in their decision-making may simply have superior management practices in place at the firm. This specification includes a Z-score for all non-data, non-performance, and non-firing questions addressed by the first section of the MOPS survey (details in Section 3); it does account separately for the dismissal/re-assignment questions (15 & 16). It also controls for the skill level at the plant in terms of the percentage of managers and non-managers with bachelor degrees. The coefficient on DDD goes down substantially – from .112 to .086 – suggesting that care is required here. However, the robustness of the DDD coefficient suggests that important differences remain between structured management, more generally, and the specific practice of DDD.

Finally, column 6 restricts the analysis to the observations in the balanced sample for comparison with the next set of results. Here, it appears that the marginal relationship between DDD and performance is somewhat lower, though the confidence intervals overlap once standard errors are taken into account.

Despite the larger and more representative sample and rich set of controls, the primary concern with this cross-sectional analysis is the likelihood that unobservable plant characteristics may affect both productivity and the use of DDD. To address this, we construct a two-period panel using the 2005 recall questions and 2005 ASM as described in Section 3. As long as the unobserved factors that are correlated with DDD do not vary over time at the establishment level (corresponding to f_i in equation (1)), they can be differenced out by running a fixed-effect panel regression. Of course this coefficient may still be upwardly biased if management practices are proxying for time-varying unobserved coefficients. At the same time, the coefficient on management may be attenuated towards zero by measurement error, and this downward bias is likely to become worse in the fixed-effect specification.

Table 7 presents the differences-in-differences results that are central to our findings. These specifications are estimated using the balanced panel of roughly 18,000 plants that report data on the relevant questions for both years and can be linked up to the appropriate ASM surveys for the performance measures and controls. Column 1 shows the DDD coefficient controlling for changes in IT capital stock but not changes in structured management; it is commensurate with the coefficient from column 6 of Table 6. This estimate is robust to excluding the IT control (available upon request). Controlling for changes in structured management practices has a large effect in column 2. Column 3 is similar, but, as in the previous table, divides the IT capital measure into indicators for being in the top and 3rd quartiles. This, our preferred specification, reports a correlation between value-added and DDD of just over 3%. To put these findings in context: to achieve the same advantage in value-added from IT investment as from DDD adoption would, at the mean of the IT distribution entail an additional $5 million in IT capital accumulation over the five-year period.[14] This has interesting implications for the relative benefit of investment in tangibles such as computer hardware and software versus investment in intangibles such as processes and practices that better leverage existing technologies.

Another insight from column 3 is that unobserved heterogeneity may have a non-trivial impact on the observed relationship between DDD and firm performance. Controlling for time-invariant plant characteristics, the magnitude of the DDD coefficient is roughly half of what it was in the cross-sectional analysis.

Columns 4 and 5 reveal that the magnitude of this relationship varies a great deal by whether the plant is a single-establishment firm or belongs to a larger, multi-plant firm. Most of the benefit appears to accrue to single-unit firms that adopt DDD, as opposed to plants belonging to larger firms. The coefficient for single-unit plants is a striking .127, and significant at the 1% level; that for multi-unit plants is only .02 and very noisy. Some of this may be due to the timing of when the establishment adopted DDD and

[14] However, as noted previously, the IT capital distribution is quite skewed. A less intuitive but nonetheless salient comparison is that the DDD coefficient is nearly identical to that for the indicator of being in the top quartile of the IT capital distribution (High IT). In rough numbers, the difference in IT capital stock between an average plant in the top quartile of the IT distribution versus an average plant in the bottom half is approximately $1 million.

the rate at which the frontier of practice progresses. As described in Table 5, multi-unit establishments are disproportionately more likely to have cleared the threshold for intensive DDD by 2010. In fact, multi-unit plants have an incidence of DDD according to our definition that is nearly three times that for single-unit firms – Table 2b shows that by 2010 they report a mean adoption of DDD of roughly 30%, while single-unit plants report an estimated mean adoption of 11%. This pattern is seen in 2005, as well: 12% of multi-unit plants used data intensively, while only 4% of single-unit firms did.

This information is useful for two reasons. The first is to point out that the loss of precision for the multi-unit sub-sample is not due to the reduction in sample size – the number of plants that change from non-DDD status to DDD status is much larger for multi-unit plants than single-unit plants. The second is to raise the possibility that what constitutes the "frontier" of practice may differ systematically between these two types of establishments.

This latter hypothesis appears to be supported by the findings reported in Table 8. This table reports the coefficient on DDD for a specification identical to that used in column 3 of Table 7. Each row, however, reports on a different combination of data-related practices in place of our core DDD measure. The first column shows that the estimates for the entire sample from Table 7 appear quite robust to perturbations in the definition of DDD. However, the next two columns split the sample by multi-unit status, and some interesting differences appear. Columns 2 and 3 of the first row show that the advantage of DDD seen in Table 7 is replicated even if we only consider the top two categories for *availability* of data (question 27). This is significant at the 5% level for the multi-unit establishments. The *use* of data in decision-making is more meaningful for the single-unit firms according to the second row of the table, where the coefficient is .049 and significant at the 10% level; it is practically zero for the multi-unit plants. Adding restrictions from the first section of the survey, which focuses on monitoring of production processes, appears to be a meaningful exercise for the single-unit plants, but not terribly informative for the multi-unit ones.[15] This could be because plants belonging to larger firms are relatively ahead of

[15] The definition of DDD becomes more restrictive as we move down Table 8, as the intersection of additional practices becomes increasingly small. We do not report on individual practices in isolation due to noisiness of the

smaller, single-establishment firms when it comes to the state of management practices, be they data-focused or not. (*Specific means for individual practices have not been cleared for disclosure*). A possible implication is that we need additional information to separate out the most advanced multi-unit plants from the rest.

Another plausible explanation exists for the single-unit/multi-unit divide, however. Not reported for disclosure avoidance reasons, a different sample split on whether a plant belongs to a firm with five or fewer units (i.e., a split which lumps single-unit and multi-unit establishments belonging to small parent firms together) yields a very similar picture to that shown in Table 7. Multi-unit establishments of small firms also show a very high correlation between DDD and value-added. This raises the question of whether larger firms face higher adjustment costs than smaller firms, which we might then observe at the plant level in the form of lower net benefits from investments in DDD despite their higher propensity to adopt.

Table 9 explores the possibility that adjustment costs can help explain why multi-unit plants belonging to larger firms do not appear to benefit significantly from DDD adoption. A useful starting point is to investigate whether the size of the organization mediates the relationship between DDD and value-added. If coordinating inputs from a greater number of employees and/or data sources increases the costs of choosing KPI's, establishing targets, or interpreting data inputs, then it might more difficult to implement DDD. Conversely, smaller organizations ought to find it easier or less costly to reach the frontier of practice. Column 1 reports the results from our two-period difference model including an interaction between an indicator of whether an establishment was relatively small in 2005 (small size is captured by an indicator of whether the plant had fewer than 100 employees in 2005, which is roughly the mean of the logged employment distribution for that year). Note that we keep the size indicator fixed at 2005 levels so as not to confound changes in size as well as changes in DDD status, thus the direct term is differenced out of the regression; this approach holds for the other specifications in this table, as well. In

estimates, probably due in part to measurement error whenever a single question from the survey is taken in isolation – a point discussed in Bloom et al. (2014) as well.

this model, *all* of the benefit of DDD accrues to the smaller establishments, with the coefficient of DDD becoming indistinguishable from zero and the interaction effect having an effect of nearly 11% (with a coefficient of 10.4) and significant at the 1% level. Column 2 suggests that younger establishments (defined as being five or fewer years old in 2005) – where legacy management practices might be less ingrained – also disproportionately benefit from DDD. Column 3 interacts DDD with an indicator of whether the establishment belonged to a parent firm with fewer than five sites in 2005. Again, the interaction effect is large and positive at nearly 9%, while the coefficient on DDD by itself is indistinguishable from zero. The pattern of results across this table is consistent with DDD practices having an impact in environments where adjustment might plausibly be lower, on average.

An outstanding question at this point is whether the relationship between DDD and performance can be interpreted as a causal one. As we lack an exogenous shock to DDD adoption, we hesitate to claim unequivocally that DDD *causes* better firm performance. There are certain trends in the data, however, that are consistent with a causal story.

The classic problem with ascribing positive performance results to any type of technology adoption is that adoption is typically voluntary, and those adopting are more likely to be those who expect to benefit from adoption (David, 1969; Rogers 2010). Good instruments for adoption are also typically difficult to find in large, heterogeneous populations. Our context is subject to these critiques as well, with the following caveats. First, DDD is an intangible practice that requires investments in managerial attention and time, but it is not a standard investment in tangibles such as various types of capital that typically requires free cash flow -- and hence it is less likely that only firms with prior good performance will be capable to invest in DDD.

Second, we are able to exploit the rich Census data to gain some insights into the timing of the effects that we observe. Although we only observe DDD and management practices for 2005 and 2010, we can observe performance and non-MOPS-based controls for every year preceding our sample and for two years afterwards. If better performance were preceding DDD adoption, a regression of DDD adoption on value-added in our pre-period should show a positive relationship. However, a probit analysis of DDD adoption in 2005 (Table 10) shows that value-added growth from 2002 – 2005 does not predict the presence of DDD in 2005. A similar regression can be run for those plants that did not have DDD in 2005, using growth in the 2005 – 2010 period as the key explanatory variable. Again, growth in value added does not predict adoption of DDD in the window between 2005 and 2010. Finally, we construct a balanced panel of ASM performance data and controls for every year from 2002 to 2012 and interact indicators of DDD by 2005 and DDD by 2010 with every year to see if performance anticipates the year in which DDD is reported to be in practice at the plant. While there is considerably more noise in this smaller sample, the timing of effects is consistent with a causal effect. Specifically, the correlation between having DDD in 2005 and value-added (controlling for the standard inputs and establishment observables) shows up right around then. (It actually begins to appear in 2004, but this would be consistent with adoption happening for some plants before 2005, which is not at all ruled out by the structure of the recall questions). For the plants that do not have DDD in 2005 but report it by 2010, the coefficient on DDD does not begin to show up until 2010, and actually increases slightly by 2011. Note that this pattern would furthermore be consistent with DDD-related adjustment costs (*Graph pending disclosure review*).

Table 11 further explores the complementarities that were hinted at in the adoption regressions. We find evidence consistent with complementarities between IT adoption and DDD. Column 1 provides additional evidence consistent with complementarities between DDD and IT investment. Again, we fix the indicator for having IT capital stock in the top quartile of the 2005 distribution, so the direct effect is differenced out. The interaction effect has a coefficient of .075 and is significant at the 1% level. The coefficient on DDD by itself is not significantly different from zero, suggesting that almost all of the difference in performance occurs in establishments that already have significant IT investments and subsequently adopt DDD. This is unsurprising to the extent that IT infrastructure is needed to collect, track, and analyze the data inputs to decision-making at the firm, a theme that occurred often in our qualitative interviews. However, we noted in our analyses that this interaction effect is strongest for the multi-unit establishments (not reported), pointing again to potentially important systematic differences in how sophisticated and "teched up" their DDD practices may be when compared to smaller, single-unit firms.

Likewise, more-skilled labor also appears to be complementary with DDD in column 2 of Table 11, with the effect driven entirely by establishments that had a greater than the median percentage of managerial and non-managerial employees with Bachelor's degrees in 2005 (again, the degree term by itself is differenced out). Plants with a greater than average value of exports in 2005 also are the main beneficiaries of DDD adoption.

Perhaps one of our most tantalizing findings is evidence that DDD is complementary to the way in which decision-making is organized throughout the firm. Recall that we created an index of delegated decision-making that combines all the different types of decisions from the second section of the MOPS and is based on whether authority rests *exclusively* with the plant, and not either with headquarters or shared between the two. We create an indicator for whether this index is above average for all the plants in 2005, and interact it with DDD adoption. In this case, DDD by itself retains a positive and significant coefficient. However, the interaction effect has a large and significant negative coefficient, suggesting that DDD is less beneficial in decentralized organizations. We speculate that this may be because central decision makers can access more of the data needed to make decisions when information is explicit rather than tacit, creating more benefits to data collection and analysis when decision-making is not distributed throughout the firm. Combine with the result on education, this suggest that there may be an increase in separation between what Piore and Sabel (1984) called "conceptualization vs. execution". The data-driven approach requires analysis by educated workers to be effective, but that seems to leave less of a role for tacit, and perhaps more informal knowledge distributed throughout the firm. [16]

Table 12 provides a few robustness checks that take advantage of the extremely rich Census data to explore specifications with additional controls. Including controls for whether or not the plant is an

[16] We explored other interactions, but found little to be statistically significant, including continuous improvement practices (Q01), firing practices (Q15 and Q16), nonunion status, vertical integration measured as the value of within-firm transfers between plants, communicating with the firm about production targets (questions 5 and 8) and frequency of review of KPIs (questions 3 and 4).

exporter, its age, and whether or not it conducts e-commerce have no effect on the core findings. Additionally, based on qualitative interviews, we checked that the results are robust to restricting the sample to managers who have been at the plant five or more years; the core findings remain unchanged.

6. CONCLUSION

Theory and case evidence suggest that productivity gains come from leveraging IT investments to collect and bring data to bear in managerial decision-making, tracking performance within the firm, and communicating about the state of the production process. This paper provides clear statistical evidence that putting data "into action" in these ways is associated with significantly higher productivity. Specifically, it is not simply the adoption of IT, but the adoption of a specific set of data-driven decision-making practices that are correlated with the greatest increases in productivity. This relationship holds across a variety of industries in the U.S. manufacturing sector and stands up to the inclusion of difficult-to-observe controls for the level of IT investment and other management practices at the organization. Our findings provide the first evidence that this correlation is robust to unobserved, time-invariant, plant heterogeneity that might otherwise threaten the validity of these findings. Finally, the timing of the adoption decision and performance effects is consistent with a causal relationship between DDD and performance.

Certain limitations of this study are worth noting. For instance, we do not observe directly what types of decision-making are being influenced by data. Based on the content of the survey directly preceding the DDD questions, one might infer that respondents have a range of types of decision-making in mind: HR decisions such as hiring and pay raises, product strategy decisions such as new product introductions and pricing, marketing decisions such as advertising spend, and financial decisions pertaining to the purchase of new capital assets.

Also, we do not know precisely what respondents have in mind when asked about "data." We infer from the priming earlier in the survey -- as well as from a limited number of qualitative interviews –

that respondents at least think about Key Performance Indicators regarding cost, waste, quality, and so on. However, there may be significant differences between firms and industries in whether respondents are thinking about automated data collection and analysis based on significant on-site IT investments, or whether they have in mind some of the relatively “low tech” data management techniques that are nevertheless effective in many lean manufacturing settings (e.g., Kanban systems for ordering parts). The results pointing to complementarities with IT investment suggests the former, but our broad-brush analysis surely misses some interesting and important nuances in how DDD plays out in firms of different types.

Finally, the potential for bias due to self-selection into DDD persists. The timing of adoption and performance we observe is consistent with a causal explanation. However, firms in our sample adopt DDD as a choice, and therefore there is no sense in which firms are exogenously “treated with DDD” – our findings can best be taken as informative about the effect of treatment on the treated. However, we hope they will also be taken as compelling evidence about the potential importance of this practice and the usefulness of distinguishing investments in data-related practices from investments in tangible IT capital. Our findings suggest that the benefits may be commensurate, while we speculate that the costs may differ significantly in nature and magnitude. In particular, the former may be more in the form of managerial time and attention, and therefore invisible on the balance sheet and difficult to observe and copy – potentially conferring some measure of competitive advantage.

We hope to spur further research into the relationship between data-driven decision-making and firm performance in manufacturing and other sectors of the economy. Given the large increases we are certain to see in both IT capabilities and the availability of digital data for use in decision-making, the effects we identify, and the role of complementary changes in organizations, may grow even more economically important in the coming years.

REFERENCES

Aral, S. and P. Weill (2007). "IT Assets, Organizational Capabilities, and Firm Performance: How Resource Allocations and Organizational Differences Explain Performance Variation." *Organization Science* **18**(5): 763-780.

Athey, S. and S. Stern. 1998. An Empirical Framework for Testing Theories About Complementarity in Organization Design. Working Paper #6600. NBER.

Bloom, N., E. Brynjolfsson, L. Foster, R. Jarm, I. Saporta-Eksten, J. Van Reenen (2014). “Management in America.” CES Working Paper 13-1 .U.S. Census Bureau’s Center for Economic Studies Research Program, Washington D.C.

Bresnahan, T. F. and S. Greenstein (1996). "Technical Progress and Co-Invention in Computing and in the Uses of Computers." *Brookings Papers on Economic Activity. Microeconomics* **1996**: 1-83.

Brynjolfsson, E. and L. Hitt (1995). "Information Technology As A Factor Of Production: The Role Of Differences Among Firms." *Economics of Innovation and New* Technology **3**(3-4): 183-200.

Brynjolfsson, E., L. M. Hitt, et al. (2011). “Strength in Numbers: How Does Data-Driven Decisionmaking Affect Firm Performance?” MIT Working Paper.

Brynjolfsson, Erik and Paul Milgrom (2013). "Complementarity in Organizations". In *The Handbook for Organization Economics*, edited by Robert Gibbons and John Roberts, 11-55. Princeton: Princeton University Press.

Dunne, T., L. Foster, et al. (2004). "Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment." *Journal of Labor Economics* **22**(2): 397-429.

Galbraith, J. R. (1974). “Organization design: An information processing view.” *Interfaces*, 4(3), 28–36.

Gartner (2014) “Survey Analysis: Big Data Investment Grows but Deployments Remain Scarce in 2014,” http://www.gartner.com/document/2841519.

Jensen, J. B. and R. H. McGuckin (1996). “Firm Performance and Evolution: Empirical Regularities in the U.S. Microdata.” Center for Economic Studies Working Paper CES 96-10. U.S. Census Bureau.

Hayes, R.H. and S.C. Wheelright (1979). “The Dynamics of Process-Product Lifecycles.” Harvard Business Review (March-April):127-136.

Henderson, R. and R. Gibbons (forthcoming). “Relational Contracts and Organizational Capabilities.” *Organization Science.*

McElheran, K. (forthcoming). "Do Market Leaders Lead in Business Process Innovation? The Case(s) of E-Business Adoption." *Management Science*.

Piore, M. and C. Sabel (1984) *The Second Industrial Divide*, New York, Basic Books.

Syverson, C. (2011). "What Determines Productivity?" *Journal of Economic Literature* **49**(2): 326-365.

Tambe, P. and L. Hitt. 2012. The Productivity of Information Technology Investments: New Evidence from IT Labor Data. Information Systems Research, 23(3-1), 599-617.

Table 1. Descriptive Statistics for non-DDD variables

Variable	Description	2010 Mean Balanced Sample (weighted*)	2010 Std. Deviation Balanced Sample (weighted*)
Log Value Added	Log of value added at the plant, which is total value shipped minus total cost of goods sold	**9.23**	**1.66**
Log Sales (TVS)	Log of total value shipped by the plant	**10.02**	**1.69**
Markup	Operating profit per sales at the plant, calculated as value added minus wages and salaries, divided by total value of shipments	**.31**	**.26**
Log Exports	Log of total value of goods exported by the plant	**3.65**	**4.23**
Log Employment	Log of total number of employees at the plant	**4.34**	**1.23**
Log K stock	Log of capital stock at the plant, calculated using the perpetual inventory method and BEA capital deflators	**8.91**	**1.61**
Log Energy Costs	Lot of total cost of both fuel and electricity consumed by the plant	**5.75**	**1.78**
Multi-Unit Status	=1 if the plant belongs to a multi-unit firm	**.71**	**.46**
HQ Status	=1 if the plant is a headquarters; equal to 1 for all single-unit firms	**.49**	**.50**
Exporter Status	=1 if the value of exports is >0	**.46**	**.50**
E-Com Dummy	=1 if the value of e-commerce at the plant is >0	**.61**	**.49**
Age	Derived from the year first observed in the LBD. Truncated at 44 due to LBD starting in 1976	**24.9**	**10.5**
Log IT Capital Stock	Log of value of hardware and software stocks at the plant, calculated using the perpetual inventory method and BEA deflators.	**3.83**	**1.96**
Structured Management Z-Score	Index created by summing up the normalized values from questions 1, 5, 8, 9, 11, 13 & 14 of the MOPS	**.61**	**.19**
% Employees with BA	% of managers and non-managers with bachelor's degrees at the plant	**.08**	**.25**
Number of sites	Number of establishments belonging to the parent firm	**107.22**	**368.69**
Number of products	Number of products produced at the plant (7-digit NAICS)	**3.55**	**2.07**

* Means and standard deviations estimated using ASM sampling weights. DDD means are reported in Table 3.

Table 2a. Adoption of Data-Related Managerial Practices by Year, all establishments

Data-Related Management Practice	**Adoption in 2005 Balanced sample** (weighted*)	**Adoption in 2010 Balanced sample** (weighted*)	**Adoption in Entire 2010 Cross-section sample** (weighted*)
Top 2 categories for "availability of data" (Q27)	.43	.69	.59
Top 2 categories for "use of data" (Q28)	.41	.66	.55
Top 2 categories for both availability and use of data (Q27 & Q28)	.32	.58	.47
Top for availability and use of data, plus tracking 10 or more KPI's (Q27, Q28, & Q2)	.14	.34	.22
"Data-Driven Decision-making (DDD)" as defined for main specifications: as above, as well as use of short-term and long-term targets (Q6)	.09	.25	.16
N	~18,000	~18,000	~34,000

*Mean adoption estimated using ASM sampling weights

Table 2b. Adoption of Data-Related Managerial Practices, split by multi-unit status

	Adoption in 2005 Balanced sample **Multi-Unit Plants** (weighted*)	**Adoption in 2010 Balanced sample** **Multi-Unit Plants** (weighted*)	**Adoption in 2005 Balanced sample** **Single-Unit Firms** (weighted*)	**Adoption in 2010 Balanced sample** **Single-Unit Firms** (weighted*)
"Data-Driven Decision-making (DDD)"	.12	.30	.04	.11
N	~14,000	~14,000	~4,000	~4,000

*Mean adoption estimated using ASM sampling weights

Table 3. Principal Factor Analysis of Data-Related Managerial Practices

Principal Factor Analysis of 2010 Balanced Sample (~18,000 obs.)	**Eigenvalue**	**Proportion of Variance**
Factor 1	2.28	.570
Factor 2	.851	.213
Factor 3	.648	.162

Polychoric Correlation Matrix and Factor Loadings	Top 2 categories for "availability" of data	Top to categories for "use" of data	Track 10 or more KPIs	Use of short-term and long-term targets	**Factor 1 Loadings**
Top 2 categories for "availability" of data	1				.870
Top 2 categories of "use" of data	.778	1			.854
Track 10 or more KPIs	.366	.385	1		.664
Use of short-term and long-term targets	.291	.328	.343	1	.595

Note: Calculated using the **polychoric** command for Stata 13.

Table 4. Adoption of Data-Driven Decision-making by Industry (3-Digit NAICS Code)

3-Digit NAICS Code	Industry	Mean DDD in 2010 Balanced Sample (weighted*)	Mean DDD change 2005 to 2010 Balanced Sample (weighted*)
311	Food	.39	.23
312	Beverage and Tobacco Products	.40	.28
313	Textile Mills	.32	.18
314	Textile Product Mills	.21	.15
315/316	Apparel and Leather	.13	.06
321	Wood Products	.22	.12
322	Paper	.37	.24
323	Printing and Related Support Activities	.19	.14
324	Petroleum and Coal Products	.32	.13
325	Chemicals	.37	.21
326	Plastics and Rubber	.31	.22
327	Non-metallic Mineral Products	.24	.14
331	Primary Metals	.33	.22
332	Fabricated Metal Products	.25	.17
333	Machinery	.24	.17
334	Computers and Electronic Products	.33	.21
335	Electrical Equipment, Appliances, and Components	.34	.24
336	Transportation Equipment	.41	.25
337	Furniture	.18	.14
339	Miscellaneous	.26	.17
321, 322, 324, 325, 327, 331	"Continuous Flow" industries	.34	.19

* Mean adoption estimated using ASM sampling weights.

Table 5. Correlates of Data-Driven Decision-making Adoption

Dependent Variable:			**DDD**			
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
IT capital stock (logged)	.007*** (.002)	.003** (.001)	.003** (.001)	.005** (.002)	.005** (.002)	.002 (.002)
% Employees with Bachelor's degrees		.117*** (.021)	.116*** (.021)	.151*** (.036)	.144*** (.036)	.143*** (.032)
Structured management		.431*** (.017)	.398*** (.016)	.692*** (.036)	.674*** (.036)	.673*** (.028)
Multi-unit status			.058*** (.006)			.058*** (.012)
High capital stock (top quartile)			.020*** (.006)	.046*** (.011)	.044*** (.071)	.031*** (.010)
Delegation				-.035 (.022)	-.031 (.022)	
HQ				-.039* (.021)	-.029 (.021)	
Number of sites at parent firm (logged)					.010*** (.003)	.009*** (.002)
Number of employees (logged)	.061*** (.003)	.034*** (.003)	.027*** (.003)	.030*** (.005)	.029*** (.006)	.036*** (.006)
Observations	~34,000	~34,000	~34,000	~22,000	~22,000	~18,000
Sample	All 2010	All 2010	All 2010	All 2010 MU	All 2010 MU	Balanced 2010
Fixed effects	NAICS4	NAICS4	NAICS4	NAICS4	NAICS4	NAICS4

Note: Weighted Maximum likelihood probit estimation of the likelihood that an establishment reports relatively intensive use of data-driven decision-making (DDD) in 2010 using ASM sampling weights. Reporting average marginal effects calculated at sample means of the covariates. Columns 1-3 report on the complete 2010 cross section, columns 4-5 restrict to plants belonging to multi-unit firms, column 6 is restricted to observations that appear in the balanced sample (both single- and multi-unit establishments). All columns include industry controls at the 4-digit NAICS level. Robust standard errors are clustered at the firm level and reported in parentheses. Statistical significance is denoted as follows: * 10%, ** 5%, *** 1%.

Table 6. Conditional Correlations between Data-Driven Decision-making (DDD) and Firm Performance

Dependent Variable:	**Logged Value Added**					
	(1) DDD	**(2) Add IT Stock**	**(3) IT Categories**	**(4) Mgmt. & Education**	**(5) Noise Controls**	**(6) 2010 Balanced Sample**
DDD	**.120*** (.015)**	**.112*** (.015)**	**.112*** (.015)**	**.086*** (.015)**	**.084*** (.015)**	**.069*** (.017)**
IT capital stock (logged)		.048*** (.004)				
High IT			.222*** (.018)	.209*** (.018)	.195*** (.018)	.192*** (.024)
Med IT			.133*** (.013)	.127*** (.013)	.118*** (.013)	.102*** (.019)
Structured management index				.239*** (.042)	.239*** (.042)	.292*** (.061)
% Employees with Bachelor's degrees				.309*** (.062)	.302*** (.062)	.408*** (.076)
Capital, Labor, and Energy inputs (logged)	Yes	Yes	Yes	Yes	Yes	Yes
Observations	~34,000	~34,000	~34,000	~34,000	~34,000	~18,000
Sample	All 2010	All 2010	All 2010	All 2010	All 2010	Balanced 2010
Fixed Effects	NAICS6	NAICS 6	NAICS 6	NAICS 6	NAICS 6	NAICS6

Note: Weighted OLS regressions using ASM sampling weights. The dependent variable is logged nominal value added at the plant. The structured management index is an unweighted average of the score for questions 1, 5,8,9,11,13 & 14, where each question is first normalized to a 0-1 scale. Unreported controls in all columns include: firing practices (questions 15 & 16), whether the establishment belongs to a multi-unit firm, logged capital stock, logged employment, winsorized logged energy expenditures, and whether or not information on employees' education is missing. Noise controls (column 5 onwards) include: (1) measures for the distance between ASM and MOPS reported employment for 2005 and 2010; (2) online filing indicator; (3) date of filing and date; (4) day of week; (5) tenure of the respondent; (6) seniority of the respondent. Industry controls are included at the 6-digit NAICS level. The sample in columns 1-5 is all MOPS observations with at least 11 non-missing responses to management questions for 2010 and a successful match to the 2010 ASM, which were also included in ASM tabulations, and have positive value added, positive employment and positive imputed capital. Column 6 restricts to the establishments that fit these criteria for 2005 as well. Robust standard errors are clustered at the firm level and reported in parentheses. Statistical significance is denoted as follows: * 10%, ** 5%, *** 1%.

Table 7. Fixed-Effects Estimation of Data-Driven Decision-making and Firm Performance

Dependent Variable:	**Log Value Added**				
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
	Control for IT	**Control for Mgmt & Educ.**	**IT Categories**	**SU Firms**	**MU Estabs**
DDD	**.073*** (.014)**	**.033** (.016)**	**.033** (.016)**	**.127*** (.038)**	**.020 (.017)**
IT capital stock (logged)	.010*** (.003)	.007*** (.003)			
High IT			.031* (.016)	.062* (.028)	.024 (.019)
Med IT			.018 (.012)	.028 (.023)	.016 (.014)
Structured management index		.284*** (.033)	.284*** (.048)	.298** (.122)	.275*** (.053)
% Employees with Bachelor's degrees		-.158 (.115)	-.159 (.115)	.095 (.256)	-.201 (.128)
Capital, Labor, and Energy inputs (logged)	Yes	Yes	Yes	Yes	Yes
# Establishments	~18,000	~18,000	~18,000	~4,000	~14,000
Sample	Balanced	Balanced	Balanced	Balanced	Balanced

Note: Two-period linear regression with establishment-fixed effects. In all columns the dependent variable is logged nominal value added. The structured management index is the unweighted average of the score for each of the first 16 questions, omitting questions 1, 5,8,9,11,13 & 14, where each question is first normalized to be on a 0-1 scale. Unreported controls in all columns include: firing practices (questions 15 & 16), whether the establishment belongs to a multi-unit firm, logged capital stock, logged employment, winsorized logged energy expenditures, and whether or not information on employees' education is missing. The sample in columns 1-3 is all MOPS observations with at least 11 non-missing responses to management questions for both 2005 and 2010 and a successful match to ASM in both years, which were also included in ASM tabulations for both 2010 and 2005, have positive value added, positive employment and positive imputed capital in the ASM for both 2005 and 2010. Columns 4 and 5 split this sample by multi-unit status. Robust standard errors are clustered by establishment and reported in parentheses. Statistical significance is denoted as follows: * 10%, ** 5%, *** 1%.

Table 8. Robustness to Definitions of DDD

	Coefficients for DDD from Table 5, Column 3 specification		
	(1)	(2)	(3)
	Full Sample	**Single-Unit Firms**	**MU Plants**
Top 2 categories for Q27	.039***	.035	.038**
	(.014)	(.027)	(.016)
Top 2 categories for Q28	.007	.049*	-.003
	(.013)	(.027)	(.015)
Top KPIs	.030**	.069**	.023
	(.014)	(.032)	(.016)
Q27 & Q28	.030**	.069***	.021
	(.013)	(.025)	(.015)
Q27, Q28, & KPI's	.040***	.104***	.029***
	(.014)	(.032)	(.016)
Q27, Q28, KPI & high frequency of review	.031*	.074**	.027
	(.017)	(.037)	(.018)
Q27, Q28, KPI, high frequency of review, and ST & LT targets	.017	.103***	.009
	(.017)	(.040)	(.019)
# Establishments	~18,000	~4,000	~14,000
Sample	Balanced	Balanced	Balanced

Note: Two-period linear regression using establishment-fixed effects. Replicating the specification in column 3 of Table 5, but with differing definitions of DDD. Columns 1 and 2 split the sample by multi-unit status. Column 3 reports the full sample (equivalent to column 3 of Table 5). Statistical significance is denoted as follows: * 10%, ** 5%, *** 1%.

Table 9. Variation in the Effects are Consistent with Adjustment Costs

Dependent Variable:	**Log Value Added**		
	(1)	**(2)**	**(3)**
	Establishment Size	**Age (Young)**	**Number of Sites**
DDD	-.028 (.019)	.024 (.016)	.012 (.018)
DDD x Small Size in 2005	.104*** (.030)		
DDD x Young in 2005		.090* (.049)	
DDD x Low Number of Sites in 2005			.085*** (.028)
Establishments	~18,000	~18,000	~18,000
Sample	**Balanced**	**Balanced**	**Balanced**

Note: Two-period linear regression with establishment-fixed effects. In all columns the dependent variable is logged nominal value added. Small Size is an indicator of whether the plant had fewer than 100 employees in 2005. Young is an indicator for whether the establishment was five or fewer years old in 2005. Low Number of Sites is an indicator for whether or not the plant belongs to a firm with fewer than 5 sites in 2005. Unreported controls include the structured management index, firing practices (questions 15 & 16), whether the establishment belongs to a multi-unit firm, logged capital stock, logged employment, winsorized logged energy expenditures, and whether or not information on employees' education is missing. The sample in all columns is all MOPS observations with at least 11 non-missing responses to management questions for both 2005 and 2010 and a successful match to ASM in both years, which were also included in ASM tabulations for both 2010 and 2005, have positive value added, positive employment and positive imputed capital in the ASM for both 2005 and 2010. Robust standard errors are clustered by establishment and reported in parentheses. Statistical significance is denoted as follows: * 10%, ** 5%, *** 1%.

Table 10. Timing of DDD Adoption is Inconsistent with Reverse Causality

Dependent Variable:	**DDD adoption**	
	(1)	**(2)**
	DDD in 2005	**Adopted DDD in 2005-2010 period**
Value-added growth 2002 - 2005	-.00004 (.00005)	
Value-added growth 2005 - 2010		.00001 (.00002)
High IT	-.0003 (.006)	.024** (.009)
Medium IT	-.006 (.006)	.013 (.008)
Total Employment (logged)	.029*** (.003)	.048*** (.004)
Multi-Unit Status	.045*** (.007)	.120*** (.010)
Structured management index		-.032 (.021)
% Bachelor's Degree		.109*** (.027)
Sample	**Balanced**	**Balanced & DDD = 0 in 2005**
Establishments	~18,000	~16,000

Note: Maximum likelihood probit estimation of DDD adoption in 2005. Reporting average marginal effects calculated at mean values of the covariates. Controls for column 1 include High and Medium IT from 2002, logged employment from 2002, whether the establishment belongs to a multi-unit firm, and the percentage of employees with Bachelor's degrees in 2005. Controls for column 2 are identical, except that they come from 2005 instead of 2002 and 2005 instead of 2010, respectively. Statistical significance is denoted as follows: * 10%, ** 5%, *** 1%.

Table 11. Interaction Tests for Complements

Dependent Variable:	**Log Value Added**			
	(1)	**(2)**	**(3)**	**(4)**
	High IT	**% Bachelor's Degrees**	**High Exports**	**Delegation**
DDD	.010 (.019)	-.001 (.024)	.004 (.022)	.042** (.020)
DDD x High IT in 2005	.075*** (.027)			
DDD x High %BA Degree in 2005		.056** (.028)		
DDD x High Export in 2005			.051* (.027)	
DDD x High Delegation in 2005				-.057* (.030)
Sample	**Balanced**	**Balanced**	**Balanced**	**Balanced MU establishments**
Establishments	~18,000	~18,000	~18,000	~14,000

Note: Two-period linear regression with establishment-fixed effects. In all columns the dependent variable is logged nominal value added. High IT in 2005 is an indicator of whether the plant was in the top quartile of the IT capital distribution in 2005. High Degree in 2005 is an indicator for whether the establishment had greater than the median percentage of employees with Bachelor's degrees in 2005. High Export in 2005 is an indicator of whether the plant was above the mean value of exports in 2005. High Delegation is an indicator of whether the plant was above the mean delegation index in 2005. Note that these values are fixed in 2005 and therefore the direct effect is differenced out. Unreported controls include the structured management index, firing practices (questions 15 & 16), whether the establishment belongs to a multi-unit firm, logged capital stock, logged employment, winsorized logged energy expenditures, and whether or not information on employees' education is missing. The sample in all columns is all MOPS observations with at least 11 non-missing responses to management questions for both 2005 and 2010 and a successful match to ASM in both years, which were also included in ASM tabulations for both 2010 and 2005, have positive value added, positive employment and positive imputed capital in the ASM for both 2005 and 2010. Robust standard errors are clustered by establishment and reported in parentheses. Statistical significance is denoted as follows: * 10%, ** 5%, *** 1%.

Table 12. Robustness to Additional Covariates

Dependent Variable:		**Log Value Added**		
	(1)	**(2)**	**(3)**	**(4)**
DDD	.033** (.016)	.030* (.016)	.031*** (.016)	.030* (.016)
High IT	. 033** (.016)	.026 (.017)	.028* (.016)	.028* (.017)
Med IT	.019 (.012)	.015 (.012)	.016 (.012)	.016 (.012)
Structured Management Index	.283*** (.048)	.261*** (.053)	.268*** (.049)	.259*** (.053)
% Employees with Bachelor's Degree	-.160 (.116)	-.174 (.116)	-.169 (.116)	-.175 (.116)
Export	.041*** (.015)			.037*** (.015)
Age		.002 (.002)		.001 (.002)
Ecommerce			.028** (.011)	.021* (.012)
Establishments	~18,000	~18,000	~18,000	~18,000
Sample	**Balanced**	**Balanced**	**Balanced**	**Balanced**

Note: Two-period linear regression with establishment-fixed effects. In all columns the dependent variable is logged nominal value added. The structured management index is the unweighted average of the score for each of the first 16 questions, omitting questions 1, 5,8,9,11,13 & 14, where each question is first normalized to be on a 0-1 scale. Unreported controls in all columns include: firing practices (questions 15 & 16), whether the establishment belongs to a multi-unit firm, logged capital stock, logged employment, winsorized logged energy expenditures, and whether or not information on employees' education is missing. The sample in columns 1-3 is all MOPS observations with at least 11 non-missing responses to management questions for both 2005 and 2010 and a successful match to ASM in both years, which were also included in ASM tabulations for both 2010 and 2005, have positive value added, positive employment and positive imputed capital in the ASM for both 2005 and 2010. Robust standard errors are clustered by establishment and reported in parentheses. Statistical significance denoted as follows: * 10%, ** 5%, *** 1%.

APPENDIX:

Table A1. Comparison Sample Descriptive Statistics

Variable	Description	2010 Balanced Sample Mean (estimated)	Entire 2010 Sample Mean (estimated)	2010 ASM Mean (estimated)
Log Value Added	Log of value added at the plant, which is total value shipped minus total cost of goods sold	9.23	8.25	7.89
Log Sales (TVS)	Log of total value shipped by the plant	10.02	8.93	8.50
Markup	Operating profit per sales at the plant, calculated as value added minus wages and salaries, divided by total value of shipments	.31	.31	.13
Log Exports	Log of total value of goods exported by the plant	3.65	2.41	1.80
Log Employment	Log of total number of employees at the plant	4.34	3.64	3.39
Log K stock	Log of capital stock at the plant, calculated using the perpetual inventory method and BEA capital deflators	8.91	8.26	N/A
Log Energy Costs	Lot of total cost of both fuel and electricity consumed by the plant	5.75	4.61	4.28
Multi-Unit Status	=1 if the plant belongs to a multi-unit firm	.71	.52	.52
HQ Status	=1 if the plant is a headquarters; equal to 1 for all single-unit firms	.49	.63	N/A
Exporter Status	=1 if the value of exports is >0	.46	.34	.25
E-Com Dummy	=1 if the value of e-commerce at the plant is >0	.61	.54	.44
Age	Derived from the year first observed in the LBD. Truncated at 44 due to LBD starting in 1976	24.9	21.3	N/A
Log IT Capital	Log of value of hardware and software stocks at the plant, calculated using the perpetual	3.83	2.89	N/A

Stock	inventory method and using BEA deflators.			
Structured Management Z-Score	Index created by summing up the normalized values from questions 1, 5, 8, 9, 11, 13 & 14 of the MOPS	.61	.54	N/A
% Employees with BA	% of managers and non-managers with bachelor's degrees at the plant	.08	.07	N/A
Number of sites	Number of establishments belonging to the parent firm	107.22	60.95	N/A
Number of products	Number of products (at the 7-digit NAICS code level) produced at the plant.	3.55	3.47	3.15

Table A2. Pairwise Correlations

	Log VA	DDD	IT Stock	Mgmt	% BA Degree	Exp	Log TE	Log K stock	Energy	MU	# Sites	# Prod	Age
Log Value-Added	1												
DDD	.294	1											
IT capital stock	.561	.183	1										
Structured Management	.447	.327	.276	1									
% BA Degree	.135	.078	.115	.119	1								
Exporter	.344	.118	.236	.205	.108	1							
Log Total Employment	.821	.266	.581	.407	.094	.324	1						
Log K stock	.725	.272	.497	.393	.110	.261	.675	1					
Energy	.763	.292	.431	.426	.076	.269	.715	.735	1				
Multi-Unit Status	.336	.211	.141	.357	.026	.103	.253	.321	.386	1			
# Sites	.134	.106	.060	.131	.023	-.011	.094	.125	.138	.195	1		
# Products produced	.116	.026	.095	.020	.021	.103	.120	.075	.072	.030	.033	1	
Establishment Age	.297	.082	.191	.099	.011	.163	.333	.262	.300	.039	.045	.109	1

NOTE: All correlations are significant at least the 5% level except the correlation between % with degree and establishment age