

NBER WORKING PAPER SERIES

ORGANIZATIONAL RESPONSES TO PRODUCT CYCLES

Achyuta Adhvaryu
Vittorio Bassi
Anant Nyshadham
Jorge A. Tamayo
Nicolas Torres

Working Paper 31582
<http://www.nber.org/papers/w31582>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2023

We would like to thank Ricardo Alonso and Kate Whitefoot for insightful discussions of the paper. We also thank Daron Acemoglu, David Atkin, Augustin Bergeron, Luis Cabral, Robert Clark, Wouter Dessein, Florian Englmaier, Miguel Espinosa, Ben Friedrich, Luis Garicano, Deepak Hedge, Mitch Hoffman, Larry Katz, Amit Khandelwal, Jin Li, Rob Metcalfe, Tommaso Porzio, Andrea Prat, Esteban Rossi-Hansberg, Raffaella Sadun, Jesse Shapiro, Chris Stanton, Eric Verhoogen, Tom Wollmann, as well as seminar and conference participants at MIT, HBS, HKU Business School, NYU, UCSC, the 2023 NBER Organizational Economics Spring Meeting, the 2023 NBER Summer Institute (Personnel group), the 2023 Barcelona Summer Forum, and the 2022 IPA-CDEP Entrepreneurship and Private Sector Development Working Group Meeting for useful comments. All errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w31582>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Achyuta Adhvaryu, Vittorio Bassi, Anant Nyshadham, Jorge A. Tamayo, and Nicolas Torres. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Organizational Responses to Product Cycles

Achyuta Adhvaryu, Vittorio Bassi, Anant Nyshadham, Jorge A. Tamayo, and Nicolas Torres
NBER Working Paper No. 31582

August 2023

JEL No. D22,M12,M53,O3

ABSTRACT

We use daily administrative data from a leading automobile manufacturer to study the organizational impacts of introducing new models to the auto assembly line. We first show that costly defects per vehicle spike when new models are introduced. As a response, the firm trains in problem-solving skills and promotes lower- and mid-level employees to solve the more complex problems that arise, thus moving to a less pyramidal knowledge hierarchy with fewer layers. We develop an extension to the classic theory of knowledge-based hierarchies that reconciles our novel empirical results by allowing the firm to also invest in its training resources.

Achyuta Adhvaryu
School of Global Policy and Strategy
University of California, San Diego
La Jolla, CA 92093
and NBER
aadhvayu@ucsd.edu

Vittorio Bassi
University of Southern California
Department of Economics
3620 South Vermont Avenue, Kaprielian Hall
Los Angeles, CA 90089
and NBER
vbassi@usc.edu

Anant Nyshadham
Stephen M. Ross School of Business
University of Michigan
701 Tappan Street
Ann Arbor, MI 48109-1234
and NBER
nyshadha@umich.edu

Jorge A. Tamayo
Harvard Business School
Soldiers Field, Morgan Hall 292
Boston, MA 02163
jtamayo@hbs.edu

Nicolas Torres
Good Business Lab
nicolas.t@goodbusinesslab.org

A data appendix is available at <http://www.nber.org/data-appendix/w31582>

1 Introduction

Industrial innovation is a key driver of economic growth (Grossman and Helpman, 1991b). Identifying which firm-level factors enable or constrain quality upgrading and process improvement is thus critical to understand why some firms increase their productivity and grow while others lag behind. As summarized in a recent review by Verhoogen (2021), studies of firm-level upgrading have emphasized the importance of several factors enabling the production of higher quality goods, such as plant size and quality of inputs (Kugler and Verhoogen, 2012), competition in input markets (Amiti and Khandelwal, 2013; Fan et al., 2015), access to export markets (Atkin et al., 2017b; Verhoogen, 2008), managerial skill and know-how (Adhvaryu et al., 2019; Bandiera et al., 2020; Bloom and Van Reenen, 2007, 2010), and knowledge spillovers (Bai et al., 2021).

One area which has received relatively little attention, however, is the study of the personnel and organizational adjustment costs firms face when upgrading. That is, how does the internal organization of the firm respond to the introduction of new products to enable upgrading in practice?¹ This question is particularly important in industries exhibiting product cycles, where product innovation can be frequent and relentless, as firms develop new generations of products to remain profitable when the production of current generations becomes imitable (Bayus, 1994; Grossman and Helpman, 1991a,b, 1994; Krugman, 1979; Vernon, 1966). Product cycles may demand substantial adaptation from upstream suppliers or subsidiaries, often in developing countries, who must deal with new and increasingly complex parts and processes.

This paper studies how firms manage production changes arising from the product cycle. We focus on the production of new models of automobiles, a prototypical example of product cycles. Leveraging daily administrative data from a leading global auto manufacturer and using event-study and discontinuity-based methods, our core contribution is to trace out the organizational responses to product cycles. We focus on two main ways in which firms may respond. First, they may increase the knowledge of their workers to manage these changes through training provision: a long literature has studied the degree to which firms engage in on-the-job training (Acemoglu and Pischke, 1998, 1999; Becker, 1962), with recent empirical studies identifying large potential returns to such training (Adhvaryu et al., 2018; Espinosa and Stanton, 2022; Sandvik et al., 2020). Second, firms may change their internal organizational structure to facilitate problem solving: a parallel literature has modeled and

¹A related literature studies the role of employee incentives and relational contracts within the firm (Amodio and Martinez-Carrasco (2018); Atkin et al. (2017a); de Rochambeau (2017); Kelley et al. (2021)), as well as the role of discrimination within the firm and biased beliefs on employee skills (Ghosh (2022); Hjort (2014); Macchiavello et al. (2020)) as organizational barriers to upgrading and productivity.

empirically tested predictions on how the distribution of knowledge within the firm might change as the firm re-optimizes under new production goals (Caliendo et al., 2020, 2015; Caliendo and Rossi-Hansberg, 2012; Garicano, 2000; Garicano and Rossi-Hansberg, 2006).

Our data comes from an Argentinian subsidiary plant of the auto manufacturer. We begin by showing that demand (planned number of vehicles produced) and number of total parts do not change in the short term after a new model is introduced; the main change is a large, discontinuous increase in *new* parts: after a model change, about 13-15% of the parts that need to be assembled are new. The production of new models thus necessitates the solving of novel problems that arise when first interacting with new parts. For example, if the buttons on the dash console or the front bumper for the new model are of a different shape or material than for the previous model, the angle and force which were used when attaching these parts in the previous model may no longer be appropriate. But the factory faces great difficulty in predicting which parts or operations may encounter these problems, let alone what the appropriate solution may be. As a result, we show that defects per vehicle (DPV) spike immediately by 2 standard deviations after the model change on average and up to 5 standard deviations after those involving a higher share of new parts. The factory is able to bring DPV back down to their prior level over approximately a three-week period on average, though it takes more than twice as long after model changes involving more new parts. This increase in defects created by new product introduction, though it is temporary, is indeed costly to the firm, as a day in which DPV is even just 1 SD over the mean results in an unrecoverable reduction in cars produced of nearly 20%.

We next ask how changes in the firm's organizational decision-making facilitate the problem-solving required to bring defects back down to a minimum in such a short time. Using granular data on training completed by all workers and on the hierarchical structure of each working group, we find that when a new model is introduced on the assembly line, the firm upskills lower and mid-level workers in problem-solving and communication skills via internal training, which increases the number of middle and high level managers in the working group who are able to solve the new problems arising on the production line. As a result, the aggregate level of skill capital within each group, and accordingly the firm, increases. The shape of working groups changes from markedly pyramidal to more rectangular, with the ratio of frontline workers to middle and upper level managers (i.e., what is referred to as the problem-solvers' span of control in Garicano (2000)) dropping significantly by .5 from a mean of 4 or 5.

Working groups also become flatter with less hierarchical knowledge layers. The firm delineates 16 possible knowledge layers reflecting the amount of cumulative training the worker has accrued. A working group has on average 20 workers spread across 6-7 of these

layers. Following a model change, we find a significant reduction in the number of hierarchical layers within working groups, which decreases by about 1 layer in the first three weeks after the model change. The firm starts to back-fill lower-level positions after defect rates recover to pre-model-change levels such that the number of hierarchical knowledge layers in working groups on average reverts, but working groups remain higher skill and more rectangular.

Finally, we use novel data from the firm’s continuous improvement system to document impacts on problem-solving activities of workers at all levels of the working groups. This system allows workers to report solutions they found to production problems to upper management for the purpose of recognition and bonus pay as well as to proliferate solutions which may have broader applications. We find, indeed, that in line with the removal of one layer of the knowledge hierarchy, the increased training in problem-solving skills, and the reduction in the span of control of mid and upper level workers we observe following the introduction of a new model, the number of reports of solutions to problems from lower and mid-level workers to upper management increases more than two to three-fold, both in total and on a per employee basis. In sum, we document that both changes in the amount of problem-solving knowledge and its distribution in teams by way of their organizational structure are important levers used by the firm to deal with the spike in new problems after model changes, bringing defects down to pre-change levels fairly quickly.

We then contrast these results with the impacts of increases in quantity produced of the same model, for which we show the organizational response is a monotonic, permanent increase in both employment and knowledge layers, consistent with prior evidence from manufacturing in high-income countries (Caliendo et al., 2020, 2015; Friedrich, 2022). Importantly, the increase in knowledge layers is not accompanied by an increase in average skill capital through training when the change involves producing more of the same model. Rather, consistent with problems arising from increased volume being routine and easily solved by frontline workers, we see that the pyramidal shape of working groups becomes more pronounced with their frontline bases growing; while middle and upper-level managers are shed to new working groups. Accordingly, average skill in working groups drops as does the need to elicit help from farther up the chain. Employment grows less than proportionately with output, as cars per worker increase due to economies of scale. Finally, consistent with the additional layer of problem-solvers in working groups and more pronounced pyramidal structure (i.e., increase in span of control) following the volume change, we find that reports of solutions to problems from lower and mid-level workers to upper management in the continuous improvement system decrease significantly.

We interpret our empirical results in light of the canonical models of knowledge hierarchies within the firm of Garicano (2000) and Caliendo and Rossi-Hansberg (2012). Exactly as

in our results, as well as the empirical results of [Caliendo et al. \(2015\)](#) and [Caliendo et al. \(2020\)](#), the canonical theory predicts that an increase in the volume produced of the same model should result in an expansion of the hierarchy through the addition of frontline workers and more knowledge layers: as production expands, adding layers allows the firm to focus the scarce time of skilled managers on solving only the most complex problems. However, this canonical theory does not predict the reduction in knowledge layers that we observe in response to quality upgrading (i.e., the introduction of new models). In the canonical model, the increase in problem complexity resulting from model changes should yield a reduction in the ratio of frontline operators to problem-solvers we find (i.e., the problem-solvers' span of control), but would not yield a reduction in hierarchical layers, but rather an increase in layers if anything, as each worker can solve a smaller fraction of the problems (which are now more complex and/or frequent).

We note, however, that both [Garicano \(2000\)](#) and [Caliendo and Rossi-Hansberg \(2012\)](#) explore comparative statics with respect to both problem complexity or frequency and cost of acquiring knowledge, showing they move the choice of hierarchical knowledge layers in opposite directions. Leveraging novel data on knowledge (i.e., completed trainings) of all workers across layers, our results show that the firm chooses the cost of knowledge acquisition (via costly investments in training efficacy) endogenously in response to the problem complexity or frequency it faces. That is, we find that when a new model is introduced, the firm invests in improving its in-house training programs by adding courses and bringing in new trainers. Accordingly, we modify the canonical theory in line with this empirical observation to allow firms to make a costly investment in increasing the efficacy of training. That is, training becomes another lever the firm adjusts alongside the distribution of knowledge across layers.

We show that when quantity is fixed in the short-run, as both our empirical results and anecdotes from the firm's upper management confirm, the firm will reduce the number of knowledge layers as a response to the increase in complexity or frequency of problems arising from the introduction of new models. Intuitively, as the firm can train its workers more efficiently in problem solving, it can have more managers farther up the hierarchy able to solve more problems, with fewer knowledge layers needed to produce. We show that allowing for endogenous investment in the productivity of training can reconcile our empirical results on the impact of model changes on the internal organization of the firm. Indeed, the increases in problem-solving activities we see at both the lower and mid-levels of the working group reflect a combination of the results from comparative statics explored by [Garicano \(2000\)](#) with respect to an increase in problem complexity and a decrease in the cost of acquiring knowledge. Importantly, we show that volume changes still lead to an increase in the number of knowledge layers in this extended version of the theory, consistent with both our findings

and previous results in the literature. Taken together, when quantity is fixed in the short run, the firm treats investments in training efficacy and the choice of layers as substitutes, but when quantity can adjust in the long run these levers are complementary.

Related literature and contribution. We aim to contribute to the literature documenting the importance of management and organizational design in determining the productivity and growth of firms (Aghion et al., 2014; Bloom et al., 2013, 2010, 2016; Bloom and Van Reenen, 2007, 2010; Frederiksen et al., 2020; Hoffman and Tadelis, 2021; Metcalfe et al., 2023; Minni, 2022). In particular, we focus on the distribution of knowledge in teams and organizational flexibility in dealing with external stimuli.² Previous work in this area has focused quite fruitfully on the effects of demand shocks (Caliendo et al., 2020, 2015; Garicano, 2000). We build on the elegant theory developed in this and earlier work (most directly, the foundational work of Garicano (2000) and Caliendo and Rossi-Hansberg (2012)) to demonstrate how the quality upgrading in product cycles – and the increased complexity and/or frequency of problems that results – generates quite clear, and distinct, changes to the knowledge hierarchy. We use granular data (at the shift-by-day as opposed to yearly level as in previous work) to document the surprising flexibility of organizational structure in response to product cycles. We uncover that organizational structure can evolve very rapidly in response to the introduction of new models, and note that these substantial changes undertaken by the firm would be overlooked when relying on less granular data if, for example, the firm experienced both model and volume changes between observations. We build an extension to the canonical theory that allows us to reconcile the theory’s predictions with the empirical patterns observed following product quality upgrading in product cycles.

Our study also adds complementary evidence to prior empirical studies of how organizations achieve volume changes in response to demand shocks (Caliendo et al., 2020, 2015; Friedrich, 2022). Specifically, though on-the-job knowledge acquisition of workers across knowledge layers in teams and the distribution of problem-solving responsibilities across these layers are at the core of the original Garicano (2000) model, empirical evidence of impacts on workers’ stock of problem-solving knowledge across the hierarchy and their actual problem-solving activities is limited by data availability. We use granular data on training completion of all workers (including the degree to which course content covers skills relevant for communicating about and solving problems), team composition, and records of individual problem-solving activities from the continuous improvement system to fill these gaps.

In documenting the role of firm-provided training in dealing with product cycles, we also

²A related literature studies the allocation of workers to managers within organizations (Adhvaryu et al., 2020; Fenizia, 2022; Limodio, 2021).

contribute to a large literature on the returns to on-the-job training within organizations. This literature has highlighted that training provision to workers can have high returns for the organization (Adhvaryu et al., 2018; Espinosa and Stanton, 2022; Hoffman and Burks, 2017; Sandvik et al., 2020), as well as that firms may face challenges to allocating training efficiently within the organization (Adhvaryu et al., 2022; Sandvik et al., 2022). Our contribution is to show how training enables firms to deal with product innovation and quality upgrading, and to highlight its interaction with changes in organizational structure (or how knowledge is distributed) in helping firms achieve product quality upgrading. In studying the stock and distribution of knowledge within the organization, our paper also relates to the literature studying knowledge spillovers between co-workers (Guillouet et al., 2021; Jarosch et al., 2021) and between workers at firms located near each other (Atkin et al., 2022).

We also add to the understanding of the personnel and organizational adjustment costs involved in product quality upgrading. Much of the literature on upgrading in quality, technology, and/or product innovation has emphasized the roles of drivers on both the output and input sides, as well as firm knowledge in the forms of managerial skill and learning (see Verhoogen (2021) for a thorough review). Less attention has been paid to how suppliers are able to adapt to consumer-demand-driven product quality upgrading; this is particularly relevant in industries exhibiting product cycles, in which product innovation is frequent and relentless (Grossman and Helpman, 1991a,b, 1994). In nearly all industries in which product cycles feature heavily, consumer-facing firms either contract with suppliers or have their own operations in the “global south.” This means that the firms that actually deal with the continual increase in complexity generated by innovation are different from the firms actually doing the R&D. Understanding the behavior of these suppliers and their ability to adapt to product cycles is thus of first-order interest to understand the global patterns of trade that are created by continual innovation.

Finally, our results documenting investments in knowledge, particularly problem-solving skills, via training and how this knowledge is distributed within teams contribute to the learning by doing patterns that have long been studied in economics (see, e.g., Adhvaryu et al. (2019); Arrow (1962); Irwin and Klenow (1994); Thompson (2010), among numerous others). Our study is closest to Levitt et al. (2013) who also study an automobile assembly plant and document similar spikes in DPV at the start of production of a new model in a team-based approach and subsequent reductions over time. They show that these learned improvements are not held by the workers themselves but rather in the form of “organizational capital”. Specifically, they describe a whiteboards system in which frontline operators reported problems to higher knowledge level workers (i.e., managers and quality control engineers) who designed and communicated solutions. Our results complement these findings in providing

novel empirical evidence of both investments in problem-solving skills along the hierarchy and changes in the shape of teams, both in terms of number of hierarchical knowledge layers and whether the distribution of workers across these layers yields more pyramidal or rectangular shaped teams. We interpret the training in problem-solving skills and changes to the shape of teams, as well as the solutions to problems that they enable, as specific examples of the “organizational capital” which [Levitt et al. \(2013\)](#) argue embodies the gains from learning.

Structure of the paper. The rest of the paper is organized as follows: Section 2 describes the setting of our study and the data used for estimation. In Section 3 we describe the empirical strategy, and Section 4 shows the results on the organizational response to discrete changes in both the models and the quantity of cars produced. In Section 5, we introduce a model to reconcile our results with standard models of knowledge hierarchies within the firm. Section 6 concludes. Additional details are in the Online Appendix.

2 Context, Data and Descriptives

In this section, we describe the setting where our study takes place and the data used for estimation.

2.1 Context, Organization of Production, and Product Cycles

We partnered with a leading global auto manufacturer subsidiary plant located in Argentina. The plant began operating in the 1990s and now produces more than 140,000 cars per year, employing over 3,400 workers. Around two-thirds of the production is exported to different countries in Latin America, and the rest is dedicated to the local market.

2.1.1 Production Process

Production takes place in a production line setting. The production line is made up of eight sectors: Press, Welding, Painting, Frame & RX Axle, Engines, Resin, Assembly, and Quality Check. The different parts of a production unit (i.e., a car) are produced in parallel by different sectors of the production line: chassis and car-body components are manufactured and connected by Press and Welding respectively, the rest of the parts are manufactured by Frame & RX Axle, Engines, and Resin. Finally, these different parts are all assembled together in the Assembly sector. Assembly is one of the most delicate phases of production, with 75% of the defects per vehicle occurring at this stage of the production process. Our

study focuses on the Assembly sector, both because it is where the majority of defects occur and because we are able to track working group composition for this department.

2.1.2 Organization of Labor, Knowledge Hierarchies and Training

Production line workers are divided into *working groups* within each sector of the line. The plant has two *shifts* per day (morning and afternoon) with each working group operating in one of the two daily shifts. Employees within each working group are organized in a clear production hierarchy, characterized by different knowledge layers. Workers are assigned to the different layers depending on the tasks they perform and the skills they have accumulated. Employees can accumulate skills by receiving in-house training programs provided by the firm. A diagram of the hierarchical structure of working groups is shown in Table 1.

There are three main employee categories within the hierarchy. The first are Front-line Operators (FL), who engage in production tasks on the production line. Above FL workers are Mid-line Operators (ML), who are tasked with identifying and solving production problems and closely supervising and helping FL workers whenever problems or bottlenecks occur. Finally, above ML workers are Superiors (S), who are in charge of supporting and managing the overall working group performance.

If an FL worker faces an issue with the assembly process, one of the ML workers makes a quick diagnosis of the issue and offers a solution to the FL worker. In case the issue exceeds the knowledge of the ML worker, the problem is then directed to the Superior.³ Accordingly, the hierarchical structure is pyramidal in shape. The typical working group has about 20 workers, divided into about 16 FL workers, 3 ML workers and 1 S worker.

Within each main employee category (i.e., within the FL, ML and S categories), employees are then divided into different *layers*, depending on their skills and training received. FL workers are divided into three possible layers of New Front-line Operators (New FL), four layers of Regular Front-line Operators (Reg FL), and one layer of Mid Front-line Operators (Mid FL). New FL are the lowest level of the hierarchy, comprising workers who have recently started and have received only basic training. Reg FL workers are more experienced and skilled, having typically spent at least a year on the production line and having completed at least five training programs, and Mid FL are the highest skilled within the FL category, having more than two years of experience and having gone through 13-14 training programs. ML and S workers are subdivided into four layers each, again based on their training. So, in total, there are 16 possible employee layers that a working group can have – eight within FL workers, four within ML workers, and four within S workers. Each of these 16 layers is associated with a different level of training and therefore knowledge acquired by the workers

³Superiors are supervised by Line Superiors, who are tasked with overseeing multiple working groups.

in that layer. A given working group does not necessarily have an employee from each of the 16 layers.

Most new employees start as FL workers with the lowest levels of basic skills (i.e., they start in the first layer as New FL Operator 1). Workers are promoted as they gain experience, complete specific training programs, and develop new job skills required for each layer, as shown in Table 1. For instance, as the New FL gain experience and skills they can get promoted first to the higher layers within the New FL category (e.g., from New FL Operator 1 to New FL Operator 2), and then across categories (e.g., from New FL Operator 3 to Reg. FL Operator 1) all the way up the hierarchy.

The firm provides training to its employees; and through such training, employees can be promoted to higher positions in the working group and advance their professional careers. Training is provided in-house through training programs administered by trainers and is primarily focused on management, problem solving and alignment of production line activities with company targets, rather than on technical skills to operate machinery on the production line.

More advanced training programs place a greater emphasis on problem solving and management. Specifically, 78% of the courses required for FL workers to be promoted to ML positions are related to identifying problems, coming up with solutions, and coordinating production to minimize waste and slow-downs. 91% of the courses required to advance from lower ML to upper ML positions are related to these topics; while 68% of the courses required to advance from ML to S positions cover these topics, with the remaining courses covering more managerial topics.⁴

2.1.3 Problem-Solving in Working Groups and Continuous Improvement System

The assembly line operates almost entirely sequentially such that issues at one process or station along the line will lead to stoppages of the entire line, slowing the flow of completed vehicles out of the end of the line. At a given set “takt” time (i.e., the planned production time), the line would generally produce a fixed number of cars each day or shift. The major variation in output or productivity (i.e., cars completed per unit time) comes from these stoppages due to problems and any defects they cause.

Accordingly, each working group’s ability to solve problems quickly is critical. Indeed, the pyramidal hierarchical structure of the working groups and the bulk of the content of the training together emphasize the importance of problem-solving in this context. Though problems may arise even during business as usual operations, from machine malfunctions

⁴These percentages reflect the share of courses with this content, but share of training hours yields similar magnitudes.

Table 1: Hierarchical Structure of Working Groups

	Employee layer	Tasks/Responsibilities	Cumulative training
S	Superior 4	Support and manage overall working group performance	17 to 20 training programs
	Superior 3		
	Superior 2		
	Superior 1		
ML	Mid-line Operator 4	Identify production problems; closely supervise front-line workers to ensure production standards are met	14 to 17 training programs
	Mid-line Operator 3		
	Mid-line Operator 2		
	Mid-line Operator 1		
	Mid Front-line Operator	Engage in production tasks; support other front-line operators	13 to 14 training programs
FL	Reg. Front-line Operator 4	Engage in production tasks	12 to 13 training programs
	Reg. Front-line Operator 3		10 to 12 training programs
	Reg. Front-line Operator 2		8 to 10 training programs
	Reg. Front-line Operator 1		5 to 8 training programs
	New Front-line Operator 3	4 to 5 training programs	
	New Front-line Operator 2	Engage in production tasks; learn basic skills	1 to 4 training programs
	New Front-line Operator 1		0 to 1 training programs

Note: Table 1 shows a diagram of the hierarchical structure of working groups. For each employee layer, the diagram shows (i) typical tasks and responsibilities; and (ii) the cumulative number of training programs typically received by workers in a given layer.

for example, problems of the greatest complexity (or novelty) and in the greatest frequency occur when the product changes and when the takt time or planned production is adjusted.

Given the importance of diagnosing and solving problems quickly throughout the year including at these times of production changes, the firm has in place a continuous improvement system in which workers are encouraged and even incentivized (both monetarily and by way of recognition) to record solutions they find to problems, with these records reviewed by upper management. The goal is both to elicit maximal effort from workers as well as to leverage opportunities to proliferate solutions which may apply more broadly across the production process. As we discuss further below, these records prove invaluable in our analysis as they allow us to study impacts on how many problems are being solved in each working group and by whom (i.e., from which knowledge layers).

Note that: (i) the structure of the working groups; (ii) the specific knowledge that higher levels obtain via required training for promotion; (iii) the fact that productivity is increased mainly by minimizing stoppages and defects from production problems; (iv) and the firm’s clear valuing of problem-solving all underscore the link between production settings of this sort and the intuition at the heart of the knowledge hierarchies models of [Garicano \(2000\)](#) and [Caliendo and Rossi-Hansberg \(2012\)](#). The features of this context, however, are not unique to this firm. These processes and systems are ubiquitous across the automobile manufacturing industry, and indeed make up the core of the mechanism for organizational learning presented

in [Levitt et al. \(2013\)](#). Furthermore, lean manufacturing and continuous improvement are regarded as frontier production practices across most manufacturing contexts. For example, [Adhvaryu et al. \(2023\)](#) study this same system in electronics manufacturing factories in Thailand. Relatedly, communication and problem-solving skills of both frontline operators and production team supervisors contribute greatly to productivity in Indian garment factories ([Adhvaryu et al., 2018, 2022, 2019](#)).

2.1.4 Product Quality Upgrading and Changes in Production Volume

In our analysis, we exploit two types of discrete changes to the production process. First, we exploit the fact that the product cycle is characterized by the frequent changes in model variants over time. Over our sample period, the company manufactures two auto models, which are continually updated through new model variants, usually every two or three years. Most changes tend to be made to the bumpers, lights, grille, wheels, and color options. As we will document in our data, such model changes result in a significant increase in the number of new parts that have to be produced and assembled, with resulting changes in the types and frequency of problems that arise on the assembly line.

One important point that was emphasized in discussions with factory management, and is reflected in the event study plots we present below, is the inability to predict which operations will face problems when the new parts for the new model are introduced. For example, the bumper on the new model may appear to be nearly identical to the old one, but if it is of a slightly different shape or weight or made of slightly different material (or even just purchased from a different supplier) the angle or force used when installing the bumper may require adjustment. Accordingly, working group Superiors cannot warn operators in advance of the problem let alone suggest an appropriate solution. Rather, the working groups must rely on the problem-solving skills that have been imparted to high level FL, ML, and S members of the team to diagnose the causes of defects (like cracked or crooked bumpers) only after they arise, and to suggest and implement solutions in a timely manner.

On a tour of the factory one such example was highlighted: a bumper attachment was being performed with minimal to no defects for the prior model, but when the new model was introduced the new bumper was heavier and attached slightly differently to account for the weight. After the switch to the new model, a high proportion of bumpers were being attached crooked or otherwise imperfectly. That this process would be one that would produce defects could not easily be predicted by any of the team members as they had never tried to attach this particular part before. In this particular instance, an ML in the working group diagnosed that the root cause was the new weight and angle such that having to hold the bumper in place while attaching it was not as easy for this new model as it had been for the prior model.

They suggested a hydraulic lift be used to hold the bumper at the correct height so that the worker could focus their time and attention on simply attaching it correctly. This simple adjustment allowed the operation to be accomplished once again with minimal to no defects.

Second, we exploit the fact that the plant expanded its production volume from around 90,000 cars per year (at a takt time of roughly 145 seconds) to around 140,000 (at a takt time of 90 seconds) in seven years, between 2011 and 2018. This expansion occurred progressively through a sequence of discrete “jumps” in the number of cars that the plant was asked to produce by the manufacturer. Such volume changes result in a substantial increase in the number of cars that have to be produced, but not in the complexity of what needs to be produced since the model does not also change at the same time.

However, despite the fact that no new parts or processes are introduced, new problems still arise; and once again, where these will arise is hard for management to predict. This is because some processes will have enough slackness to easily keep up with the new faster takt time, while others will have to be re-engineered to achieve the new pace (often by splitting the operation into multiple steps to be done by additional workers or even additional working groups). Which operations will be challenging and produce defects at the new pace is not easily predicted by any of the team members as they have never tried to perform the operation (or seen it performed) at the new pace before.

It is important to note that the plant has no discretion as to whether or when to implement model and volume changes: such decisions are made by the manufacturer. The plant is tasked only with executing production. Our interest is precisely in studying the organizational responses put in place by the plant as a response to such product cycles and volume changes.

2.2 Data and Descriptives

We use data from the Assembly sector as that is the most labor-intensive sector and accounts for most of the defects per vehicle in the plant. Our data comes from four main sources. The first is data on productivity, employment, and training received by employees at the shift-day level. The second is daily data on the composition of each working group, in terms of the number of employees by knowledge layer. The third is daily records of problems identified and solved by each employee. The fourth is a record of the exact dates when the introduction of new model variants and changes in the production volume took place.

2.2.1 Productivity, Employment and Training

We have access to daily productivity data for each shift in the Assembly sector from January 2012 to February 2019. The data include the number of cars produced per day for each model

and the number of defects per 100 vehicles produced (DPV). The DPV is a key performance indicator that the plant uses to monitor productivity for each shift-day.

Panel A of Table 2 presents summary statistics for our two key productivity variables of interest at the daily level.⁵ The plant is large, producing around 410 cars per day. However, there is significant variation in the number of cars produced; across days due to stoppage and defects and dramatically from the start of our observation period to the end due to several discontinuous drops in takt time. We standardize the number of defects per vehicle using the mean and standard deviation of the full sample, so that our DPV measure has mean 0 and standard deviation 1.⁶

Our data shows that the incidence of DPV is associated with a sizeable decrease in the number of cars produced per day, thus confirming that DPV is an appropriate measure of productivity in this context. More precisely, Appendix Figure A2 displays the coefficients of a Distributed Lag Model where we regress the number of cars produced per day on lags of an indicator variable equal to 1 if the plant experienced a DPV 1 standard deviation above the mean on a given day. The figure shows a contemporaneous decrease of roughly 50 cars the same day the plant experiences high DPV. The negative effect persists the next day with a loss of another 10-15 cars and fades quickly thereafter. Importantly, there is no evidence of a significant increase in number of cars produced in subsequent days, thus confirming that high DPV on a particular day is associated with an overall permanent (i.e., unrecoverable) reduction in the total number of cars produced.

We also have data on employment in the Assembly sector between January 2017 to February 2019. The data includes information on the date of hiring (and exit) from the plant, whether the worker is present at the factory on any given day, current position and historical data on previous positions, as well as completion date and course type of each training program received by each worker. Panel B of Table 2 presents basic summary statistics on the number of employees in Assembly, where information on employment is averaged at the weekly level to minimize measurement error related to short-term absences from the plant.⁷ There are on average 1,187 employees each week: 966 front-line employees, 170 mid-line employees and 51 superiors, confirming the pyramidal structure described above. The average

⁵Productivity data is originally at the day-shift level. To go to the daily level, we sum the number of cars and average the defects per vehicle across the two shifts within the same day.

⁶To protect the confidentiality of the plant, we were not allowed to disclose mean and standard deviation of the DPV measure. Appendix Figure A1 shows the distribution of (standardized) DPV-day observations, confirming that there is substantial variation in defects per vehicle across days.

⁷Note that such measurement error is not an issue with DPV measures as they are key daily KPIs for the plant and so are carefully recorded in the administrative data. Our data on employment instead records the number of employees present at the factory on a given day, and so is subject to potential measurement error driven by short-term absenteeism.

Table 2: Descriptive Statistics on Productivity, Employment and Training

	Mean	SD
<i>Panel A: Productivity, daily level</i>		
Number of Cars	413.32	110.86
Defects Per Vehicle (DPV)	0.00	1.00
Number of observations	1590	
<i>Panel B: Employment and training, weekly level</i>		
Number of Employees	1187.07	224.05
Front-line Employees (FL)	965.82	195.10
Mid-line Employees (ML)	169.81	26.73
Superior Employees (S)	51.44	3.37
Number of Working Groups	57.77	5.92
Avg Completed Training Programs	11.56	0.74
Avg Accumulated Training Hours	54.50	7.23
Number of observations	146	
<i>Panel C: Training courses, weekly level</i>		
Number of Courses	35.65	15.51
Number of Teachers	38.95	16.81
Number of observations	370	

Note: Panel A shows descriptive statistics on productivity at the daily level, Panel B and C on employment and training averaged at the weekly level, all for the Assembly sector. DPV is standardized using the mean and standard deviation of DPV of the full sample.

employee has received 11.6 training programs, totalling about 55 hours of company-sponsored on-the-job training.

Finally, we have data on the number of training courses provided by the plant each week for the entire sample period (January 2012 to February 2019). Panel C shows that the company engages in continual on-the-job training: in the average week the plant is carrying out 36 separate training courses taught by about 39 trainers, although again there is substantial week-to-week variation in the amount of training provided.

2.2.2 Composition of Working Groups

Our data allows us to track the exact composition of each working group for two years, between January 2017 to February 2019. That is, we know the precise layer each employee

belongs to (out of the possible 16 layers described in Table 1) and their accumulated training. This enables us to identify promotions across layers within each employee category (e.g., New FL 1 to New FL 2) and across categories (e.g., Mid FL Operator to ML Operator 1), as well as moves across working groups over this period. Our data on the production hierarchy within the firm is thus even more granular than the data used in recent papers studying production hierarchies in French and Portuguese firms (Caliendo et al., 2020, 2015).⁸

Table 3 shows descriptive statistics on the structure of working groups at the weekly level (so that the unit of observation is a working group in a week).⁹ The average size of each working group is about 20, which includes around 16 FL workers, 3 ML workers, and 1 S worker.¹⁰ We note that each of the 16 layers need not be, in fact are rarely, present in a given working group at a given time. On average there are about eight layers of skills observed across groups: five within FL, two within ML, and one within the S worker category.

Panel B reports the mean and standard deviation of the week on week change in the number of layers of each working group over the sample period. Working groups are not static: while the number of layers and ratio of FL to ML and S workers both tend to increase on average over the sample period, there is substantial variation in the evolution of the hierarchical structure of working groups. Working groups both add and remove layers and each layer can both expand and contract over our sample period. Consequently, it is precisely this organizational flexibility that we seek to analyze in the rest of the paper, exploiting our granular data on the production hierarchy of the firm and training provision.¹¹

⁸For instance, Caliendo et al. (2015) identify five separate hierarchical layers in their French data, and Caliendo et al. (2020) identify eight hierarchical layers in Portuguese data. In contrast, we identify 16 layers. Working with only one firm allows us to gather very granular data including employee-specific training and problem-solving as discussed more below, although of course this comes at the expense of not being able to analyze responses across many firms, as these other papers do.

⁹In the data, a working group is defined as a set of workers with the same Superior each day. We again average at the weekly level to smooth out measurement error arising from short-term worker absences.

¹⁰There are some periods in which Mid-line workers act as Superiors. Such cases are limited however to about 11 percent of working group-shift observations, which do not have a Superior.

¹¹In the Appendix we show the evolution of average size of working groups (Appendix Figure A3a), average number of layers in each working group (Appendix Figure A3b) and average number of working groups (Appendix Figure A3c) over our sample period. Working groups grow from less than 17 workers to roughly 22 workers on average. Correspondingly, the number of layers in each working group climbs by roughly 1.5; and the ratio of FL workers to ML and S workers grows from less than 4 to more than 5 on average. The number of working groups also grows from roughly 95 to more than 125 on average. Note that employee absenteeism and turnover contribute to creating some of the high frequency variation in these Appendix Figures. Over the same period (from January 2017 to February 2019), the takt time falls from 110 to 90 seconds, corresponding to more than 20% growth in annual planned production from 115 to 140 thousand cars per year. Our interest is in documenting systematic changes in these working-group level outcomes as a response to the discrete introduction of new model variants or increases in the volume of production (i.e., reductions in takt time) over this period.

Table 3: Descriptive Statistics on Working Groups and Employee Hierarchies

	Mean	SD
<i>Panel A: Composition, weekly level</i>		
Number of Employees	20.25	8.23
Number of FL Employees	16.47	7.56
Number of ML Employees	2.89	1.27
Number of S Employees	0.89	0.31
FL/(ML+S)	4.56	2.23
Number of Layers	7.80	1.99
Number of FL Layers	5.02	1.65
Number of ML Layers	1.89	0.76
Number of S Layers	0.89	0.31
Number of observations	13,280	
<i>Panel B: Change in layers, weekly level</i>		
Change in Number of Layers	0.015	0.349
Change in Number of FL Layers	0.011	0.303
Change in Number of ML Layers	0.003	0.141
Changes in Number of S Layers	0.002	0.035
Changes in FL/(ML+S)	0.005	0.050
Number of Observations	13,060	

Note: The information presented in Panel A is averaged at the working group-shift-week level. The changes in layers in Panel B are computed as the average change in the number of layers in the working group at the weekly level.

2.2.3 Problem-Solving Records from Continuous Improvement System

As discussed above, the firm places a strong emphasis on problem-solving skills, particularly at upper frontline and middle layers of working groups, and encourages workers to participate in problem-solving by way of recognition and monetary rewards. To implement these incentives, the firm has a continuous improvement system in which workers log any problem they solved. The firm not only leverages these records to reward workers for their efforts, but also looks for opportunities to proliferate solutions which may be more broadly applicable.

These records are ideal for an analysis of how problem-solving responsibilities evolve with training and group compositional changes, as we can compute the daily number of problems solved by each worker for 2018 and 2019, the period for which this data is available. Note that the ease of using the system to record problems and the clear incentives to do so yield a likely

complete record of the universe of problem-solving activities. Appendix Table A1 reports the share of reports of problems solved by each knowledge layer as well as the average reports per employee by layer. FLs contribute the largest share of problems-solved, particularly the most knowledgeable layer among FLs (FL4). However, accounting for the number of workers in each layer, we see that ML2s and ML3s solve the most problems per worker.

2.2.4 Model and Volume Changes

We observe seven discrete changes in the model variants produced during our sample period, as a result of product quality upgrading spurred by automotive industry product cycles. We call these “Model changes.” We also observe five discrete changes in production volume, as a result of the overall expansion of production and reduction in takt time described above. We call these “Volume changes.”

On average, with each Model change, 14% of the car parts are modified, so that the production line needs to deal with such new parts. This share of new parts can be as high as 88% for major generational model changes, and we use this heterogeneity below in our analysis. However, the number of cars the plant plans to produce does not change. In contrast, Volume changes represent a discrete increase in the number of cars the factory plans to produce, which jumps up by 16% on average, again with substantial heterogeneity in the magnitude of the increase which we exploit in our analysis below. However, Volume changes do not bring about new parts, as the model that is produced does not get updated. As such, the production line needs to deal with more of the same tasks. Additional summary statistics on the Model and Volume changes that we exploit are reported in Appendix Table A2.¹²

Model and Volume changes are therefore very different in nature: Model changes require dealing with new parts, but without a change in overall quantity produced; Volume changes require producing a higher quantity of the same product. Our data allows us to study organizational responses to both types of events within the same context.

3 Empirical Strategy

Our aim is to study how training provision and the internal organization of the firm change in response to the Model and Volume changes described above. We begin the empirical analysis by showing how Model changes impact the share of new parts and, as a result, productivity (measured by the occurrence of DPV). We then study how the stock of knowledge and the shape of working groups change in response. Then, we contrast the impacts of Model changes

¹²Two of the five Model changes and two of the seven Volume changes take place in the period between January 2017 to February 2019 when we have data on employment and the composition of working groups.

to those of Volume changes to highlight the different effects of needing to deal with new tasks, as opposed to needing to conduct more of the same tasks. In doing so, we highlight the unique implications of product quality upgrading for the internal organization of the firm.

To answer these questions, we exploit the high-frequency nature of our data to implement event studies and regression discontinuity specifications around the exact time of such changes. The event studies allow us to trace out the dynamics of the impacts. The discontinuity in time specifications allow us to provide a useful summary point estimate of the impacts averaged over the time period under consideration. Next, we describe both sets of specifications.

3.1 Event Studies

For our key productivity outcome of DPV, which is available at the level of each shift in each day, we estimate the following specification:

$$Y_{st} = \theta_s + \gamma_m + \gamma_y + \sum_{k \geq -b, k \leq b, k \neq -1} \beta_k D_t^k + \delta X_t + \epsilon_{st} \quad (1)$$

where Y_{st} is the outcome measured in shift s and day t . D_t^k is an indicator of the distance with respect to the (Model or Volume) change, where k indicates the distance to the event on day t . In our dataset we know the exact day when the event takes place. b represents the bandwidth (window of time) considered in the event study. Following [Calonico et al. \(2014\)](#) we find that the optimal bandwidth for Model changes when using the DPV measure as the outcome is 40 days (8 five-day working weeks) before and after the event. We use this bandwidth throughout for both Model and Volume changes and so set $b = 40$.¹³ Our specification controls for shift fixed effects (θ_s) and year and month fixed effects (γ_m and γ_y) to net out any shift-specific effects and to control flexibly for time trends and seasonality. Additionally, we control for linear distance to other events (e.g., when looking at a given Model change, we control for distance to all other Model and Volume changes) and a linear time trend (X_t). Finally, standard errors are clustered by week-shift level.

For the outcomes at the level of individual working groups (e.g., composition of the working groups), we estimate a specification similar to equation 1 but with information at the working group-week level:

$$Y_{iw} = \theta_i + \gamma_m + \gamma_y + \sum_{k \geq -b, k \leq b, k \neq -1} \beta_k D_w^k + \delta X_w + \epsilon_{iw} \quad (2)$$

where Y_{iw} is the outcome of group i in week w . D_w^k is an indicator for the distance with

¹³We use the same bandwidth for Model and Volume changes for ease of comparison of the results across the two types of events. The optimal bandwidth is similar for Volume changes (25 days).

respect to the event, where k indicates the distance to the event at week w . We use the same optimal bandwidth as in equation 1, and so in this case $b = 8$ as this analysis is at the weekly level. The specification further controls for working group fixed effects (θ_i), month and year fixed effects (γ_m and γ_y), linear distance to other Model and Volume changes in weeks and a linear time trend (X_w). Standard errors are clustered at the week-working group level.

Our parameters of interest are the β_k . For $k \geq 0$ (i.e., for time periods corresponding to the time of the event or later) we interpret β_k as the causal effect of the Model or Volume change on the different outcomes at time k . All β_k parameters are relative to the outcomes in the time period immediately before the event $k = -1$, which is the excluded category in these regressions.

Identification relies on the assumption that, after controlling for the linear time trend and for month and year fixed effects, any discrete changes in the outcomes observed just after the event are not due to underlying (residual) trends in the outcome variables. The high-frequency nature of our data and our ability to focus on a narrow time window around the events reassures us regarding the validity of this assumption. The focus on a narrow time window reduces the possible concern that changes in the outcome variables after the Model or Volume changes are due to other events affecting trends in productivity or organizational structure (and we are controlling for distance from other Model and Volume changes). As described in Section 2, the plant has no discretion over whether to implement changes in models or volume; it simply has to execute production mandates handed down from the global brand headquarters. Of course, there is still the possibility that such changes are communicated to the firm in advance, which may lead to anticipation effects. The availability of high-frequency data before the event allows us to test for the presence of such possible anticipation effects. Lack of significance of the estimated β_k in the time periods before the event (i.e., $k < -1$) would provide evidence in support of the limited role of any anticipation effects or other trends in potentially biasing our results.

3.2 Regression Discontinuity in Time

In addition to the event study specifications described above, we also estimate regression discontinuities in time. Doing so allows us to estimate the average effect of the event in our time window, thus providing a useful summary point estimate. This method also has the advantage of pooling coefficients across time periods, thus improving precision.

For the DPV outcome, we estimate the following specification:

$$Y_{stw} = \theta_s + \gamma_m + \gamma_y + \beta_1 I(0 \leq dis_w \leq 3)_w + \beta_2 I(4 \leq dis_w \leq 7)_w + f(dis_t) + \delta X_t + \epsilon_{stw} \quad (3)$$

where Y_{stw} is the outcome measured in shift s , day t and week w . $I(0 \leq dis_w \leq 3)_w$ is an indicator for being four weeks (we define a week as five working days) after the event occurring at time t . $I(4 \leq dis_w \leq 7)_w$ is an indicator for being more than four weeks after the event at day t . Recall that the optimal bandwidth for the event studies was 40 days, which guides our definition of the indicator functions. We control for shift fixed effects (θ_s), month and year fixed effects (γ_m and γ_y). Additionally, we control for a function of the distance to the event ($f(dis_t)$) in days, where we choose the linear function, and for the linear distance in days to all other Model and Volume changes (X_t). Standard errors are clustered by distance to the event-shift level.

For working-group level outcomes, we estimate a specification similar to equation 3 but where the level of observation is a working group in a week:

$$Y_{iw} = \theta_i + \gamma_m + \gamma_y + \beta_1 I(0 \leq dis_w \leq 3)_w + \beta_2 I(4 \leq dis_w \leq 7)_w + f(dis_w) + \delta X_w + \epsilon_{iw} \quad (4)$$

where Y_{iw} is the outcome for working group i in week w . $I(0 \leq dis_w \leq 3)_w$ is an indicator for being within four weeks after the event at week w . $I(4 \leq dis_w \leq 7)_w$ is an indicator for being more than 4 weeks after the event. We control for working group fixed effects (θ_i), month and year fixed effects (γ_m and γ_y), and for a linear function of the distance to the event ($f(dis_w)$) as well as distance to all other events (X_w) in terms of weeks. Standard errors are clustered by distance to the event-working group level.

Our parameters of interest are β_1 and β_2 , which capture the average effect of each event in the two narrow windows of time. This allows us to speak to the dynamics of the Model and Volume changes, while still providing summary point estimates for the average effect of each event.

4 Results

We present the impacts of Model and Volume changes, in turn. For both, we first validate the nature of the event by studying impacts on planned number of cars to be produced and corresponding parts and share of new parts, when relevant. Then, we report impacts on defects per vehicle, to confirm that both types of changes have real implications for productivity, showing heterogeneity by share of new parts and number of new hires, respectively, to link defects to novelty. Next, we present impacts on knowledge (via training) and its distribution within working groups, as well as the shape of working groups, to study how the organizational responses the plant puts in place to bring down defects per vehicle differ across the two types

of changes. Finally, we study problem-solving activities at the worker level across knowledge layers in working groups to validate the interpretation of impacts as being driven by both stock and distribution of problem-solving knowledge and responsibility. The full set of Model changes results are summarized in a schematic in Appendix Figure A4; and the full set of Volume changes results are summarized in a schematic in Appendix Figure A5.

4.1 Product Cycles

4.1.1 Validating Model Changes

Appendix Table A3 reports the results of regressions like equation 3 but where the outcome variables are total number of cars planned for assembly per day, corresponding total number of parts per day, and the share of parts that are new relative to the last model variant produced just before the Model change.¹⁴ Columns 1 and 2 show that Model changes do not result in a significant change in the number of cars produced nor in the total number of parts used. The coefficients for both outcomes are small in magnitude relative to the mean. Column 3 instead shows that Model changes result in an increase of at least 14% in the share of parts that are new over the two months after the change. The increase in the share of new parts remains positive and significant persistently after the Model change; however the magnitude fluctuates due to variation in product mix.¹⁵ This confirms that Model changes result in an increase in complexity of production as a substantial amount of new parts need to be used to assemble the new model.

4.1.2 Impacts on Productivity

Figure 1 reports the results of an event study specification following equation 1 with (standardized) DPV as the outcome. The figure shows that there is a discrete jump in defects per vehicle right after the Model change. Defects increase by about 2 SD on average right after the Model change, and start to gradually come back down the next day, reverting back to the pre-shock level after about three-four weeks.¹⁶ The figure also shows that daily DPV does not exhibit any significant pre-trend in the period up until the *exact day* of the Model change.


¹⁴The level of observation in Table A3 is a day and the sample size is 567 because there are 7 Model changes and for each event we choose a bandwidth of 40 days around the event, so 81 days in total for each event.

¹⁵For each model, the factory makes several different trim levels. Higher level trims usually see a higher share of new parts to keep up with evolving tastes of customers willing to pay for luxury features. The share of production devoted to each trim can vary based on orders.

¹⁶Appendix Table A4 confirms the results in Figure 1 by reporting that Model changes lead to an increase in DPV of about 0.75 SD in the first four weeks after the shock, with the effect coming down to zero after that. This pattern holds, although point estimates are larger, when we restrict attention to the period 2017 to 2019 coinciding with the working group composition database. Appendix Table A5 confirms this.

This is consistent with the intuition shared by factory management that which specific parts and processes will face problems causing defects and stoppages is hard or even impossible to predict. As such, this pattern also confirms that any anticipation effects are not first order, thus validating the identification assumptions.

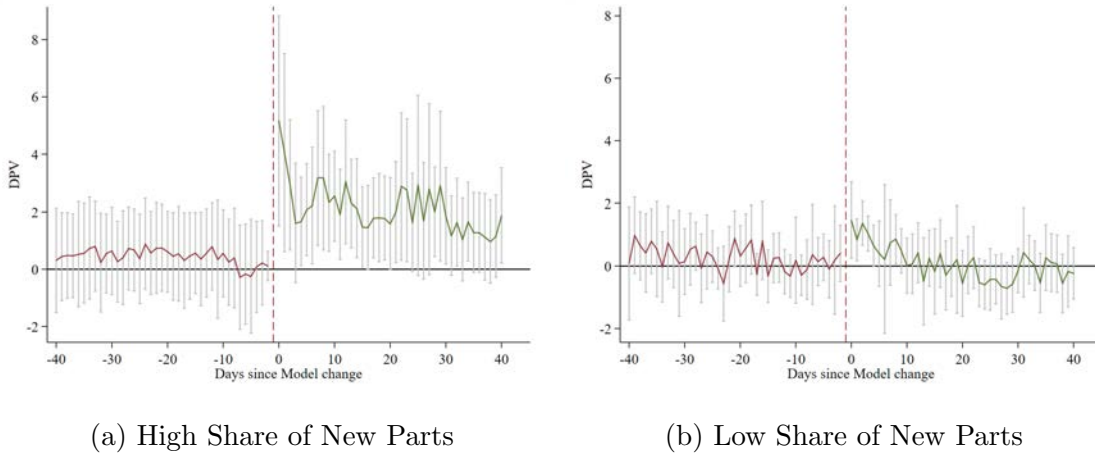
Figure 1: Event Study of Model Changes on Productivity



Note: Figure 1 shows the effect of Model changes on DPV in a time window running from 40 days before the event to 40 days after the event. DPV is computed as number of defects per 100 vehicles, and is standardized using the mean and standard deviation of the full sample. We control for month, year, and shift fixed effects. We also control for a linear function of distance to the Model change and to all other Model and Volume changes and a linear time trend. Standard errors are clustered by distance to the event-shift level. 95% confidence intervals are reported. Number of observations: 2 shifts x 81 days x 7 events.

Figure 2 presents the same event studies, but with the sample split between Model changes involving high (above median) and low share of new parts. The initial spike in DPV is much larger (roughly 5 SDs in magnitude) and the subsequent decay slower (DPV remains significantly elevated for more than 6 weeks) for changes involving the larger shares of new parts. This heterogeneity analysis lends strong support to the interpretation that the spike in DPV following model changes is driven by the novelty of many parts and corresponding processes.

Figure 2: Event Study of Model Changes on Productivity by Share of New Parts



Note: Figure 2 shows the effect of Model changes on DPV in a time window running from 40 days before the event to 40 days after the event split between changes with high share of new parts (above the median) and changes with low share of new parts (below the median). DPV is computed as number of defects per 100 vehicles, and is standardized using the mean and standard deviation of the full sample. We control for month, year, and shift fixed effects. We also control for a linear function of distance to the Model change and to all other Model and Volume changes and a linear time trend. Standard errors are clustered by distance to the event-shift level. 95% confidence intervals are reported. Number of observations for Panel (a): 2 shifts x 81 days x 3 events. Number of observations for Panel (b): 2 shifts x 81 days x 4 events.


The results in Figure 1 and 2 show that there are real negative productivity implications of the introduction of new model variants as a result of product quality upgrading. At the same time, the plant is able to bring these back down fairly quickly on average. Despite the temporary nature of the rise in DPV, the costs to the firm are substantial and unrecoverable. Appendix Figure A2 plots the coefficients from a Distributed Lag Regression model in which the number of cars produced each day is regressed on a dummy for whether the factory experienced DPV 1 SD above its mean for that day as well as 9 daily lags. We see that such a spike in DPV reduces output the same day by roughly 50 cars and a total of nearly 80 cars (or 20% from the sample mean of cars of 413) over the 10 day period, with no sign of any of that lost production being recovered. Note that in Figure 1 DPV stays elevated by 1 SD or more for more than a week following a Model change on average; and by more than six weeks in Figure 2 for Model changes with a high share of new parts.

4.1.3 Knowledge Responses to Product Cycles

We first investigate how the average stock of knowledge, as measured by the cumulative number of completed courses and hours of training for the average worker in the working group, changes after a Model change. Figure 3 reports the results of an event study specification following equation 3 with the average number of cumulative completed courses per worker as dependent variable. The figure shows that in response to the Model changes, the firm

increases the stock of knowledge in each working group to deal with the more frequent arrival of more complex problems. Note that in Figure 3, the impacts persist over time — more than seven weeks following a Model change. The point estimates in Appendix Table A6 show that the stock of knowledge in terms of cumulative completed courses goes up by 3% and in total hours of training by 6%, such that a working group will accrue roughly seven more completed courses or 60 more hours of training across its 20 workers.

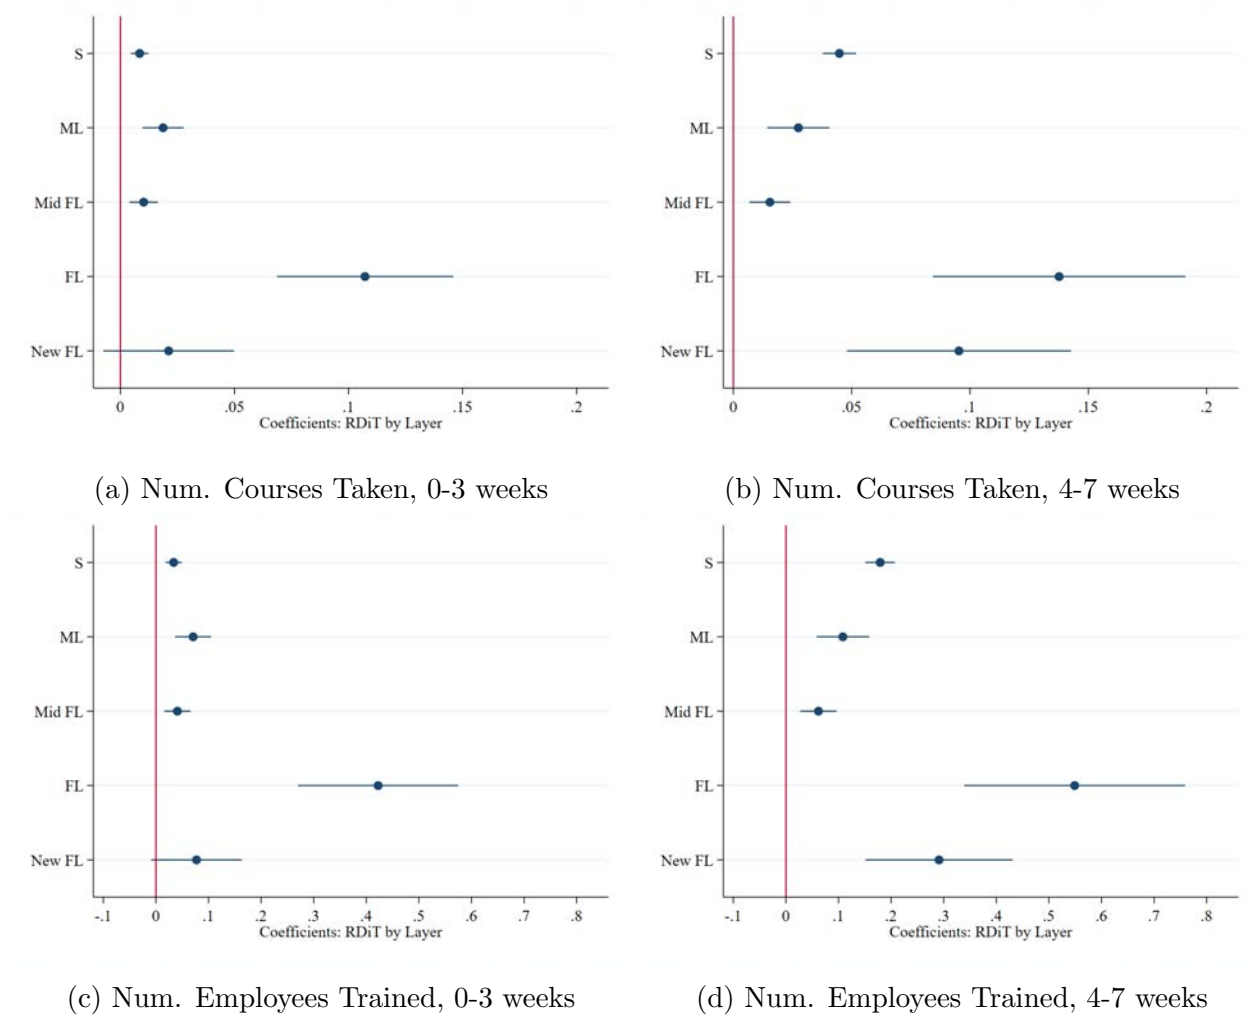
Figure 3: Event Study of Model Changes on Stock of Knowledge in Working Groups



Note: Figure 3 shows the effect of Model changes on the average completed training programs within working groups in a time window running from 8 weeks before the event to 8 weeks after the event (where the week of the Model change is labelled as week 0 on the x-axis). We control for month, year, and group fixed effects. We also control for a linear function of distance to the Model change and to all other Model and Volume changes and a linear time trend. Standard errors are clustered at week-working group level. 95% confidence intervals are reported. Number of observations: 220 working groups x 16 weeks x 2 events.

Figure 4 shows that the incremental training investments – in terms of both number of courses taken by workers and number of employees trained – are concentrated primarily among FL workers (in fact upper FL on the margin of promotion to ML as seen in the more disaggregated results in Appendix Figure A6, panels (a) and (c)), with magnitudes consistent with roughly 4-6 FL workers each completing the additional course they need to reach ML level skill and the remaining 1-3 courses spread across the other workers in the group. Figure A6 in the Appendix also presents this same result, but restricting attention to just courses covering problem-solving and communication content (panels (b) and (d)). The near identical pattern and similar order of magnitude in these results confirm that training investments in these communication and problem-solving skills are primarily driving the knowledge results.

Figure 4: Impact of Model Changes on Incremental Knowledge Provision



Figures 4a and 4b show the effect of Model changes on the average number of courses taken by workers by layer at 0-3 weeks and 4-7 weeks post Model change, respectively. Figures 4c and 4d show the effect of Model changes on the number of employees trained by layer at 0-3 weeks and 4-7 weeks post Model change, respectively. For more details on the definition of the layers see Table 1. Each coefficient is estimated from a separate regression. We control for month, year, and group fixed effects. We also control for a linear function of distance to the Model change and to all other Model and Volume changes. Standard errors are clustered by distance to Model change and working group. 95% confidence intervals are presented in the figure. Number of observations: 220 working groups x 16 weeks x 2 events.

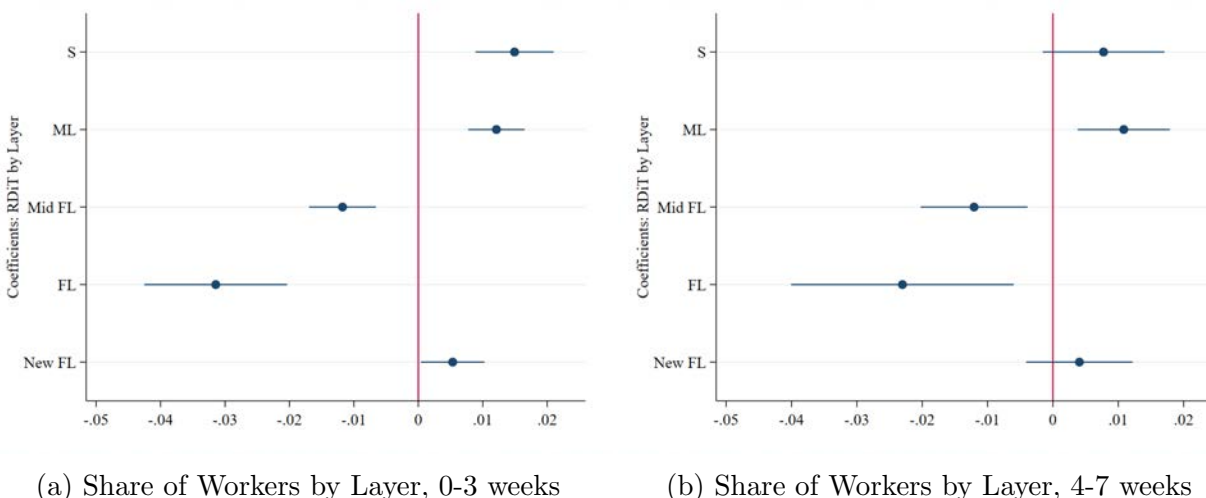
4.1.4 Organizational Responses to Product Cycles

Next, we study which organizational responses to the Model changes the plant puts in place and how these enable the firm to bring back down the level of defects and increase productivity. Appendix Table A7 reports results of regressions like equation 4 at the shift-week level, with various outcome variables related to employment. We find no evidence that Model changes lead to a change in the number of employees in the Assembly sector, nor in the number of hires or separations, in the eight weeks after the event: all coefficients are small in magnitude

relative to the mean and far from being significant. Appendix Table A8 confirms also that the number of working groups and the size of the average group are unaffected by Model changes. This indicates any organizational adjustments we find next in response to Model changes must be taking place as a result of the *reallocation* of existing workers across layers, rather than by increasing or downsizing the workforce. This result is sensible in that the total number of cars that need to be produced has not changed.

Now, to check if indeed the firm is reallocating existing workers, in Figure 5 we study the impact of the Model changes on the shape of the working groups. We do so by running a series of regressions like equation 4 with as dependent variables shares of workers in the working group by layer, in the four weeks (Panel (a)) and eight weeks (Panel (b)) after the Model change. This figure shows a recomposition of working groups with a reduction in upper-FL layers in exchange for an increase in the share of ML and S layer workers. Appendix Figure A8 shows a more disaggregated version of Figure 5, confirming that the recomposition is mainly driven by upper FL workers being upskilled and promoted to higher level positions. These results are consistent with Figures 4 and A6, as promotions in the plant are associated with receiving additional training.

Figure 5: Impact of Model Changes on Working Group Structure and Knowledge Hierarchies



Note: Figures 5a and 5b show the effect of Model changes on the share of workers in the working group by layer at 0-3 weeks and 4-7 weeks post Model change, respectively. For more details on the definition of the layers see Table 1. Each coefficient is estimated from a separate regression. We control for month, year, and group fixed effects. We also control for a linear function of distance to the Model change and to all other Model and Volume changes. Standard errors are clustered by distance to Model change and working group. 95% confidence intervals are presented in the figure. Number of observations: 220 working groups x 16 weeks x 2 events.

Figures 5 and A8 show that working groups are becoming flatter with less distinct knowledge layers, and more rectangular in shape. These results are confirmed in the corresponding regression discontinuity in time estimates reported in Table 4: column 1 shows a negative

and significant effect of about 0.7 layers in the first four weeks after the shock, corresponding to a 10% decrease from the mean number of layers across working groups. The negative effect on layers means that the plant is reducing the number of separate skill levels within working groups after a Model change. Column 2 shows that the ratio of FL to ML workers also decreases by 10%, while Column 3 shows that the ratio of FL to ML and S decreases by 12%. That is, the span of control of ML and S workers (or number of FL workers from which each fields problems) is going down.¹⁷

Table 4: Impact of Model Changes on Working Group Structure

	(1) Num. Layers	(2) FL/ML	(3) FL/(ML+S)
0 to 3 weeks	-0.666*** (0.124)	-0.552*** (0.117)	-0.448*** (0.085)
4 to 7 weeks	-0.285 (0.190)	-0.346* (0.183)	-0.319** (0.130)
Observations	7,040	7,040	7,040
Obs. Level	Group-Week	Group-Week	Group-Week
Mean	6.373	5.026	3.753

Note: Standard errors clustered by distance to event and working group. Number of observations: 220 working groups x 16 weeks x 2 events. Number of Layers is defined as the number of separate positions present in a working group. FL/ML is the ratio between Num. of Front-Line workers and Mid-Line workers. FL/(ML+S) is the ratio between Num. of Front-Line workers and Mid-Line workers plus Superiors. We use as controls month, year and group fixed effects. We also control for a linear function of distance to the Model change and to all other Model and Volume changes. * p<0.1, ** p<0.05, *** p<0.01

Finally, given which workers are being trained and the content covered in those trainings, Table 5 explicitly looks for effects on problem-solving activity. Indeed all members, at all layers, of the working group are solving more problems, consistent with both the removal of a layer and the upskilling of low-level FLs to upper level FLs and MLs with more problem-solving skills. We also see that per employee, the effects are mainly driven by FLs and MLs; though consistent with more complex problems arising after a Model change, we see that each

¹⁷Figure A7 reports the estimation results of equation 2 with the number of layers per working group (averaged at the weekly level) as the dependent variable. We see that Model changes result in a sudden *decrease* in the number of layers in the weeks immediately following the event. More precisely, the number of layers decreases on average by about one layer, and this effect is stable and significant in the first three weeks after the shock, after which the number of layers reverts back to pre-shock levels. The dynamics of the impacts show a remarkable similarity to Figure 1 in that the effect on the number of layers dissipates after about four weeks since the Model change. Note that though the working groups start to back-fill to some degree the layer that was removed, the rectangular shape persists even 2 months after the model change (as shown in Panel (b) of Figure 5 and in columns 2 and 3 of Table 4), consistent with working groups having a permanently higher stock of knowledge.

S worker is indeed fielding more problems than before. Note that both Table 5 and Figure 5 show that though the effect on number of layers is short-lived (as seen in Appendix Figure A7), the effects on working group structure and span of control of complex problem-solvers (i.e., the more rectangular shape of groups) is persistent.

Table 5: Impact of Model Changes on Problem-Solving Activity by Layer

	(1) All	(2) FL	(3) ML	(4) S
Panel A: Number of reports				
0 to 3 weeks	0.735*** (0.148)	0.607*** (0.125)	0.120*** (0.0301)	0.00777 (0.00524)
4 to 7 weeks	2.326*** (0.279)	1.932*** (0.238)	0.373*** (0.0544)	0.0213** (0.00846)
Mean	0.826	0.661	0.159	0.006
Panel B: Share of reports per employee				
0 to 3 weeks	0.0935*** (0.0152)	0.0945*** (0.0158)	0.119*** (0.0242)	0.0241* (0.0141)
4 to 7 weeks	0.231*** (0.0250)	0.239*** (0.0264)	0.278*** (0.0376)	0.0537*** (0.0205)
Mean	0.072	0.072	0.096	0.014
Observations	3,520	3,520	3,520	3,520
Obs. Level	Group-Week	Group-Week	Group-Week	Group-Week

Note: Standard errors clustered by distance to event and shift. Number of observations: 2 shifts x 16 weeks x 2 events. We control for month, year, and shift fixed effects. We also control for a linear function of distance to the Model change and to all other Model and Volume changes. * p<0.1, ** p<0.05, *** p<0.01

In sum, the plant is increasing the overall stock of human capital of working groups, by upskilling FL workers in problem solving. Roughly 4-6 FL workers are completing the additional course they need to reach ML level skill (Figure 4). The plant shifts working group composition away from upper-FL layers and towards ML and S level positions (Figure 5). In doing so, the plant is compressing the hierarchy and reducing the distance, in terms of knowledge layers, between workers farther up the hierarchy, who have been trained to possess the necessary knowledge to solve complex problems, and front-line workers working on physical production, who are coming against new problems as a result of the introduction

of new parts. As described in Section 2, training focuses primarily on communication and problem solving, especially for workers farther up the hierarchy. These results then indicate that through such training provision, the firm increases the ability of upper level workers to help front-line workers solve the new complex problems they come up against, but the span of control of each of these complex problem-solvers must be reduced to compensate such that the working groups become more rectangular in shape.

These results are notable in that they show how the firm combines both in-house training and changes in organizational structure to adapt to product quality upgrading. Our granular data allows us to document the surprising flexibility of the organization and to uncover a substantial and immediate response to the introduction of new models, which would largely be missed with less granular data aggregated at the yearly level, for example.

These results highlight how an increase in the complexity of the problems that need to be solved can lead to a compression of the knowledge hierarchy through a reduction in organizational layers. This is a novel result, which stands in contrast to the literature studying how the hierarchical structure of firms changes with an expansion of production. As we will discuss in more detail in the next section, this literature tends to find that as firms grow larger and there is a larger number of problems to be solved, the number of knowledge layers *increases*, as this allows groups to focus the scarce time of skilled managers only on the most complex problems (Caliendo et al., 2020, 2015; Friedrich, 2022). Our novel data allows us to highlight instead organizational responses to an increase in the complexity of what needs to be produced, while keeping total quantity produced constant. We show how the organizational response is very different when the complexity of production increases, but quantity produced does not change.

In the next subsection, we exploit discrete increases in the volume of production to show that when production expands but the complexity of problems does not change, this leads to an increase in overall employment and the number of layers in each working group as well as an intensifying of the pyramidal shape. Thus, we can replicate the finding in the literature regarding organizational responses to the expansion of production, but also add novel and complementary empirical findings for stock of knowledge and span of control, both core to the intuition of models of knowledge hierarchies (Caliendo and Rossi-Hansberg, 2012; Garicano, 2000). In the final section, we then develop a model which extends these classic models of hierarchies to help reconcile why the number of layer decreases with Model changes but increases with Volume changes.

4.2 Contrasting Product Cycles with Volume Changes

We now contrast the organizational response to product quality upgrading documented in the previous subsection with the response of the firm to sudden and sharp increases in the volume of cars that need to be produced. In short, we find that while Volume changes lead to a quantitatively similar spike in the incidence of defects per vehicle, the organizational response is very different: the plant hires more entry-level workers and adds more layers to working groups, so that the distance – in terms of knowledge layers – between front-line workers and supervisors farther up the hierarchy increases, the span of control of complex problem-solvers increases, and as a result the average skill level in the firm decreases. These results for volume changes are in line with the literature on knowledge hierarchies (Caliendo et al., 2020, 2015; Friedrich, 2022), which tends to find that as firms expand production, new layers are added to the hierarchy, because the higher volumes allow the firm to better focus the talent of highly skilled managers on solving only the more complex problems. However, the analysis of worker-specific stocks of knowledge and problem-solving activities are, to our knowledge, novel to the literature.

We first confirm in Table A9 that Volume changes result in a sudden and sizeable increase in the planned number of cars to be produced and, consequently, in the number of parts that have to be assembled. The total number of cars planned per day jumps up by about 9% and the number of parts increases by 7% after the Volume change; and these effects on production are long-lasting. Since by definition the model produced on either side of Volume changes does not change, the types of parts that have to be assembled are the same as before. That is, the share of new parts is consistently 0 through this window so we do not analyze this outcome for Volume changes.


The change in volume leads to a sudden and large increase in defects per vehicle, as shown in Figure 6. DPV shoots up by about 3 SD on average in the immediate aftermath of the change, and then comes down to pre-shock levels in about three weeks on average.¹⁸ The similarity with Figure 1 is remarkable: both Volume and Model changes lead to a similar short-term reduction in productivity.¹⁹ Appendix Figure A9 shows that the initial spike in DPV is fairly similar, but the length of time over which DPV remains elevated is longer for Volume changes for which the factory hires more new workers. This pattern of heterogeneity supports the interpretation that the rise in DPV following Volume changes is due to new workers and new groups making mistakes. That is, the spike in DPV after Volume changes still derives from novelty, but the parts and processes are new to the new workers and groups,

¹⁸This pattern holds, although less precise, when we restrict attention to the period 2017 to 2019 coinciding with the working group composition database. Appendix Table A11 confirm this.

¹⁹Appendix Table A10 shows the corresponding discontinuity in time estimates.

rather than being new to the model.

Figure 6: Event Study of Volume Changes on Productivity



Note: Figure 6 shows the effect of Volume changes on DPV in a time window running from 40 days before the event to 40 days after the event. DPV is computed as number of defects per 100 vehicles, and is standardized using the mean and standard deviation of the full sample. We control for month, year, and shift fixed effects. We also control for a linear function of distance to the Volume change and to all other Volume and Model changes and a time linear-trend. Standard errors are clustered by distance to the event-shift level. 95% confidence intervals are reported. Number of observations: 2 shifts x 81 days x 5 events.


However, the organizational response of the firm to the increase in the volume of production is very different. Figure 7 reports the results of an event study specification following equation 3 with the stock of knowledge in the working group as measured by the average number of cumulative completed courses per worker as the outcome. The figure shows that in response to the Volume changes, the average stock of knowledge in each working group *decreases* immediately and persistently.²⁰ We next ask how the decrease in the average knowledge is generated. Table 6 shows that as a response to Volume changes the firm hires about 24 additional workers per shift-week in the first four weeks after the Volume change (column 2) while separations do not increase significantly (column 3). The result is that overall shift-level employment increases by about 80 workers over the two months following the Volume change, or a 7.5% increase over a mean of about 1,100 workers per shift (column 1).²¹ This hiring is

²⁰Appendix Table A12 reports the corresponding regression discontinuity in time results and shows that Volume changes lead to a persistent reduction in average completed training programs of about 7% and a reduction in average cumulative training hours of about 11% after eight weeks.

²¹Appendix Figure A12 reports the corresponding event studies on employment, showing that Volume changes lead to a persistent increase in employment over at least an eight-week period. Appendix Table A13 shows that the increase in employment is associated with an increase in the number of working groups in the

concentrated among front-line layers with a low level of skill. Column 1 of Table 7 confirms that the way this stock of knowledge reduction is being achieved is by adding layers and increasing the relative size of the base of the pyramid and consequently the span of control of complex problem-solvers (column 2 and 3 of Table 7). That is, as a result of Volume changes, the firm moves to a more pronounced pyramidal structure of working groups with a thicker base and a larger number of separate knowledge layers, resulting in an overall reduction in the average level of knowledge of working groups.²²

Figure 7: Event Study of Volume Changes on Stock of Knowledge in Working Groups



Note: Figure 7 shows the effect of Volume changes on the average completed training programs within working groups in a time window running from 8 weeks before the event to 8 weeks after the event (where the week of the Model change is labelled as week 0 on the x-axis). We control for month, year, and group fixed effects. We also control for a linear function of distance to the Volume change and to all other Model and Volume changes and a linear time trend. Standard errors are clustered at week-working group level. 95% confidence intervals are reported. Number of observations: 220 working groups x 16 weeks x 2 events.

Assembly sector, but not with a significant increase in the size of pre-existing working groups. In addition, Appendix Table A14 shows that despite the increase in employment, Volume changes still lead to an increase in the number of cars and number of parts per employee as well as per working group. This then justifies why we see an increase in the number of layers within working groups: the firm reorganizes production to deal with the increase in the volume of work per employee, achieving greater economies of scale. In line with this, we still find a positive and persistent effect on the number of layers when restricting the sample to pre-existing working groups only (Appendix Figure A11).

²²Appendix Figure A10 shows that volume changes lead to a sudden *increase* in the number of layers within working groups, which go up by just under 1 layer (from a mean of about 5 layers). This increase is persistent for at least 8 weeks after the volume change, which is consistent with the persistent increase in the volume that needs to be produced, as shown in Table A9. Appendix Figure A13 reports impacts of Volume changes on the share of workers by layer, confirming that the share of FL workers in working groups increases, so that groups become heavier in the lower part of the pyramid.

Table 6: Impact of Volume Changes on Employment

	(1) Employment	(2) Hires	(3) Separations
0 to 3 weeks	29.46 (34.32)	24.08** (9.63)	1.67 (1.50)
4 to 7 weeks	82.08*** (28.60)	-0.58 (2.48)	0.25 (0.36)
Observations	64	64	64
Obs. Level	Shift-Week	Shift-Week	Shift-Week
Mean	1094	10	1

Note: Standard errors clustered by distance to Volume change and shift. Number of observations: 2 shifts x 16 weeks x 2 events. We use as controls month, year, and shift fixed effects. We control for a linear function of distance to Volume changes and distances to the other Volume and Model changes. * p<0.1, ** p<0.05, *** p<0.01

Table 7: Impact of Volume Changes on Working Group Structure

	(1) Num. Layers	(2) FL/ML	(3) FL/(ML+S)
0 to 3 weeks	0.850*** (0.098)	0.610*** (0.094)	0.453*** (0.067)
4 to 7 weeks	0.706*** (0.117)	0.498*** (0.116)	0.398*** (0.082)
Observations	7,040	7,040	7,040
Obs. Level	Group-Week	Group-Week	Group-Week
Mean	5.284	4.158	3.094

Note: Standard errors clustered by distance to event and working group. Number of observations: 220 working groups x 16 weeks x 2 events. Number of Layers is defined as the number of separate positions present in a working group. FL/ML is the ratio between Num. of Front-Line workers and Mid-Line workers. FL/(ML+S) is the ratio between Num. of Front-Line workers and Mid-Line workers plus Superiors. We use as controls month, year and group fixed effects. We also control for a linear function of distance to the Volume change and to all other Model and Volume changes. * p<0.1, ** p<0.05, *** p<0.01

Finally, Table 8 shows that given that there are now more FLs to solve the higher number of simple problems corresponding to an increase in volume, each FL and each ML is solving fewer problems. This result reflects both a gain in economies of scale by way of an increase in the span of control of MLs and Ss and perhaps some learning in the cumulative quantity produced as well (Levitt et al., 2013). This comparison again highlights how different and unique are organizational responses to product quality upgrading vis a vis responses to

Volume changes, which have been extensively studied in the literature (Caliendo et al., 2020, 2015; Friedrich, 2022).

Table 8: Impact of Volume Changes on Problem-Solving Activity by Layer

	(1) All	(2) FL	(3) ML	(4) S
Panel A: Number of reports				
0 to 3 weeks	-1.964*** (0.220)	-1.630*** (0.191)	-0.321*** (0.0400)	-0.0123* (0.00713)
4 to 7 weeks	-2.064*** (0.378)	-1.677*** (0.326)	-0.374*** (0.0740)	-0.0125 (0.0116)
Mean	0.705	0.588	0.111	0.006
Panel B: Share of reports per employee				
0 to 3 weeks	-0.174*** (0.0172)	-0.181*** (0.0188)	-0.210*** (0.0247)	-0.0286* (0.0158)
4 to 7 weeks	-0.195*** (0.0300)	-0.202*** (0.0319)	-0.252*** (0.0477)	-0.0319 (0.0259)
Mean	0.070	0.073	0.082	0.016
Observations	3,520	3,520	3,520	3,520
Obs. Level	Group-Week	Group-Week	Group-Week	Group-Week

Note: Standard errors clustered by distance to event and shift. Number of observations: 2 shifts x 16 weeks x 2 events. We control for month, year, and shift fixed effects. We also control for a linear function of distance to the Volume change and to all other Model and Volume changes. * p<0.1, ** p<0.05, *** p<0.01

5 Interpretation

The findings in Section 4 show that there is a discrete jump in defects per vehicle right after the Model or Volume change. However, we find that a) when workers have to deal with new tasks, the firm trains existing workers and compresses the hierarchy, and b) Volume changes lead to a different organizational response: the plant hires more entry-level workers and adds more layers to working groups. In this section, we discuss how our results can be reconciled with the literature on hierarchies in organizations (Caliendo et al., 2020, 2015; Caliendo and Rossi-Hansberg, 2012).

5.1 Interpreting Our Results Through Existing Models

We consider as benchmark a simplified version of the model in [Caliendo and Rossi-Hansberg \(2012\)](#), which draws in turn from [Garicano \(2000\)](#).²³ We briefly summarize the key elements of the model, and then discuss whether this model can explain our pattern of empirical results.

In this model, organizations are composed of two types of agents: workers and managers. To generate output, workers in the organization need to solve a problem drawn from a cumulative distribution F , with decreasing density ($f' < 0$). Solving problems requires knowledge. A realization z of a problem implies that in order to solve that problem, the worker needs to have acquired a set of knowledge that includes z as an element. If the worker solves the problem, the production possibility becomes A units of output.

If the worker cannot solve the problem, she asks a manager one layer above for a solution. Then, the manager spends h units of her time listening to the worker's problem (i.e., h is the cost of communication) and solves the problem if her set of knowledge includes z . If the manager cannot solve the problem, she can communicate it to another manager one layer above her. This process continues until the problem is solved or the problem reaches the only agent in the highest layer of the organization, namely the entrepreneur. To achieve a set of knowledge $[0, z]$ for a given agent (which means that the agent can solve any problem within the interval $[0, z]$), the firm must pay wcz , that is, the cost of knowledge is wc per unit of knowledge, where w is the wage and c the training cost.

For simplicity we assume that problems are drawn from the exponential distribution; that is, $F(z) = 1 - e^{-\lambda z}$ for a given $\lambda > 0$. Note that as $f(z) = \lambda e^{-\lambda z}$ is strictly decreasing in z , agents at the bottom of the organization learn the most common problems, while agents in higher layers learn rarer problems. Also note that as λ decreases, the frequency of complex problems increases.

Given a production level q , a wage w and a training cost c , each firm solves an organizational problem that involves choosing the optimal number of layers L , the optimal amount of workers/managers at each layer, and the amount of knowledge these workers and managers acquire. The firm does so by solving two minimization problems, given q , w and c :

1. (Layers problem) It chooses the optimal number of layers L in order to minimize the cost of producing q units of output while paying a wage w to each worker and manager.

²³We note that there is another class of models of the organization of production based on the intuition of delegating responsibilities from the top down rather than passing problems from the bottom up (see, e.g., [Alonso et al. \(2008\)](#)). Given the focus on problem-solving in the training content in this setting and the emphasis on recording and encouraging problem-solving effort on the part of workers in this firm and industry context, we believe these models of knowledge hierarchies describing the distribution of problem-solving skills across the hierarchy are particularly appropriate for analyzing the changes we study in this paper.

2. (Workers and Knowledge problem) Given the number of layers L , it chooses the optimal amount of workers/managers at each layer and the optimal amount of knowledge they acquire to produce q units of output, given a training cost, c .²⁴

Using this general framework, we study how the organization of the firm changes endogenously in response to Model and Volume changes. We begin by considering Volume changes, as this maps closely to related studies that consider how the organization of the firm changes in response to an increase in the quantity produced. The model predicts that an increase in the volume produced should result in an increase in the number of layers and in the number of employees per layer, which is exactly what we see in the data.

The intuition is that when production increases, the number of problems increases too, and so each worker solves a smaller fraction of the problems. As a response, the firm can either: (i) hire new workers and increase the number of layers, so that new managers deal with rarer problems, or (ii) increase the number of employees and knowledge in all layers without an increase in the number of layers. Note that since the firm needs to produce more, the model predicts that new workers will be hired regardless of the impact on the hierarchical structure. Our empirical evidence is consistent with (i) in that we document a sharp and sustained increase in the number of layers and employees as a response to volume changes. Therefore, the model in [Caliendo and Rossi-Hansberg \(2012\)](#) can reconcile the impact of Volume changes that we document empirically.

Turning to Model changes, we note that this is not something that the literature on hierarchies in organizations has explored before. Nevertheless, we start by discussing whether the model in [Caliendo and Rossi-Hansberg \(2012\)](#) can reconcile the results on Model changes that we document. As shown before, every time a new model is introduced, the complexity of the problems increases as the share of new parts in the car increases. We follow [Garicano \(2000\)](#) in modeling this as a reduction in λ . A lower λ decreases the production level q_L at which the firm would find it optimal to move from L to $L + 1$ layers, thus *increasing* the number of layers and workers per layer for a fixed level of production.

Intuitively, if problems become more complex, given a stock of knowledge, each worker solves a smaller fraction of the problems (see Proposition 1 in Appendix B)). Thus, to keep production constant, there are two options: (i) increase the number of layers and workers to

²⁴We note that in deciding the number of layers, how many workers should be in each layer, and how much knowledge they should acquire the firm effectively solves a top down delegation problem in this model as well. We therefore do not see this knowledge hierarchy class of models as necessarily distinct in intuition from the delegation class of models like that in [Alonso et al. \(2008\)](#). Rather, we note practically that both assignment of responsibility and acquisition of skill must occur for gains from specialization to be realized. Indeed, if a manager is given the responsibility to solve a problem or complete a task without the necessary skill, or conversely acquires the skill but does not have the authority to use the skill to complete a task or solve a problem, then the organization cannot benefit from the presence of this manager in the hierarchy.

deal with the more complex problems, or (ii) increase the number of employees and knowledge in all layers, without an increase in the number of layers. That is, the model predicts that Model changes should result in an increase in the number of workers and either no impact or a *positive* impact on the number of layers (depending on the size of the shock). We instead find no impact on the number of workers and a *negative* impact on the number of layers. The model of [Caliendo and Rossi-Hansberg \(2012\)](#) therefore cannot reconcile the impact of Model changes that we document.²⁵

Next, we discuss a possible model extension that could help reconcile our results on Model changes, and present supporting empirical evidence to motivate the extension.

5.2 Extending Existing Models: The Role of Training

Training plays a key role in determining the equilibrium levels of knowledge at different levels of the hierarchy in our setting and in many similar manufacturing settings. We know from anecdotal evidence that every time the complexity of the routine tasks increases due to a Model change, the partner firm *improves* the training programs (in terms of both content and trainers) to train workers more efficiently in problem solving and managerial skills. We interpret this as an *increase in the productivity of training*, which in the model would correspond to a reduction in the training cost c . One advance we make to the current state of the theory that is relevant for our setting is to incorporate the role of training and endogenizing the equilibrium level of training the firm offers at different levels of the hierarchy.

We first explore empirically how training investments by the firm differ for Volume and Model changes, to substantiate the claim that the firm increases the productivity of training in response to Model changes. Table 9 shows that the firm increases substantially the number of *different* courses provided *in house* and the number of trainers after a new model is introduced (Panel A). These results are consistent with the anecdotal evidence provided by the partner company that the firm makes investments to improve the *in house* training academy. It also again highlights training as an important margin that enables the firm to adapt to Model changes and resolve defects quickly, consistent with the empirical evidence documented in Section 4. In contrast, we do not see any impact when the firm changes the scale of production (Panel B). This is expected as the complexity of the problems does not

²⁵For simplicity, in this section, we consider a static model. However, a dynamic version of the model of [Caliendo and Rossi-Hansberg \(2012\)](#) where workers learn by working on new tasks would still not reconcile the empirical results presented in Section 4: a Model change would still lead to a (temporary) increase in either the number of workers or the number of layers. Then, as workers learn the new tasks, the stock of knowledge would increase in each layer, leading to a subsequent reduction in the number of layers or the number of workers. Instead, we document an immediate *decrease* in the number of layers after the Model change, with no impact on the number of workers.

change with Volume changes.

Table 9: Impact of Model and Volume Changes on Training Investments

	(1)	(2)
	Number of Courses Provided	Number of Trainers
Panel A: Model Changes		
0 to 3 weeks	2.743*** (0.443)	3.173*** (0.668)
4 to 7 weeks	4.311*** (1.123)	5.345*** (1.469)
Observations	112	112
Obs. Level	Weekly	Weekly
Mean	6.255	8.055
Panel B: Volume Changes		
0 to 3 weeks	-1.588 (0.943)	-2.213 (1.668)
4 to 7 weeks	-0.466 (1.535)	-0.607 (2.169)
Observations	80	80
Obs. Level	Weekly	Weekly
Mean	6.100	8.175

Note: Standard errors clustered by distance to event. Number of Courses Provided is the number of courses provided each week and Number of Trainers is the number of trainers teaching each week. We use as controls month and year fixed effects. We also control for a linear function of distance to the event (i.e., Model or Volume change) and to all other events in the data. Panel A shows the effect for Model changes and Panel B shows the effect for Volume changes. Number of observations in Panel A: 7 events x 16 weeks. Number of observations in Panel B: 5 events x 16 weeks. * p<0.1, ** p<0.05, *** p<0.01

To reconcile the above results with the previous theoretical literature we start by asking: *What are the admissible ratios $\Delta c : \Delta \lambda$ that imply that the firm does not optimally increase the number of layers when the complexity of tasks increases?* In Appendix B, we show that the partial derivatives $(\partial q_L / \partial \lambda)|_{\bar{p}} > 0$ and $(\partial q_L / \partial c)|_{\bar{p}} < 0$ quantify the opposing effects on the intersecting point q_L (i.e., the quantity q at which the firm finds it optimal to switch from L to $L + 1$ layers) of reducing λ or c . The former constitutes the negative impact of the firm having to face more complex problems – which pushes towards an increase in number of layers for a given quantity of production – while the latter is the favorable scenario where the firm can train its workers more efficiently – which pushes towards a *reduction* in the number

of layers. In that sense, Proposition 2 in Appendix B provides the minimum investment in reducing training costs c required to balance the effect of problem complexity.

Motivated by the observations discussed above, but being mindful that firms usually can not freely change the cost of training c , in Appendix B we develop a variation of the model in Caliendo and Rossi-Hansberg (2012) where the firm can endogenously invest in reducing the training cost c but this comes at a fee, or penalty \mathcal{P} .²⁶ Intuitively, this extension captures the idea that when the complexity of production increases, the firm may find it optimal to invest in its training programs to improve the productivity of training, consistent with the evidence in Table 9.

In this version of the model with endogenous training cost, for each c the firm has to establish the optimal distribution of knowledge $\{z_L^l\}_{l=0}^L$, and then it selects the optimal training cost c , given the fee \mathcal{P} . Proposition 3 in the Appendix shows that in this version of the model there is a neighborhood of λ where an increase in complexity (i.e., a marginal reduction in λ) compresses knowledge layers in the organization.

In Appendix B3, we provide an example to illustrate the results presented in Proposition 3. Our simulations show that there is a set of production levels where the firm reduces the amount of layers every time the complexity of the tasks increases (e.g., a Model change), increasing the amount of knowledge within each layer.²⁷ Note that this is not the case for Volume changes as the number of layers generally increases for large volume increases, even when we allow the option to invest in training by reducing the training cost c (see Appendix Figure B2).

Intuitively, there are dynamic elements and adjustment costs at play here. In the short run, firms are confronted with these changes in complexity arising from Model changes that require new knowledge. However, it is inefficient for the firm to hire new workers who would require even more training to deal with this rise in complexity, especially given that quantity is fixed in the short run. Instead, the firm can first quickly retrain existing workers and reorganize teams to deal with the higher complexity at the current quantity produced.

To make this intuition clear, rather than propose an explicitly dynamic framework, we differentiate fixed factors in the short run and long run adjustment. With some abuse of notation (since L is not continuous), we see:

$$\frac{dq}{d\lambda}|_{q=\bar{q}} = 0 = \frac{\partial q}{\partial z} \frac{\partial z}{\partial \lambda} + \frac{\partial q}{\partial L} \frac{\partial L}{\partial \lambda} \quad (5)$$

and hence

²⁶We assume that \mathcal{P} is decreasing with respect to c and independent of the production levels q .

²⁷In other cases, the firm operates with the same amount of layers when there is a Model change. Note that this result depends on the magnitude of the increase in complexity.

$$\frac{\partial q}{\partial z} \frac{\partial z}{\partial \lambda} = - \frac{\partial q}{\partial L} \frac{\partial L}{\partial \lambda}. \quad (6)$$

We know q increases in z and L ; so equation 6 shows that an increase in training (knowledge z) to deal with the new problem distribution will go together with a reduction in the number of layers. After this phase, the firm can choose to expand production in the longer term if it wants. That is, when there is this adjustment cost to changing quantity, knowledge (training) and layers are substitutes in short run, but complements in the long run.

This suggests a cycle of growth for the firm in which the firm adds knowledge with a relatively flatter and more rectangular hierarchy to allow for product quality upgrading, and then shifts to a relatively taller and more pyramidal hierarchy with a larger and less skilled on average workforce when expanding production. If the firm continues to grow, the cycle starts back at the beginning but with ever an increasing training productivity and size of the workforce. Indeed in Appendix Figures A14, A15, A3a and A3c, we see that the number of courses and trainers continues to rise over time, as does the stock of knowledge in and number and size of working groups, in stages.

This discussion shows that taking into account the “technology of training” and how firms might endogenously decide to invest in improving it can be important for understanding organizational responses to the introduction of more complex problems such as those resulting from product quality upgrading. In considering the role of the training technology, we advance the recent theoretical literature on knowledge hierarchies, which has mostly considered organizational responses to changes in quantity produced rather than in the complexity or quality of what needs to be produced (Caliendo et al., 2020, 2015).

6 Conclusion

Focusing on the automotive sector as a prototypical example of an industry that experiences frequent product quality upgrading due to product cycles, we study how the stock of knowledge and organization of the firm change in response. To do so, we combine granular administrative data on production, employee hierarchies and training provision from an Argentinian subsidiary plant of a leading global auto manufacturer with event study and discontinuity-based methods.

We find that demand (planned number of vehicles) and number of total parts do not change in the short term after a new model is introduced. The main change is a large, discontinuous increase in new parts. Accordingly, the production of new models necessitates

dealing with new complex problems. Indeed, we show that defects per vehicle increase substantially after the production change, most dramatically when new models involve a higher share of new parts. Defects decrease to their prior level over a period of about 3 weeks on average, though they remain elevated for more than 8 weeks following more complete redesigns of models.

We then show that combatting this rise in defects requires an increase in the stock of knowledge and a change in the shape of working groups: the knowledge hierarchy becomes flatter and the groups become more rectangular. The firm accomplishes this by training and promoting lower-level employees to mid and top level positions that specialize in solving the complex problems arising from the use of new parts and processes in assembly. In doing so, the firm reduces the distance – in terms of knowledge layers – between front-line workers, who are dealing with these new tasks, and managers farther up the hierarchy, who have the necessary knowledge to solve the new complex problems that arise as a result. The firm also reduces the span of control of complex problem-solvers by decreasing the ratio of frontline operators to mid and high level managers. Though the hierarchy starts to heighten relatively soon after defects return to their pre-change level, the more rectangular shape and higher stock of knowledge persist.

These results for model changes contrast interestingly with the impacts of an increase in quantity produced (i.e., volume changes). In this case we show the organizational response is a monotonic, permanent increase in both employment and knowledge layers, consistent with prior evidence from manufacturing in high-income countries ([Caliendo et al., 2020, 2015](#); [Friedrich, 2022](#)). The pyramidal shape of working groups becomes more pronounced after volume changes and the stock of knowledge goes down as more lower-skilled workers are hired at the bottom of the pyramid.

In addition to these novel empirical results on the stock of knowledge and the shape of working groups, we leverage unique data from the factory’s continuous improvement system to show that problem-solving activity by all workers at all levels goes down following volume changes, but up after model changes. The increases in problem-solving activities we see following model changes at both the lower and mid-levels of the working group reflect a combination of the results from comparative statics explored by both [Garicano \(2000\)](#) and [Caliendo and Rossi-Hansberg \(2012\)](#). That is, they show that an increase in problem complexity and a decrease in the cost of acquiring knowledge move the choice of hierarchical knowledge layers and span of control of complex problem-solvers in opposite directions.

Motivated by these insights, as well as additional empirical results showing the firm invests in improving its in-house training programs by adding courses and bringing in new trainers, we modify the canonical theory to allow the firm to make a costly investment in

increasing the efficacy of training. In this sense, training becomes another lever the firm adjusts alongside the distribution of knowledge across layers. We show that when quantity is fixed in the short-run, as both our empirical results and anecdotes from the firm’s upper management confirm, the firm treats investments in training efficacy and the choice of layers as substitutes, but when quantity can adjust in the long run these levers are complementary.

We believe these insights, however, are not unique to this firm or industry context. The emphasis on problem-solving skills, processes and systems is ubiquitous across the automobile manufacturing industry, and indeed makes up the core of the mechanism for organizational learning presented in [Levitt et al. \(2013\)](#). Furthermore, lean manufacturing and continuous improvement are regarded as frontier production practices across most manufacturing contexts. For example, [Adhvaryu et al. \(2023\)](#) study this same system in electronics manufacturing factories in Thailand; and communication and problem-solving skills of both frontline operators and production team supervisors have been shown to contribute greatly to productivity in Indian garment factories ([Adhvaryu et al., 2018, 2022, 2019](#)).

We provide what is to our knowledge the first study of the way in which organizational structure responds to product quality upgrading, as occurs regularly in product cycles. Our findings show how large suppliers in the “global south” are highly flexible in their internal organization of labor, and how this allows them to adapt and respond to the ever increasing complexity of production arising from the frequent and relentless product quality upgrading necessary to remain competitive in global product markets.

References

- Acemoglu, D. and Pischke, J.-S. (1998). Why do firms train? theory and evidence. *The Quarterly journal of economics*, 113(1):79–119.
- Acemoglu, D. and Pischke, J.-S. (1999). Beyond becker: Training in imperfect labour markets. *The economic journal*, 109(453):112–142.
- Adhvaryu, A., Bassi, V., Nyshadham, A., and Tamayo, J. A. (2020). No line left behind: Assortative matching inside the firm. Technical report, National Bureau of Economic Research.
- Adhvaryu, A., Kala, N., and Nyshadham, A. (2018). The skills to pay the bills: Returns to on-the-job soft skills training. Technical report, National Bureau of Economic Research.
- Adhvaryu, A., Molina, T., and Nyshadham, A. (2023). Problem-solving skills training and continuous improvement in electronics manufacturing.
- Adhvaryu, A., Murathanoglu, E., and Nyshadham, A. (2022). On the allocation and impacts of managerial training.
- Adhvaryu, A., Nyshadham, A., and Tamayo, J. A. (2019). Managerial quality and productivity dynamics. Technical report, National Bureau of Economic Research.
- Aghion, P., Bloom, N., and Van Reenen, J. (2014). Incomplete contracts and the internal organization of firms. *The Journal of Law, Economics, & Organization*, 30(suppl_1):i37–i63.
- Alonso, R., Dessein, W., and Matouschek, N. (2008). When does coordination require centralization? *American Economic Review*, 98(1):145–179.
- Amiti, M. and Khandelwal, A. K. (2013). Import competition and quality upgrading. *Review of Economics and Statistics*, 95(2):476–490.
- Amodio, F. and Martinez-Carrasco, M. A. (2018). Input allocation, workforce management and productivity spillovers: Evidence from personnel data. *The Review of Economic Studies*, 85(4):1937–1970.
- Arrow, K. J. (1962). The economic implications of learning by doing. *The Review of Economic Studies*, 29(3):155–173.
- Atkin, D., Chaudhry, A., Chaudry, S., Khandelwal, A. K., and Verhoogen, E. (2017a). Organizational barriers to technology adoption: Evidence from soccer-ball producers in pakistan. *The Quarterly Journal of Economics*, 132(3):1101–1164.
- Atkin, D., Chen, M. K., and Popov, A. (2022). The returns to face-to-face interactions: Knowledge spillovers in silicon valley. Technical report, National Bureau of Economic Research.

- Atkin, D., Khandelwal, A. K., and Osman, A. (2017b). Exporting and firm performance: Evidence from a randomized experiment. *The quarterly journal of economics*, 132(2):551–615.
- Bai, J., Barwick, P., Cao, S., and Li, S. (2021). Quid pro quo, knowledge spillover, and industrial quality upgrades: Evidence from the chinese auto industry. Technical report, Working Paper.
- Bandiera, O., Prat, A., Hansen, S., and Sadun, R. (2020). Ceo behavior and firm performance. *Journal of Political Economy*, 128(4):1325–1369.
- Bayus, B. L. (1994). Are product life cycles really getting shorter? *Journal of product innovation management*, 11(4):300–308.
- Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of political economy*, 70(5, Part 2):9–49.
- Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013). Does management matter? evidence from india. *The Quarterly Journal of Economics*, 1(51):51.
- Bloom, N., Sadun, R., and Van Reenen, J. (2010). Does product market competition lead firms to decentralize? *American Economic Review*, 100(2):434–38.
- Bloom, N., Sadun, R., and Van Reenen, J. (2016). Management as a technology? Technical report, National Bureau of Economic Research.
- Bloom, N. and Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *The Quarterly Journal of Economics*, pages 1351–1408.
- Bloom, N. and Van Reenen, J. (2010). Why do management practices differ across firms and countries? *The Journal of Economic Perspectives*, pages 203–224.
- Caliendo, L., Mion, G., Opromolla, L. D., and Rossi-Hansberg, E. (2020). Productivity and organization in portuguese firms. *Journal of Political Economy*, 128(11):4211–4257.
- Caliendo, L., Monte, F., and Rossi-Hansberg, E. (2015). The anatomy of french production hierarchies. *Journal of Political Economy*, 123(4):809–852.
- Caliendo, L. and Rossi-Hansberg, E. (2012). The impact of trade on organization and productivity. *The quarterly journal of economics*, 127(3):1393–1467.
- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.
- de Rochambeau, G. (2017). Monitoring and intrinsic motivation: Evidence from liberia’s trucking firms. *Unpublished Manuscript*.
- Espinosa, M. and Stanton, C. T. (2022). Training, communications patterns, and spillovers inside organizations. Technical report, National Bureau of Economic Research.

- Fan, H., Li, Y. A., and Yeaple, S. R. (2015). Trade liberalization, quality, and export prices. *Review of Economics and Statistics*, 97(5):1033–1051.
- Fenzia, A. (2022). Managers and productivity in the public sector. *Econometrica*, 90(3):1063–1084.
- Frederiksen, A., Kahn, L. B., and Lange, F. (2020). Supervisors and performance management systems. *Journal of Political Economy*, 128(6):2123–2187.
- Friedrich, B. U. (2022). Trade shocks, firm hierarchies, and wage inequality. *Review of Economics and Statistics*, 104(4):652–667.
- Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of political economy*, 108(5):874–904.
- Garicano, L. and Rossi-Hansberg, E. (2006). Organization and inequality in a knowledge economy. *The Quarterly journal of economics*, 121(4):1383–1435.
- Ghosh, A. (2022). Religious divisions and production technology: Experimental evidence from india. *Available at SSRN 4188354*.
- Grossman, G. M. and Helpman, E. (1991a). Endogenous product cycles. *The Economic Journal*, 101(408):1214–1229.
- Grossman, G. M. and Helpman, E. (1991b). Quality ladders and product cycles. *The Quarterly Journal of Economics*, 106(2):557–586.
- Grossman, G. M. and Helpman, E. (1994). Endogenous innovation in the theory of growth. *Journal of Economic Perspectives*, 8(1):23–44.
- Guillouet, L., Khandelwal, A., Macchiavello, R., and Teachout, M. (2021). Language barriers in multinationals and knowledge transfers. Technical report, National Bureau of Economic Research.
- Hjort, J. (2014). Ethnic divisions and production in firms. *The Quarterly Journal of Economics*, 129(4):1899–1946.
- Hoffman, M. and Burks, S. V. (2017). Training contracts, employee turnover, and the returns from firm-sponsored general training. Technical report, National Bureau of Economic Research.
- Hoffman, M. and Tadelis, S. (2021). People management skills, employee attrition, and manager rewards: An empirical analysis. *Journal of Political Economy*, 129(1):243–285.
- Irwin, D. A. and Klenow, P. J. (1994). Learning-by-doing spillovers in the semiconductor industry. *Journal of political Economy*, 102(6):1200–1227.
- Jarosch, G., Oberfield, E., and Rossi-Hansberg, E. (2021). Learning from coworkers. *Econometrica*, 89(2):647–676.

- Kelley, E. M., Lane, G., and Schönholzer, D. (2021). Monitoring in target contracts: Theory and experiment in kenyan public transit.
- Krugman, P. (1979). A model of innovation, technology transfer, and the world distribution of income. *Journal of political economy*, 87(2):253–266.
- Kugler, M. and Verhoogen, E. (2012). Prices, plant size, and product quality. *The Review of Economic Studies*, 79(1):307–339.
- Levitt, S. D., List, J. A., and Syverson, C. (2013). Toward an understanding of learning by doing: Evidence from an automobile assembly plant. *Journal of political Economy*, 121(4):643–681.
- Limodio, N. (2021). Bureaucrat allocation in the public sector: Evidence from the world bank. *The Economic Journal*, 131(639):3012–3040.
- Macchiavello, R., Menzel, A., Rabbani, A., and Woodruff, C. (2020). Challenges of change: An experiment promoting women to managerial roles in the bangladeshi garment sector. Technical report, National Bureau of Economic Research.
- Metcalf, R. D., Sollaci, A. B., and Syverson, C. (2023). Managers and productivity in retail. Technical report, National Bureau of Economic Research.
- Minni, V. (2022). Making the invisible hand visible: Managers and the allocation of workers to jobs. *Job Market Paper*.
- Sandvik, J., Saouma, R., Seegert, N., and Stanton, C. (2022). Should workplace programs be voluntary or mandatory? evidence from a field experiment on mentorship.
- Sandvik, J. J., Saouma, R. E., Seegert, N. T., and Stanton, C. T. (2020). Workplace knowledge flows. *The Quarterly Journal of Economics*, 135(3):1635–1680.
- Thompson, P. (2010). Learning by doing. *Handbook of the Economics of Innovation*, 1:429–476.
- Verhoogen, E. A. (2008). Trade, quality upgrading, and wage inequality in the mexican manufacturing sector. *The Quarterly Journal of Economics*, 123(2):489–530.
- Verhoogen, E. A. (2021). Firm-level upgrading in developing countries. *NBER Working Paper*, (w29461).
- Vernon, R. (1966). International investment and international trade in the product cycle. *The Quarterly Journal of Economics*, 80(2):190–207.