MASTER PROJECT

Industrial Robots and Where to Find Them: Evidence and Theory on Derobotization

Amil Camilo, Doruk Gokalp, Daniil Iurchenko, Julian Klix, and Jeremy Rubinoff

Barcelona Graduate School of Economics
Master Degree in Economics and Finance
2019-2020

July 22, 2020

Abstract

Using firm-level data from Spain, we investigate robot abandonment, a phenomenon neglected by the literature, and find that a substantial proportion of robot adoption is non-permanent. We also find that (i) firms are most likely to derobotize shortly after robotization; (ii) derobotization rates are higher among smaller firms; and (iii) labor demand falls after derobotization. We develop a model of reversible automation in which firms learn the costs of using robots only after they first adopt them. We simulate a panel of firms that match the data and demonstrate that analyses of automation which ignore abandonment can overestimate the effects of automation.

Contents

1 Introduction

Robots have emerged as a remarkable feature of modern production. Around the world, and especially in high-tech economies, the demand and adoption of industrial robots have increased dramatically (Autor, 2015; Brynjolfsson & McAfee, 2014; Koch et al., 2019). However, the abandonment of robots in the production process (henceforth referred to as *derobotization* or *deautomation*) has been less discussed. We find that a substantial proportion of manufacturing firms deautomate, a fact which has been overlooked by the literature. We propose a model of automation that allows firms to not only adopt robots but also to abandon them. We calibrate our model and closely match the behavioral distribution of automation in Spanish firm-level data. In our simulations, most firms never automate, but of those which do, approximately two-fifths stay automated, two-fifths deautomate (and never automate again), and one fifth switch back and forth.

In our investigation, we use data from SEPI Foundation's *Encuesta sobre Estrategias Empresariales* (ESEE), which annually surveys over 2000 Spanish manufacturing firms on business strategies, including on whether the firms adopt robots in their production line. Extending the analysis of Koch et al. (2019), we document three major facts on derobotization. First, firms that derobotize tend to do so quickly, with over half derobotizing in the first four years after adoption. Second, derobotizing firms tend to be relatively smaller. Third, firms that abandon robots demand less labor and experience an increase in their capital-to-labor ratios. We hypothesize that the prompt abandonment of robots is indicative of a learning process in which firms robotize production with expectations of higher earnings, but later learn information which causes them to derobotize and adjust their production accordingly.

With this in mind, we propose a dynamic model of manufacturing firms that allows firms to both adopt robots and later derobotize their production. In our model, firms face a sequence of optimal stopping problems where they consider whether to robotize, then whether to derobotize, then whether to robotize again, and so on. The production technology in our model is micro-founded by the task-based approach common in this literature, where firms assign tasks to workers of different occupations as well as robots (as opposed to the factor-augmenting approach). We assume two occupations, that of low-skilled and high-skilled workers, where the latter workers are naturally more productive than the former. When firms robotize, both the firm's overall productivity and the relative productivity of high-skilled workers increases, but the relative productivity of low-skilled workers decreases. In addition to observing this productivity increase, firms which robotize also learn the total cost of maintaining robots in production which may be greater than their initial expectations. At any point in time, firms can return to the derobotized regime

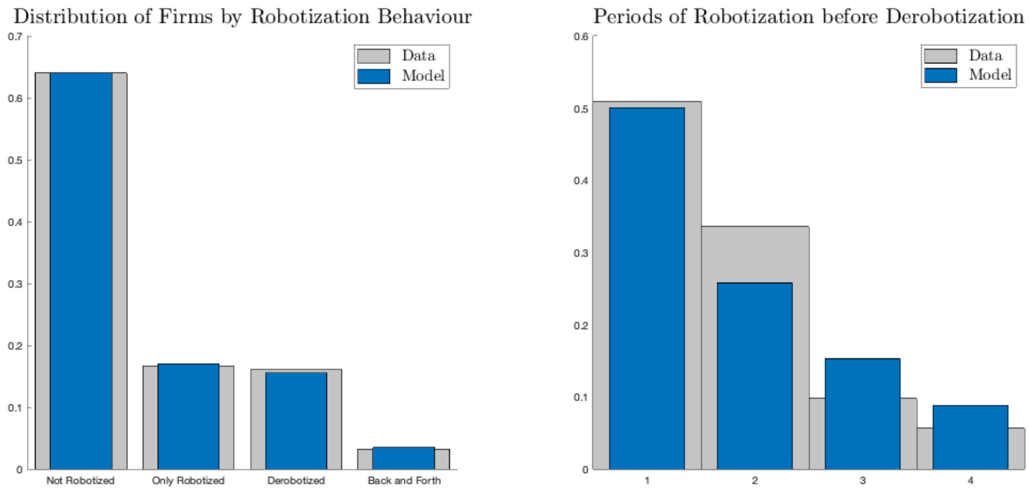

Figure 1: The distribution (left) of firms by robotization behavior, and the distribution (right) of firms' derobotization behavior (if they derobotize). The data from ESEE is in gray, while the calibrated simulation of our model is in blue.

but now with the knowledge of the costs of robotization, and a lower cost of reautomation since we assume firms retain the infrastructure of operating robots in production.

We simulate a panel of manufacturing firms in the model closely matching our findings in the data. Indeed, our simulations show that larger and more productive firms are more likely to robotize and that the firms which derobotize tend to be less productive. Our calibration of the model can accurately explain and reproduce the behavioral distribution of automation across firms in the data (see Figure 1). We conclude that omitting derobotization from our analysis can lead to an overestimation of the effects of automation.

The structure of the paper is straightforward. In Section 2, we offer a brief overview of the latest literature on robots and automation in economics. In Section 3, we detail our look into the ESEE dataset and elaborate on the trends of derobotization across Spanish manufacturing firms. In Section 4, we formally introduce our model and explicitly define the structure of firms and robot adoption. In Section 5, we analyze our model through the results of our calibrated simulation and compare the outcomes of our model with that of a permanent adoption baseline, as well as layout crucial parametric assumptions particular to Spanish firms. Finally, in Section 6, we discuss the policy implications of our results and conclude by offering insight on future research as well as the shortcomings of our investigation. Appendices are available to the reader, the contents of which will be referenced throughout the paper.

2 Literature review

Like most recent papers in the literature, the firm's production technology in our model comes from Acemoglu and Autor's (2011) paper on task-based production. In that paper, the authors argue that the neoclassical factor-augmenting approach to production is insufficient to explain key dynamics in the labor market, specifically how technological innovations may displace low-skilled workers and expand the offshoring of labor. They propose modeling production instead through a task-based approach, where firms assign a set of tasks to be completed by workers and capital (in this case, through automation), while other tasks can be completed by labor alone. This approach is able to capture these key dynamics and, after being standardized by Acemoglu and Restrepo (2018a), has become the workhorse model of the robots and automation literature in macroeconomics. Our project closely follows a number of papers building from this approach, including recent work from Acemoglu and Restrepo (2018b) and Acemoglu et al. (2020), though mainly papers from Koch et al. (2019) and Humlum (2019). Most models tend to make two simplifying assumptions, the first of which both these papers address: (i) robot adoption is homogeneous across firms which implies that when a task becomes possible and profitable to complete with only capital instead of labor, all firms immediately reassign that task to capital; (ii) robot adoption is permanent which means that once a firm completes a task with robots, it will continue to do so in the future.

Koch et al. (2019) investigate heterogeneity in adoption by making minor modifications to the workhorse task-based model and using a richer panel from the same ESEE dataset we use.[1] They show that larger and more productive firms adopt robots more frequently; upon adoption, the differences between larger and smaller firms widen across time, with significant job losses in non-robotized firms. They also find that robot adoption leads to net job creation for both low- and high-skilled workers, a fact which we exploit to motivate our model. Although not central to their paper, they also find that 38% of adopters in their sample derobotize at some point, which directly motivates our investigation. Empirically, our paper continues their work since we extend their analysis, focusing on abandonment instead of adoption.

The paper from Humlum (2019) is the closest to our paper in terms of theory. He simulates a dynamic general equilibrium setting where firms behave as if they solve an optimal stopping problem, choosing whether to adopt robots or proceed with production in a non-robotized state. The basic trade-off for firms in his paper is similar to ours: robotization has an adoption cost which is stochastic, but guarantees certain improvements in produc-

[1] Due to the budget constraints of the Master Project, we based our analysis on a data-constrained panel with fewer variables.

tivity. The production function in his model is tractable to value function iterations and is micro-founded by the task-based approach from Acemoglu and Restrepo (2018a), which allows him to estimate labor market dynamics in Denmark. The benchmark model from which we extend is a special case of his model, where we narrow the number of occupations from three to two. We perform ostensibly the same simulations, in addition to the simulations in our derobotization framework. The main difference between his paper and ours is that automation is reversible in our model.

Recent works which are worth mentioning include papers by Acemoglu and Restrepo (2017), Ocampo (2018), Frey and Osborne (2017), Brynjolfsson et al. (2018), and again, Acemoglu et al. (2020). Acemoglu and Restrepo (2017) apply their task-based approach to robots and find robust results that suggest robot adoptions are correlated with job loss in commuter zones. Ocampo (2018) generalizes the theory of task-based production to allow for an arbitrary number of occupations and tasks by solving an optimal transport problem, which is an especially useful framework for understanding task assignment models. Frey and Osborne (2017) estimate the propensity for future jobs to be replaced by capital (in their own words, for jobs to be "computerized") by using classification algorithms.[2] Brynjolfsson et al. (2018) cover similar ground by estimating the propensity for future jobs to be replaced by machine learning algorithms and artificial intelligence. Acemoglu et al. (2020) look at French manufacturing firms and find some of the same results from Koch et al. (2019), confirming that robot adoption widens the gap between large and small firms, but finding that the overall impact of robot adoption on industry employment is negative.

3 Firm-level facts on derobotization

The ESEE (Survey on Business Strategies, in English) is collected by the SEPI Foundation each year, surveying Spanish manufacturing firms on details regarding their operations. A more extensive survey is administered every four years, and between 2,000 and 2,600 firms participate in it. Crucially, it is this extensive form which tracks whether firms employ robots in their production process.

We acquired data from the extensive survey each time it was administered, first in 1991 and then every four years between 1994 and 2014. The dataset is an unbalanced panel which features both left- and right-censoring. Initially the dataset contains 15,929 observations across 5,588 firms, which reduces to 14,119 observations and 3,987 after exclusions. These observations capture 958 robotizations and 755 derobotizations. A more detailed description of the dataset and the variables used, including exclusion criteria,

[2] Interestingly enough, they estimate that economists are 43% computerizeable.

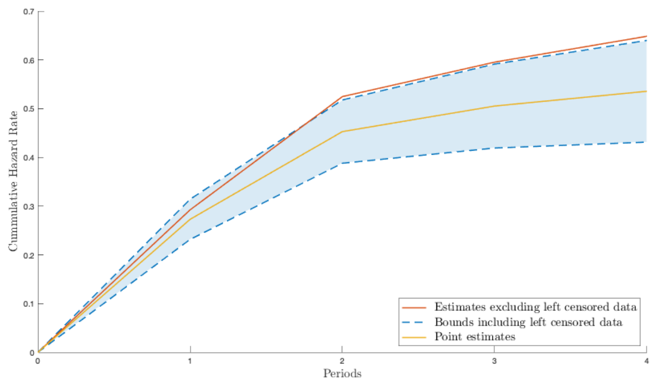

Figure 2: Kaplan-Meier estimates of derobotization rates

The blue area provides the upper and lower bounds when incorporating left-censored data, while the red line is the estimate when excluding left-censored data.

counts of behavior, and summary statistics, can be found in Appendix A.

3.1 Kaplan-Meier estimation

The duration of robotization states can be right-censored due to firms liquidating, ceasing to participate in the survey, or because of the data collection ending in 2014. As a result, directly counting the number of firms which report changing their robot use, as done in Koch et al. (2019), will underestimate the true values of derobotization. In order to determine which firms derobotize, conditional on the firms' continued reporting, we use a Kaplan-Meier survival estimator (KM) to evaluate the probability of derobotization in every period.

Our KM estimates incorporate left-censored data to provide bounds on the cumulative derobotization rate. We get an upper bound by assuming no gap between robotization and entry into the data set, and a lower bound by assuming an infinite gap. Derobotization rates are estimated only for the first four periods (sixteen years) after robotization, as observations are too sparse to estimate further. These results are shown, along with an estimate that excludes all left-censored data, in Figure 2.

The estimate which excludes left-censored data fluctuates outside the upper bound, suggesting that non-left-censored data does not adequately capture derobotization rates. Therefore, we use the midpoint between the upper and lower bounds as our best estimate for true derobotization rates. We also conduct this analysis for robotization rates over five periods. Details of the robotization and derobotization analyses can be found in Appendix A.

Using the KM estimates for rates of robotization and derobotization, we calculate the percentage of firms expected to engage in different types of robotization behavior over seven periods, and compare these in Table 1 to the percentages expected from naive counting. The impact of right-censoring on suppressing derobotization rates stands out as a major difference. These results indicate the prevalence of abandonment; ex-ante we expect it to occur among 19.28% of all firms and 53.56% of firms who robotize:

Fact 1. *A substantial proportion of robot adoption is non-permanent.*

Table 1: Robotization Behaviour

	KM Estimates	Naive Estimates
Never Automated	64.01%	58.52%
Only Automated	16.71%	26.97%
Deautomated	16.05%	11.20%
Reautomated	3.23%	3.31%

The KM estimates are also used to calculate the expected timing of derobotization, shown in Table 2 alongside an estimate without the right-censoring adjustments. Not only does the right-censoring suppress derobotization rates, it also underestimates the expected time before derobotization among the firms that do derobotize. Even for the KM estimates, over half of all derobotizations occur in the period immediately following robotization, and more than five in six occur within the first two periods.

Fact 2. *Derobotization is most likely in the first periods after adoption.*

Table 2: Derobotization Timing

Periods after Robotization	Derobotization % (KM)	Derobotization (Naive)
1	50.97%	64.64%
2	33.58%	26.62%
3	9.79%	5.56%
4	5.66%	3.18%

3.2 Derobotization and firm size

To document the profile of derobotizing firms, we rely on survival analysis techniques which account for both the time- and state-dependence of firms' decisions. We restrict our sample to firms which were not robotised when entering the survey and aggregate the data across firms. As a result, we obtain data on the duration of robot use as well as the median characteristics of the firms over the timeline.

To estimate the duration of robot usage, we use the Cox (1972) proportional hazards model and define derobotization as an exit event. Explicitly, the hazard function for firm

i takes the form

$$\lambda(t_i|X_i) = \lambda_0(t_i) \times \exp\left(\beta_1 \text{size}_i + \beta_2 \text{prevrob}_i + \beta_3 log(K/L)_i\right), \qquad (1)$$

where t_i is the duration of firm i's robot usage, $\lambda_0(t_i)$ is the baseline hazard rate. size_i is a categorical variable measuring firm size by the number of employees, prevrob_i is a dummy variable indicating whether the firm has robotized before, and $\log(K/L)_i$ is the logarithm of the firm's capital-to-labor ratio (summarized by X_i).

The resulting estimates are presented in Table 3. Column 1 presents a bivariate regression against firm size. Columns 2 to 3 have specifications which include previous robotization and capital-to-labor ratios. Columns 4 and 5 have specifications where the firm's size is measured by total labor hours or capital, respectively.

Table 3: Duration of robot usage and firm size

	(1)	(2)	(3)	(4)	(5)
size					
From 100 to 500	-0.42***	-0.42***	-0.35**		
	(0.15)	(0.15)	(0.15)		
From 500	-0.47**	-0.46**	-0.44**		
	(0.21)	(0.21)	(0.22)		
prevrob		-0.54*	-0.47	-0.48	-0.50*
		(0.30)	(0.30)	(0.30)	(0.30)
$\log(K/L)$			-0.13*		
			(0.07)		
log(capital)				-0.12***	
				(0.03)	
log(hours)					-0.15***
					(0.05)
Observations	786	786	777	777	777

The sample includes duration of robot usage for firms that: (a) were robotized at least once during 1991-2014; (c) did not enter the sample being robotized. Baseline category of size is "Less than 100 employees". Capital is adjusted to 1991 prices. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

In each specification, our estimates consistently show a negative duration dependence as measured by the baseline hazard function (see Figure 10 in Appendix A.5). In the first specification, the effect of the firm's size on the hazard rate is negative and statistically significant. Adding the indicator on previous robotization, the coefficients on size still remain negative and significant. Moreover, there is also a negative effect of previous robotization on the hazard rate. Including the log capital-to-labor ratio decreases the magnitude of the size coefficient but does not affect its significance level; on the other hand, the coefficient on previous robotization becomes indistinguishable from zero. Finally, when we add alternative measures of size, the results remain unchanged. Overall, our findings suggest a statistically significant positive correlation between firm size and robot usage duration.

Fact 3. *Larger firms are less likely to derobotize.*

In addition, there is mixed evidence that firms which previously robotized are less likely to derobotize. The intuition of Fact 2 may suggest the presence of a learning process in which firms, after initial robot adoption, may learn new information which affect their future robotization decisions. However, the inclusion of the capital-to-labor ratio absorbs this effect and makes it statistically negligible.

We also provide several robustness tests (see Appendix A.3-A.4). To account for potential selection bias, we extend our sample and add firms which were already robotized while entering the survey. With this extended sample, we are also able to control for heterogeneity across industries.[3] Likewise, we also measure the size of firms before derobotization events and estimate a pooled regression with a richer set of controls. All in all, we conclude that the negative correlation between size and derobotization is robust across different samples and estimation strategies.

3.3 Firm outcomes and event studies

Given the limitations of our data, we adopt a basic event study methodology to measure changes in capital, total work hours, and the capital-to-labor ratio after derobotization. We consider firms which robotize but do not switch back and forth, restricting our sample to the post-robotization period. It is important to highlight that this analysis is conditional on firms already having robotized.

We estimate the following equation:

$$y_{it} = \beta_1 \text{derobotization}_{it} + \mu_i + \tau_t + \varepsilon_{it}, \tag{2}$$

where y_{it} is the outcome for firm i in time t, $\text{derobotization}_{it}$ is a dummy variable for derobotization, μ_i is the firm's fixed effect over time, τ_t is a time fixed effect, and ε_{it} is an

[3] In our initial sample, there are, on average, 30 observations for each industry (20 industries), which may suggest overfitting on our part.

error disturbance. We estimate the response for when y_{it} is the logarithm of respectively capital, total work hours, and capital-to-labor ratio of the firm. The fixed effect from μ_i accounts for post-robotization time-invariant firm characteristics which may include the costs and efficiency of robot usage.[4]

Since our data has a four-year frequency, some firms could derobotize by the end of the period which may affect the response with a short lag. To account for this, we modify equation (2) and add lagged derobotization, so

$$y_{it} = \beta_1 \text{derobotization}_{it} + \beta_2 \text{derobotization}_{it-1} + \mu_i + \tau_t + \varepsilon_{it}. \tag{3}$$

The results of our estimation are displayed in Table 4. For each outcome and specification, we report the coefficients for current and lagged derobotization and compute the cumulative effect of derobotization. In all regressions, standard errors are clustered at the firm level. The specifications for capital report small and insignificant coefficients of derobotization, while the specifications for total work hours report negative and statistically significant coefficients. In particular, the total number of work hours employed by firms shrinks by approximately 9.5-15% in the eight years following derobotization.[5] The combination of both these responses induce the positive effect on capital-to-labor ratios observed in specifications 5 and 6. Specifically, the capital-to-labor ratio of firms rises by approximately 9.5-16% in the eight years following derobotization. We thus conclude:

Fact 4. *Derobotization is correlated negatively with labor inputs and positively with capital-to-labor ratios.*

In general, these results are consistent with Koch et al. (2019) who find that robot adoption leads to considerable productivity gain and net job creation for both low- and high-skilled workers. Since their result imply short-term complementarity between labor and capital, one can expect that firms would decrease the amount of labor after derobotization. On the other hand, the lack of adjustment in the capital margin is less intuitive. This may be explained by a change in capital structure after derobotization which cannot be captured with the low-frequency data we have. We leave this avenue, together with examining labor changes across occupations, for future research.

[4] To estimate the effects on labor and capital more accurately, it would be better to use the lagged response as an instrument as it may provide information about shocks which affect selection in derobotization. However, this is not feasible due to the significant loss of observations needed to use this instrument.

[5] We calculate the approximate effect of derobotization as $e^{\beta} - 1$.

Table 4: Derobotization and Firm Outcomes

	log(*capital*)		log(*hours*)		log(*K*/*L*)	
Specification	(1)	(2)	(3)	(4)	(5)	(6)
derobotization$_t$	-0.00	-0.00	-0.09**	-0.08**	0.09*	0.09*
	(0.06)	(0.06)	(0.04)	(0.04)	(0.05)	(0.05)
derobotization$_{t-1}$		0.02		-0.06		0.06
		(0.05)		(0.03)		(0.04)
cumulative effect	0.00	0.02	-0.09**	-0.14**	0.09*	0.15**
p-value of Wald test	0.98	0.88	0.03	0.01	0.05	0.02

Notes. The sample includes firms that: (a) are observed at least for 3 periods; (c) adopted robots over 1991-2014; (d) did not switch back and forth. The first observation of a firm corresponds to a year when it was observed using robots for the first time. The number of firms and observations are equal to 647 and 1822, respectively. All regressions include time and firm fixed effects. Standard errors are clustered on firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4 Model

As in Humlum (2019), we consider a dynamic partial equilibrium model of manufacturing firms. At every period, firms select the quantities of labor and intermediate inputs to use in production as well as whether to use robot technology.[6] Since our model is an extension of his permanent adoption setting, we begin by introducing a stylized version of his model and later allow firms to revert their automation in ours.

4.1 Baseline model with permanent adoption

Consider a manufacturing firm which employs capital and workers of different occupations. Time is discrete ($t = 0, 1, 2, \dots$) and future profits are discounted by the factor $\beta < 1$. At every period, firms can produce without robots or adopt robots into production. The firm's objective is to maximize expected profits.

The firm's output Y_t is given by a task-based CES production function $F(\cdot|R, \varphi)$;[7] explicitly,

$$Y_t \equiv F(M_t, L_t|R_t, \varphi_t) := z_{Ht}(\varphi_{Ht}, R_t)\Big(M_t^{\frac{\sigma-1}{\sigma}} + \sum_{o\in O} z_{ot}(\varphi_{ot}, R_t)^{\frac{1}{\sigma}} L_{ot}^{\frac{\sigma-1}{\sigma}}\Big)^{\frac{\sigma}{\sigma-1}}. \tag{4}$$

The firm demands labor $L_t \in \mathbb{R}_+^{|O|}$ and intermediate inputs $M_t \in \mathbb{R}_+$. The set of occupa-

[6]We will use intermediate inputs interchangeably with capital.

[7]The task-based micro-foundation of this production function is available in Appendix B.

tions $O = \{1,2\}$ includes low-skilled ($o = 1$) and high-skilled workers ($o = 2$). R_t is a binary variable which is 1 if the firm is using robots, and 0 if not. The baseline productivities of the firm are given by a vector $\varphi_t = (\varphi_{Ht}, \varphi_{1t}, \varphi_{2t}) \in \mathbb{R}^{|O|+1}$ of mutually independent exogenous stochastic processes. The firm's Hicks-neutral productivity z_{Ht} and the relative productivity of workers of different occupations z_{ot} are given by

$$z_{Ht}(R_t, \varphi_t) := \exp(\varphi_{Ht} + \gamma_H R_t), \tag{5}$$

$$z_{ot}(R_t, \varphi_t) := \exp(\varphi_{ot} + \gamma_o R_t), \tag{6}$$

where γ_H captures the effect of robot technology on the firm's productivity, while γ_o captures the effect of robot technology on a worker of occupation o. The parameter σ has the traditional interpretation as the elasticity of substitution between factors. The firm sells its output in the market at an iso-elastic price $P_t = P_M(Y_M/Y_t)^{1/\varepsilon}$, where P_M is the manufacturing price index, and Y_M is the aggregate manufacturing demand.[8]

The firm takes the vector of factor prices $w \in \mathbb{R}_+^{|O|+1}$ as given. To shorten notation, define $X_t \equiv (M_t, L_t)$ as a vector of inputs and $\pi_t(R, \varphi)$ as the firm's per-period profits from choosing X_t optimally, given the state of robotization and productivity.[9] The firm's decision to robotize (or derobotize) production takes one period to be implemented; consequently, at every period t, firms receive $\pi_t(R_{t-1}, \varphi_t)$. We assume that $R_0 = 0$ so that firms always start in a non-robotized stage. The firm can robotize production by purchasing robots for a price $p_R \in \mathbb{R}_+$ which is constant over time. Upon operation, robots are costly to maintain with a per-period cost of $C_R + \varepsilon_R$, where C_R is the expectation of the cost and ε_R is an independent zero-mean random variable, representing a firm-characteristic, time-invariant cost shock that remains unknown until the firm robotizes.[10] It captures the firm's inherent suitability for robotization, which is *ex ante* unknown to the firm. Once the firm is robotized, it receives only robotized profits $\pi(1, \varphi)$ for all periods forward, as is the case in Humlum (2019).[11]

[8] In Humlum (2019), these parameters are time-varying processes. For simplicity (in the notation of that paper), we set $P_{Mt} = P_M$ and $Y_{Mt} = Y_M$.

[9] Formally, $\pi_t(R, \varphi) := \max_{X \in \mathbb{R}_+^{O+1}} \left\{ P_M Y_M^{\frac{1}{\varepsilon}} F(X|R, \varphi)^{\frac{\varepsilon-1}{\varepsilon}} - w^T X \right\}$.

[10] The cost shocks in this stylized version are markedly different than the ones in Humlum's. In his framework, firms only face cost shocks for initial adoption whereas in our model, the initial adoption cost is captured entirely by p_R. In his model, firms can costlessly operate robots once adoption costs are paid. Even so, before we introduce reversibility of automation, these differences are negligible. The fact that firms in Humlum's paper can never revert robotization means that the revealed information from our cost shocks would not change the behavior of firms, since they only act on their ex-ante expectations. If we were to include the cost shocks from Humlum in our permanent adoption framework, this would increase the lumpiness of adoption but leave the incentives unchanged.

[11] We can unify this baseline with the model presented in Section 4.2 by assuming $\delta_P \longrightarrow \infty$ and $\delta_S < \infty$ in the extended framework, effectively removing the firm's incentive to derobotize.

The value of the firm when non-robotized is $V_I \equiv \max\{V_I^a, V_I^i\}$ where V_I^a is the value when $R_t = 1$ while V_I^i is the value when $R_t = 0$. If the firm robotizes at period t, at period $t+1$ the value of the robotized firm is $V(\varphi_{t+1}, \varepsilon_R)$. For every subsequent period, the value of the firm then only varies with its productivity since robotization is irreversible.

Explicitly, the value functions of the non-robotized firm are described by the following Bellman equations,

$$V_I^a(\varphi_t, \varepsilon_R) = \pi_t(0, \varphi_t) - p_R + \beta \mathbb{E}[V(\varphi_{t+1}, \varepsilon_R)|\varphi_t] \tag{7}$$

$$V_I^i(\varphi_t, \varepsilon_R) = \pi_t(0, \varphi_t) + \beta \mathbb{E}[V_I(\varphi_{t+1}, \varepsilon_R)|\varphi_t], \tag{8}$$

while the value function of the robotized firm are described by

$$V(\varphi_t, \varepsilon_R) = \pi_t(1, \varphi_t) - (C_R + \varepsilon_R) + \beta \mathbb{E}[V(\varphi_{t+1}, \varepsilon_R)|\varphi_t, \varepsilon_R]. \tag{9}$$

4.2 Robotization and derobotization (reversible adoption)

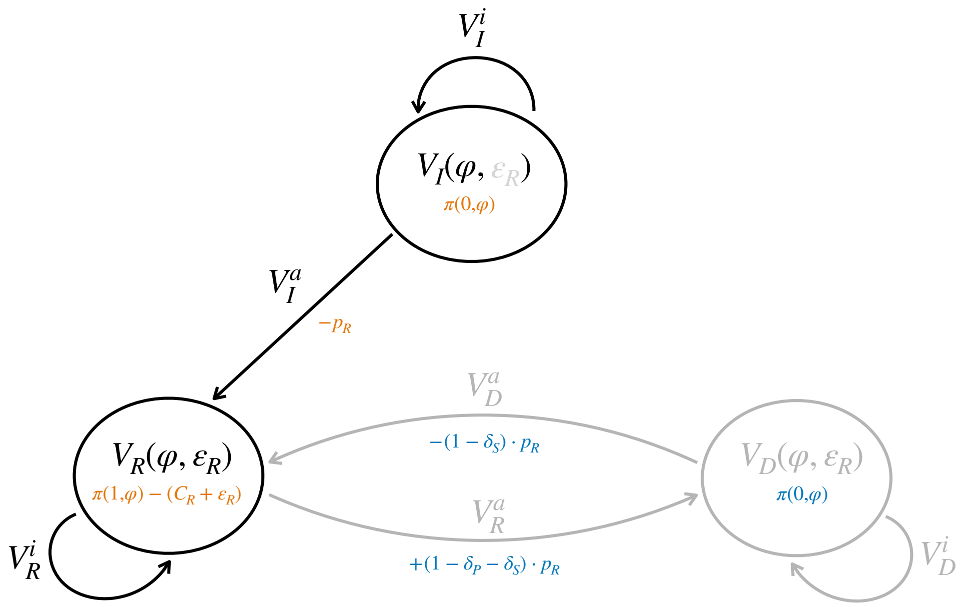

Figure 3: Overview of the structure and timing of the firm's value functions.

The black section depicts the value function transitions of the baseline model, the gray section those of the extensions of the full model. V_R collapses to the simplified V in the baseline when $\delta_P \longrightarrow \infty$ and $\delta_S < \infty$.

We now enable firms to abandon robots by selling them at a lower price than the one they paid. Indeed, the resale price is discounted by $\delta_P + \delta_S$, with $\delta_S, \delta_P \in \mathbb{R}_+$, so that robotized firms sell at the price $(1 - \delta_P - \delta_S)p_R > 0$. The parameter δ_P is the depreciation factor of used robots and the parameter δ_S is the fraction of setup costs of robot adoption which is paid only when firms robotize for the first time. It captures the initial sunk costs of robot adoption, for instance, the training of workers and other intangible investments, that cannot be resold, but in contrast to regular depreciation δ_P, remain in place when rerobotizing and thus need not be paid again.[12] Therefore, when firms sell their robots, the reselling price $(1-\delta_P-\delta_S)p_R$ is reflective of both aspects of depreciation. The firm's incentive to derobotize is straightforward: if cost shocks are significantly higher than expected (or productivity shocks are significantly lower), then maintaining robots may not be profitable, even with improved average productivity, and so firms would prefer to revert to derobotized production.

Once firms are derobotized, they return to receiving derobotized profits but now with the knowledge of the cost shock and the initial setup already paid. This latter fact is crucial as it makes rerobotization less costly, and consequently, when firms rerobotize, they only pay $(1-\delta_S)p_R$. These savings allow firms with high cost shocks and high productivity shocks to switch back and forth between robotized and derobotized production.

Our model is summarized by Figure 3. The value of the robotized (or derobotized) firm at period t is given by $V_R \equiv \max\{V_R^a, V_R^i\}$ (or $V_D \equiv \max\{V_D^a, V_D^i\}$) where V_R^a (or V_D^a) is the value when $R_t \neq R_{t-1}$ (*action*) while V_R^i (or V_D^i) is the value when $R_t = R_{t-1}$ (*inaction*). If the firm changes their regime at period t, they transition from V_D to V_R (or V_R to V_D). They can do so every period, earning higher output profits when robotized but saving costs when derobotized.

Explicitly, the value functions of the robotized firm are described by the following Bellman equations,

$$V_R^a(\varphi_t, \varepsilon_R) = \pi_t(1, \varphi_t) - (C_R + \varepsilon_R) + (1-\delta_P-\delta_S)p_R + \beta\mathbb{E}[V_D(\varphi_{t+1}, \varepsilon_R)|\varphi_t, \varepsilon_R], \quad (10)$$

$$V_R^i(\varphi_t, \varepsilon_R) = \pi_t(1, \varphi_t) - (C_R + \varepsilon_R) + \beta\mathbb{E}[V_R(\varphi_{t+1}, \varepsilon_R)|\varphi_t, \varepsilon_R] \quad (11)$$

while the value functions of the derobotized firm are described by

$$V_D^a(\varphi_t, \varepsilon_R) = \pi_t(\varphi_t) - (1-\delta_S)p_R + \beta\mathbb{E}[V_D(\varphi_{t+1}, \varepsilon_R)|\varphi_t, \varepsilon_R] \quad (12)$$

$$V_D^i(\varphi_t, \varepsilon_R) = \pi_t(\varphi_t) + \beta\mathbb{E}[V_R(\varphi_{t+1}, \varepsilon_R)|\varphi_t, \varepsilon_R] \quad (13)$$

[12] Alternatively, setup costs can be thought of as installing the infrastructure to operate robots, assuming it cannot (or will not) be resold and is costless to maintain.

5 Simulation

We analyze the model introduced in Section 4 in two steps. First, we numerically solve the Bellman equations (7) - (13) using value function iteration and obtain the policy functions for robotization R as well as the input vector X. Second, we simulate a panel of firms by using the policy functions on multiple time series of exogenous shocks to examine the behavior of firms. Using these simulations, we match the findings in Section 3 and generate event studies for robotization and derobotization. Furthermore, we contrast our findings with the behavior and predictions of the baseline permanent adoption model.

5.1 Value function iteration

We begin by discretizing the exogenous state φ to Markov processes using the Rouwenhorst (1995) method, as described by Kopecky and Suen (2010). We choose the support of φ_o so that $\varphi_1 < \varphi_2$. Likewise, we set $\gamma_1 < 0 < \gamma_2$, since robotized firms in the data (as analyzed by Koch et al.) employ more capital and high-skill labor (relative to low-skill labor) than firms which are not robotized; our interpretation of this fact is that low-skilled workers become less productive (relative to both high-skilled workers and intermediate inputs) after robotization.[13]

One of the backbones of our calibration is the fact that we set $\sigma > 1$. In Humlum's simulations to match firm-level Danish data, he calibrates σ to be 0.493. However, in the data from Spain, we find that robotized firms not only increase their overall labor demand, but they also increase their low-skill labor demand, a fact which is also documented by Koch et al. Therefore, high- and low-skilled labor are complements according to the micro-data we observe, and it is paramount to set $\sigma > 1$.[14]

The resulting policy functions for R from value function iteration are shown in Figure 4 as action-inaction regions, dependent on the cost-type ε_R and the productivity profile φ. With this figure, we are able to concisely capture the complete robot adoption behavior of firms in the multi-dimensional state space. We use the Hicks-neutral baseline productivity φ_H as an indicator for the size of the firm on the grounds that (as Humlum argues) higher productivity, ceteris paribus, allows firms to employ more labor and capital and hence leads to larger firms. Likewise, the converse can be argued so that large firms tend to exhibit higher productivity.

According to the action regions of initially unrobotized firms V_I, robot adoption happens

[13]Implicitly, the productivity of intermediate inputs is normalized to one and incorporated into $z_H(\cdot)$, making z_1 and z_2 represent the productivity of low- and high-skilled labor, relative to the intermediate.

[14]It's worth remarking that, as Humlum documents, there are no good estimates (for now) of the micro elasticity of substitution between workers in the task-based approach.

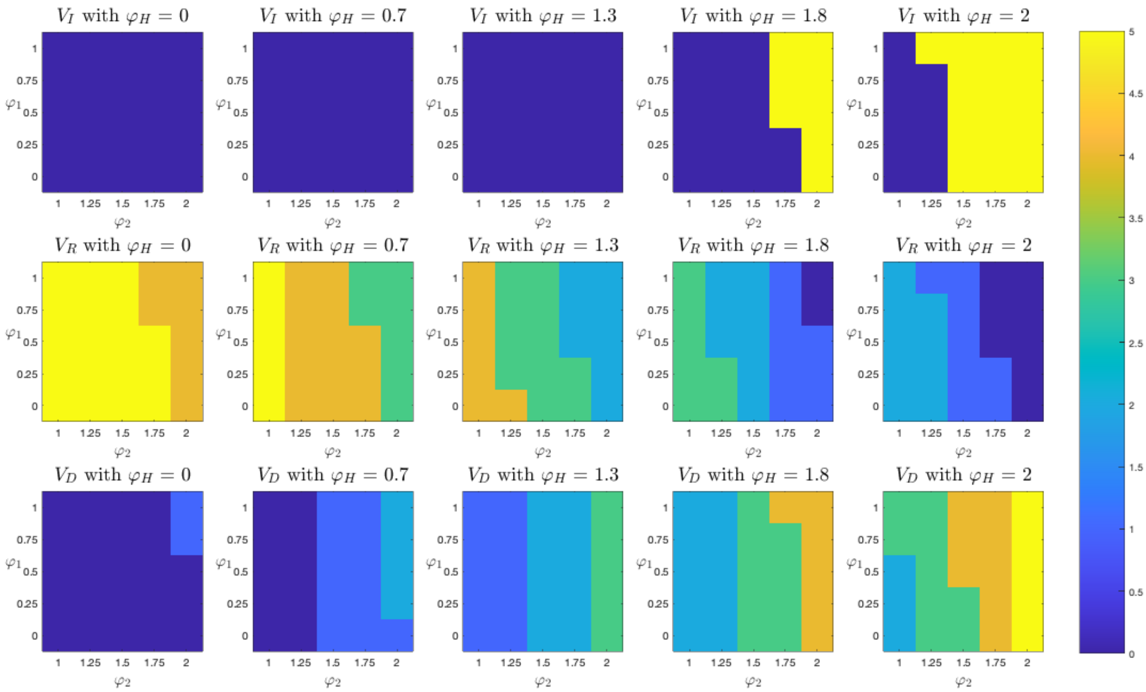

Figure 4: Heatmaps of action-inaction regions

This figure illustrates the action-inaction regions implied by the value functions. The coloring corresponds to the number of different cost types taking action in the associated state profile. The first row represents the robotization decision of a firm that never automated (V_I), the second the derobotization decision (V_R) and the last the rerobotization decision (V_D). Note that the abandonment (readoption) action region of V_R (V_D) for firms with cost shock realization ε_R' encompasses the action region for all firms with $\varepsilon_R < \varepsilon_R'$ ($\varepsilon_R > \varepsilon_R'$).

in states with high productivity, and thus in particular among large firms.[15] As a corollary, firms are most likely to derobotize when productivity is low:

Observation 1. *Less productive firms are more likely to derobotize.*

Given the link between firm size and productivity, derobotization is more frequent among small firms, which is consistent with Fact 3. We also observe, as intuition would suggest, that firms with a greater characteristic cost ε_R are more likely to derobotize; this is additionally illustrated in Figure 5.

Our simulations exhibit two driving forces of derobotization: a *productivity effect* arising from a sequence of low productivity shocks φ, and a *revelation effect* arising from learning about high characteristic costs ε_R of robotization.

We leave the heatmaps corresponding to the policy functions for inputs in Appendix C. In short, our simulation finds conventional comparative statics consistent with our assumptions.

[15] This matches the already documented observation in the literature (Koch et al., 2019), that adoption is positively correlated with firm size.

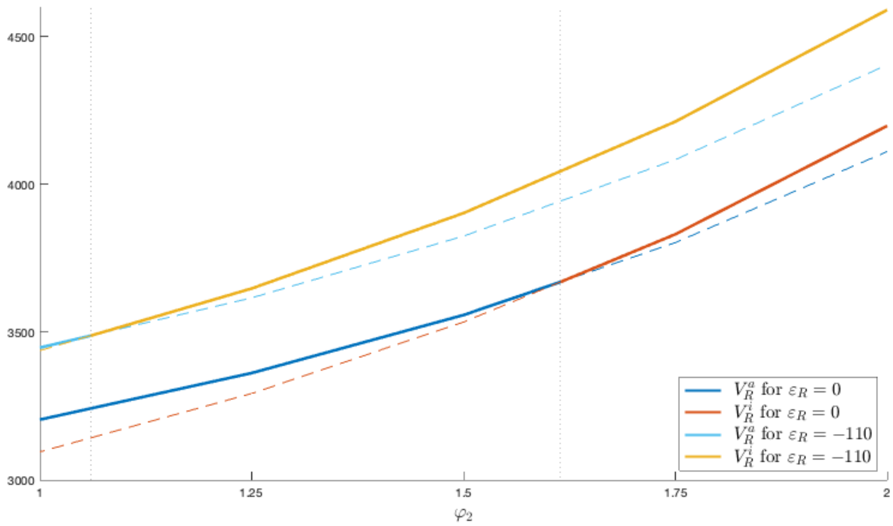

Figure 5: Action-inaction region of cost-differing automated firms
The solid line represents the value of being robotized V_R and comprises an action (left) and inaction region (right). The action region for derobotization increases with realized costs ε_R.

Observation 2. *Robotization increases the demand for all factor inputs and the following facts are true:*

1. *An increase in φ_H increases the demand for all factor inputs.*
2. *An increase in φ_1 increases the demand for low-skilled workers L_1 and has no effect on the demand for high-skilled workers L_2 or the demand for intermediate inputs M.*
3. *An increase in φ_2 increases the demand for high-skilled workers L_2 and has no effect on the demand for low-skilled workers L_1 or the demand for intermediate inputs M.*

Furthermore,

1. *The labor ratio L_2/L_1 is increasing with robotization, increasing in φ_2, constant in φ_H, and decreasing in φ_1.*
2. *The intermediate to labor ratio $M/(L_1+L_2)$ is constant in robotization, constant in φ_H, and decreasing in φ_o for all $o \in O$.*

5.2 Time series simulation and event studies

After calibrating parameters, both our model and the permanent adoption model can almost perfectly match the fraction of automating firms in the data (see Figure 6). In order to achieve the same outcome in both models, we reduce the price of robotization p_R in the permanent adoption model to account for the increased risk of the cost shock (since deautomation is not feasible). In addition, our reversible adoption model is able to replicate the respective proportion of robotized firms that stay automated, deautomate, or automate repeatedly. Remarkably, our simulations can also generate the derobotization patterns consistent with Fact 2, with half of all derobotization occurring one period after robotization. Since the action regions for robotization and derobotization are on opposite ends of the state space (cf. Figure 4) and the persistence of shocks on φ_H is relatively high, the main drive for this prompt abandonment of robots must arise from the realization of high robotization costs. Therefore, the revelation effect of derobotization is able to not only plausibly explain our observations from the data qualitatively, but also closely match those observations quantitatively.

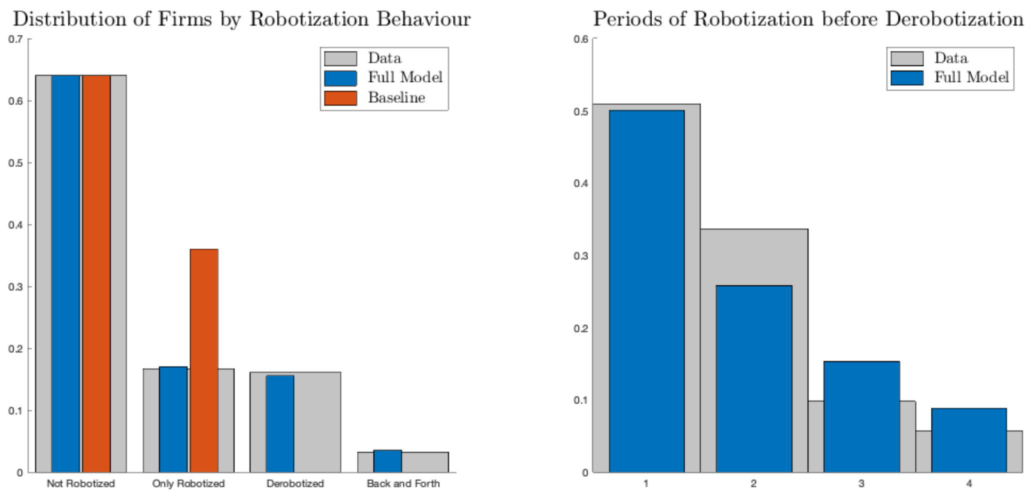

Figure 6: Robotization and derobotization behaviour of firms in simulated time series.

5.2.1 The robotization event

Having established the suitability of our model to account for the robotization and derobotization behavior of firms, we turn to analyzing the effects of a robotization event.

As evident from Figure 7, there are two crucial differences between the reversible and permanent adoption models:

Observation 3. *The demand for factor inputs steadily declines in the periods following robotization. If robot adoption is reversible, this decline is deeper.*

Even though the robotization event increases the demand for all inputs (Observation 2), it is preempted by periods of above-average productivity from robot adoption which regress toward periods of average productivity and, as a result, lower input demands. Furthermore, the deepening effect of reversibility arises because firms which face high cost re-

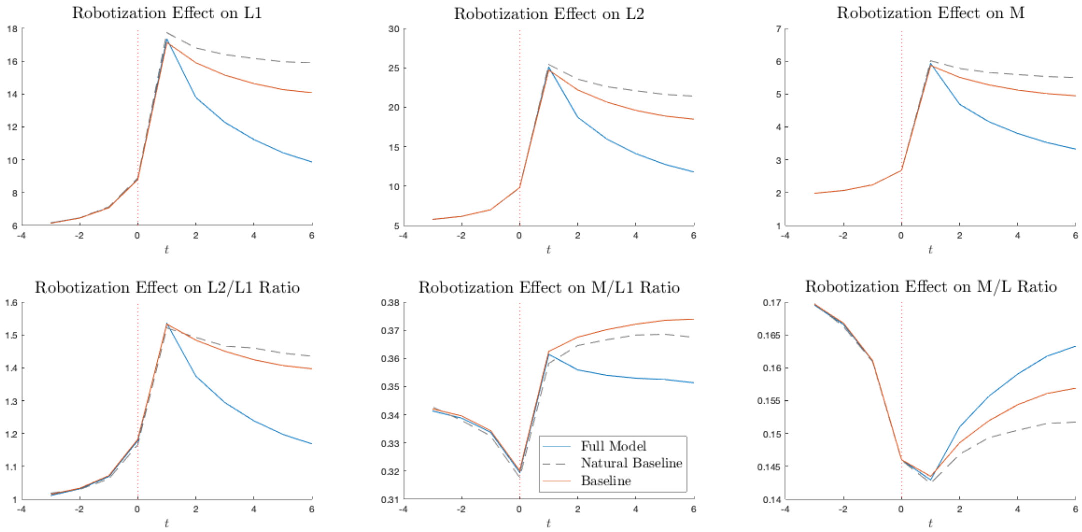

Figure 7: Robotization event in the simulated time series

Values represent the mean of those firms that robotized at least once. Firms that robotize either in the first or last periods are excluded to properly trace the behaviour before and after the automation event. The period the robotization happens are normalized to 0 for each firm.

alizations (or a sequence of low productivity shocks) would be better off derobotizing, which causes further decline in productivity, and then, by Observation 2 again, causes a decrease in demand.

Observation 4. *If robot adoption is reversible, only the more productive firms stay robotized.*

We will refer to firms which stay robotized under the reversible adoption model as *natural adopters*, as represented by the dashed gray line in Figure 7. The intuition of Observation 4 is straightforward: less productive firms are more likely to derobotize, leaving a positively biased (in demand) subsample of firms. Therefore, natural adopters have a higher input demand compared to the setting in which all firms, regardless of their productivity, are permanent adopters.

These differences have significant implications for relative factor demands. Although relative demand for high-skill labor increases following robotization, the decline in the labor ratio in the succeeding periods is sufficiently steep to return to prerobotization values (see Section 5.2.2 for a longer discussion).

Furthermore, the permanent adoption model overestimates the increase of the intermediate-low (M/L_1) ratio following robotization. Notice that in our model the drop in low-skilled labor (approximately 37.5%) over the five periods after robotization is less than the drop in intermediate inputs (approximately 42%). In the permanent adoption model, these differences are much smaller, with drops of approximately 12.5% and 12% for low-skilled labor and intermediate inputs respectively. A similar observation follows for the intermediate-labor ratio.

5.2.2 The derobotization event

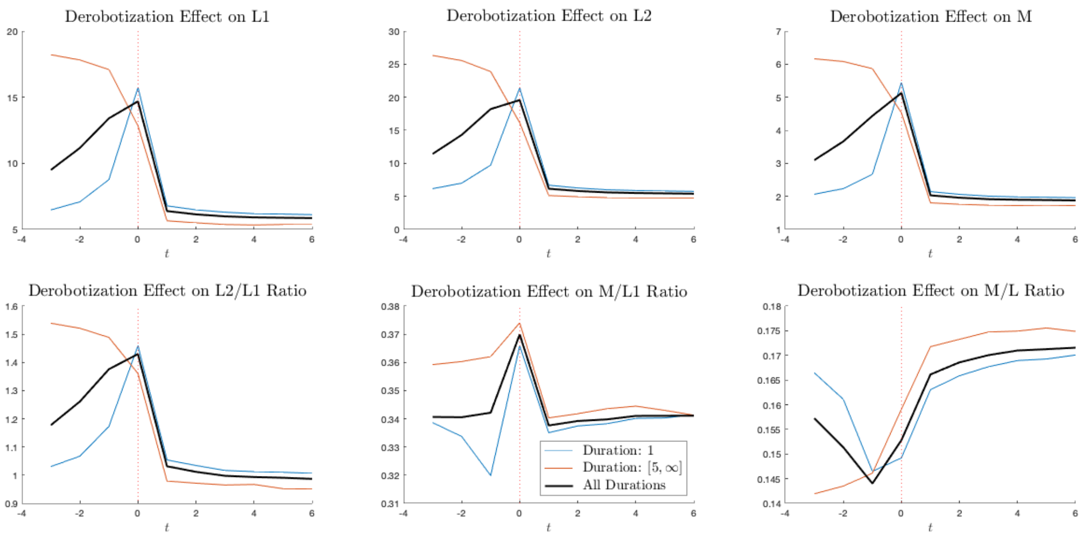

Figure 8: Derobotization event in the simulated time series

Values represent the mean of those firms that derobotized, differentiating across how long it took for them to do so. Black lines represent the mean of all the firms that derobotized at least once.

We extend our analysis to look at the firm's dynamic behavior around derobotization, summarized in Figure 8.

It should be noted that different firms in our model, hand-in-hand with firms in the data, derobotize at different periods (see Figure 6 again); firms which abandon promptly do so due to the revelation effect, which implores us to explore why firms derobotize in the second period onward.

Observation 5. *If firms take longer to derobotize, they do so following a sequence of low productivity shocks.*

This observation is clearly visible in the simulation but is masked by the behavioral distribution of derobotized firms. The firms which derobotize after the first period have a decline in demand which is indicative of low productivity. On average, the revelation effect overcrowds the reaction to the productivity effect; indeed, see that the mean firm follows the trends of prompt abandonment. It is also evident that the derobotization event in our model comes with a negative shock to input demand, confirming Fact 4. We conclude with a final observation:

Observation 6. *Prior to the derobotization event, firms which do not promptly abandon* *exhibit lower L_2/L_1 ratios. Furthermore, derobotization causes a drop in labor demand and (therefore) an increase in the relative productivity of intermediates.*

For firms which robotize for longer periods, the decline in low-skill productivity is weaker. In addition, our simulation also confirms Fact 4 as all firms which derobotize see significant increases in their intermediate-labor ratio (note that the productivity effect has crowded out the revelation effect past the first period). Therefore, the sequence of low productivity shocks which trigger derobotization increase the relative productivity of intermediate inputs.

6 Conclusion

Contributing to the growing literature of automation and robotization, our paper analyzes one aspect of adopting robots generally neglected, namely their abandonment.

As our first contribution, we empirically investigate this phenomenon using Spanish firm-level data and document a series of facts on derobotization. In contention with the literature, we find that derobotization is an inherent feature of robot adoption as more than two-fifths of firms abandon robots during their lifetime. Furthermore, the distribution of derobotization is unevenly distributed across firm size and time: when derobotization happens, it usually occurs within 8 years following adoption and affects small- to medium-sized firms. We also find that once a firm derobotizes, it faces a drop in labor demand, and as a result, an increasing in the intermediate-labor ratio.

We proposed a model of reversible robot adoption that was able to match the facts of the data and captured reasons why firms both adopt and abandon their robots. In addition to the productivity effect we documented, we observe in simulations a revelation effect brought on by *ex-ante* uncertainty regarding robotization. Our model is the first to document both effects of derobotization by simply allowing firms to abandon in an environment of stochastic productivity and learning-by-doing costs. Our simulations are able to match the data, especially in generating the correct proportion of derobotizing firms.

Finally, simulations of our model of robot adoption hint at an overestimation of demand effects from the current literature. In conjunction with the fact that firms derobotize in the data, we argued that derobotization is a non-negligible behavior of firms. It is necessary, at least to us, to take this phenomenon into account when analyzing automation and its effect on the labor market and policy.

Particularly regarding policy, the conversation on automation has often focused on the propensity (and sometimes inevitability) of workers' displacement, something which our paper takes no stance on but provides some context. Clearly, the setting in which robots outright replace workers is incomplete; the existence itself of firms derobotizing shows that the existence of technology does not necessitate its widespread adoption.[16] This insight is valid for any adoption of technology, as has been obvious in the literature of other sub-disciplines of economics, particularly in industrial organization and operations research. Even if the idea of learning (as it is implemented in our model) is novel in this particular literature, other papers have recognized that the matter is at least nuanced (see Acemoglu and Restrepo's (2018b) discussion on labor displacement).

It is apparent to us that any future research which seeks to model automation should also model deautomation. Crucially, capturing prompt abandonment is necessary, but models could incorporate other explanations. An improvement to our paper could include modeling the revelation effect to include uncertainty over productivity instead of costs. Moreover, further research is needed to analyze the full implications of derobotization, its determining factors, and its implications in a general equilibrium framework in a fashion similar to Humlum (2019). This will give greater insight on the effects of automation in the labor market, especially how wages evolve with robotization.

[16]This is a basic economic insight, since our story suggests the rudimentary idea that costs are also part of the equation.

References

[1] **Acemoglu, Daron, and David Autor.** 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." *Handbook of Labor Economics.*

[2] **Acemoglu, Daron, Claire LeLarge, and Pascual Restrepo.** 2020. "Competing with Robots: Firm-level Evidence from France." National Bureau of Economic Research, Working Paper 26738.

[3] **Acemoglu, Daron, and Pascual Restrepo.** 2018a. "Modeling Automation." National Bureau of Economic Research, Working Paper 24321.

[4] **Acemoglu, Daron, and Pascual Restrepo.** 2018b. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review.*

[5] **Acemoglu, Daron, and Pascual Restrepo.** 2017. "Robots and Jobs: Evidence from US Labor Markets." National Bureau of Economic Research, Working Paper 23285.

[6] **Autor, David.** 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives.*

[7] **Brynjolfsson, Erik, and Andrew McAfee.** 2014. "The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies." *W. W. Norton Company.*

[8] **Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock.** 2018. "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings 2018, Economic Consequences of Artificial Intelligence and Robotics.*

[9] **Cox, D.R.** 1972. "Regression Models and Life-Tables." *Journal of the Royal Statistical Society.*

[10] **Frey, Carl Benedikt, and Michael A. Osborne.** 2017. "The future of employment: How susceptible are jobs to computerisation?" *Technological Forecasting and Social Change.*

[11] **Humlum, Anders.** 2019. "Robot Adoption and Labor Market Dynamics." https://static1.squarespace.com/static/5d35e72fcff15f0001b48fc2/t/5dcf78576d59eb44eac86f63/1573877848584/humlumJMP.pdf. Accessed July 22, 2020.

[12] **Koch, Michael, Ilya Manuylov, and Marcel Smolka.** 2019. "Robots and firms." CESifo Working Paper 7608.

[13] **Kopecky, Karen, and Richard Suen.** 2010. "Finite state Markov-chain approximations to highly persistent processes." *Review of Economic Dynamics.*

[14] **Ocampo, Sergio.** 2018. "A Task-Based Theory of Occupations with Multidimensional Heterogeneity." https://ocamp020.github.io/JMP_Draft.png. Accessed July 22, 2020.

[15] **Rouwenhorst, Geert K.** 1995. "Asset pricing implications of equilibrium business cycle models." *Frontiers of Business Cycle Research.*

Appendix

A Firm-level facts on derobotization

A.1 Firm data and summary statistics

The ESEE asks firms to

> State whether the production process uses any of the following systems:
>
> 1. Computer-digital machine tools;
> 2. Robotics;
> 3. Computer-assisted design;
> 4. Combination [sic] of some of the above systems through a central computer (CAM, flexible manufacturing systems, etc.);
> 5. Local Area Network (LAN) in manufacturing activity.

The firm's positive (or negative) response to "2. Robotics" determines whether it is identified as using robots.

Before any further adjustments, the dataset contains 15,929 observations across 5,588 firms. We first exclude observations in which robot use was not provided, which reduces the number of observations to 15,730 across 5,511 firms. Then, we exclude firms with less than two periods of data or with a gap on robotization, whether this is from a lack of reporting or from the first exclusion we made. As our analysis focuses on the change in robotization and the time between such changes, not knowing if firms switched during a particular period makes it difficult to make conclusions about periods before and after. This reduces the number of observations to 14,119 across 3,987 firms. Further exclusions were made in specific analyses, which we point out where applicable.

Table 5: Starting robotization state and state changes.

State Changes	Unrobotized Start	Unrobotized left-censored	Robotized Start	Robotized left-censored
0	1,343 (33.5%)	977 (24.9%)	388 (9.7%)	134 (3.3%)
1	231 (5.8%)	218 (5.4%)	180 (4.5%)	89 (2.2%)
2	77 (1.9%)	118 (2.9%)	52 (1.3%)	66 (1.6%)
3+	15 (0.4%)	57 (1.4%)	9 (0.2%)	33 (0.8%)

We partition the firms between those which were robotized in the first year they appear in the data, and those that were not. We count a state change when firms go from robotizing to derobotizing, or from derobotizing to robotizing. Some of the firms' state changes are left-censored by the survey starting in 1991. We further partition the panel of firms between those which are left-censored and those who are not. A breakdown of state changes is given by Table 5.

We also generate descriptive statistics in Table 6, which summarize the distribution of capital, labor, and capital-to-labor ratios between firms of different robotization behaviors.

Table 6: Descriptive Statistics

Variable	All Firms	Not Robotized	Only Robotized	Derobotized	Back & Forth
Hours Worked (1,000s, Log)	4.78 (1.50)	4.12 (1.22)	5.68 (1.40)	5.08 (1.41)	5.77 (1.45)
Capital (Log)	14.77 (2.36)	13.66 (2.05)	16.14 (1.82)	15.40 (2.16)	16.38 (2.06)
Capital per 1,000 Hours (Log)	9.99 (1.33)	9.54 (1.35)	10.46 (1.06)	10.32 (1.25)	10.61 (1.10)

This table shows the average values and standard deviations (in parentheses) of firm-year observations for the variables listed.

A.2 Kaplan-Meier estimation

The exact values for the KM estimation of cumulative derobotization rates can be found in Table 7.

Table 7: Kaplan-Meier Estimates - Cumulative Derobotization Rate

Robotization	Non-Censored Data Only	Minimum Overall Derobotization	Maximum Period Derobotization	Midpoint of Interval
1	29.2% (45.1%)	23.2% (53.8%)	31.4% (49.1%)	27.3 (51.0%)
2	23.3% (36.0%)	15.6% (36.2%)	20.4% (31.9%)	18.0 (33.6%)
3	7.0% (10.8%)	3.1% (7.2%)	7.4% (11.6%)	5.2 (9.7%)
4	5.3% (8.2%)	1.2% (2.8%)	4.8% (8.5%)	3.0 (5.6%)

The first term is the percentage of robotized firms estimated to derobotize in that period, while the term in brackets is the percentage of firm who derobotize that are estimated to derobotize in that period.

Table 8 provides the exact KM estimates of the cumulative robotization rate, while Figure 9 plots the results graphically.

Table 8: Kaplan-Meier Estimates - Cumulative Robotization Rate

Period	Non-Censored Data Only	Minimum Overall Robotization	Maximum Overall Robotization	Midpoint of Interval
1	14.6% (27.2%)	8.8% (45.0%)	12.3% (24.6%)	10.6 (30.3%)
2	14.2% (26.4%)	6.5% (32.9%)	12.8% (25.5%)	9.6 (27.6%)
3	10.7% (20.0%)	2.9% (14.6%)	9.2% (18.4%)	6.0 (17.4%)
4	11.1% (8.2%)	1.3% (2.8%)	9.4% (18.8%)	5.4 (15.4%)
5	3.1% (5.8%)	0.2% (0.9%)	6.3% (12.7%)	3.3 (9.3%)

The first term is the percentage of unrobotized firms estimated to robotize in that period, while the term in brackets is the percentage of firm who robotize that are estimated to robotize in that period.

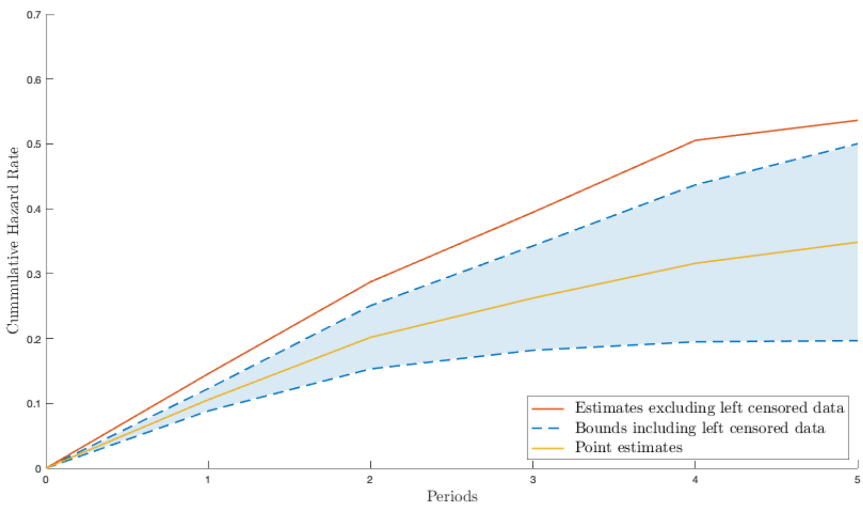

Figure 9: Kaplan-Meier estimates of robotization rates

A.3 Robustness test: duration dependence

As pointed out in Section 3, we restricted our sample to the firms for which the start of duration is known. However, this approach decreases the number of observations and thus does not allow us to introduce a richer set of controls. As a robustness check, we extend the sample to firms for which the start date of robotization may not be observed.

We assume that if a firm enters the survey being robotized, the date of entry and robotization coincide. We estimate equation (1) and also add controls for industries.

The results of our estimation are shown in Table 9. Compared to our previous estimates, the coefficient on firm size increases significantly in specification 1 to 3. Once a control for industry is included, the coefficient on size decreases in magnitude but remains higher than the baseline estimates. In particular, the probability that larger firms derobotize is much smaller than in our estimates from Table 3. A possible explanation behind this may be the longer lifetime of larger firms and the selection of smaller firms in our initial sample. At the same time, the coefficients on previous robotization remain indistinguishable from zero in each specification.

Table 9: Robustness test: Inclusion of firms robotized at t_0

	(1)	(2)	(3)	(4)	(5)	(6)
Size						
From 100 to 500	-0.53***	-0.53***	-0.41***	-0.35***		
	(0.10)	(0.10)	(0.10)	(0.10)		
From 500	-0.92***	-0.92***	-0.81***	-0.72***		
	(0.15)	(0.15)	(0.16)	(0.16)		
prevrob	-0.06	0.06	0.07	0.04	-0.03	
		(0.15)	(0.15)	(0.15)	(0.15)	(0.15)
$\log(K/L)$			-0.23***	-0.30***		
			(0.05)	(0.05)		
$\log(capital)$					-0.22***	
					(0.02)	
$\log(hours)$						-0.23***
						(0.04)
Industry FE				yes	yes	yes
Observations	1792	1792	1769	1769	1769	1769

The sample includes duration of robot usage for firms that: (a) were robotized at least once during 1991-2014; (b) did not abandon robots more than 2 times. Baseline category of size is "Less than 100 employees". Capital is in prices of 1991. Standard errors in parenthesis. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A.4 Robustness test: disaggregated data and firm size

So far, we studied the effect of firm size on derobotization using the aggregate data on duration and the firms' characteristics over the robotization period. Since the firm's size may be subject to productivity shocks, there are several caveats to this approach: first, the median size of firms may not align with the size of firms with longer periods before derobotization. Second, firms facing a higher volatility of shocks are more likely to both derobotize and change their size.

To overcome this problem, we restrict our sample to firms which robotize only once and measure firm size by its lagged value instead. Then, we estimate a linear probability model on pooled data, explicitly

$$\text{derobotization}_{it} = \beta_1 \text{size}_{it-1} + \beta_2 \log(K/L)_{it-1} + \beta_3' \times \text{industry}_{it} + \tau_t + \varepsilon_{it}. \qquad (14)$$

The result of our estimation is presented in Table 10. In each specification, the coefficient of size on the probability of derobotization is negative and statistically significant at a 5 percent level. In comparison to smaller firms, the probability of derobotization for medium- and big-sized firms is maller by roughly 8 and 16%, respectively. We conclude that this finding is robust across different estimation strategies and samples.

Table 10: Robustness test: disaggregated data

	(1)	(2)	(3)	(4)
Size_{t-1}				
From 100 to 500	-0.09***	-0.08***	-0.08***	-0.06***
	(0.02)	(0.02)	(0.02)	(0.02)
From 500	-0.17***	-0.16***	-0.16***	-0.14***
	(0.02)	(0.02)	(0.02)	(0.03)
$\log(K/L)_{t-1}$		-0.02*	-0.02	-0.02**
		(0.01)	(0.01)	(0.01)
Time FE			yes	yes
Industry FE				yes
Observations	1560	1538	1538	1529

The overall number of firms is 646. The sample includes firms that: (a) are observed at least for 3 periods; (b) were not robotized in 1991; (c) adopted robot over 1991-2014; (d) did not switch back and forth. Baseline category of size is "less than 100 employees". Standard errors are clustered on firm level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A.5 Additional figures and tables

Table 11: Summary statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
size	1.85	0.72	1	3	2211
$\log(K/L)$	10.62	1.05	4.16	13.87	2185
$\log(hours)$	5.64	1.46	1.39	10.18	2188
$\log(capital)$	16.28	1.97	8.39	21.92	2208
derobotization	0.22	0.41	0	1	2503
industry	11.36	5.24	1	20	2211

The sample provides information about variables used in Section 3 and includes firms that: (a) are observed at least for 3 periods; (b) were not robotized in 1991; (c) adopted robot over 1991-2014; (d) did not switch back and forth.

Figure 10: Baseline Hazard: Cox proportional hazards regression

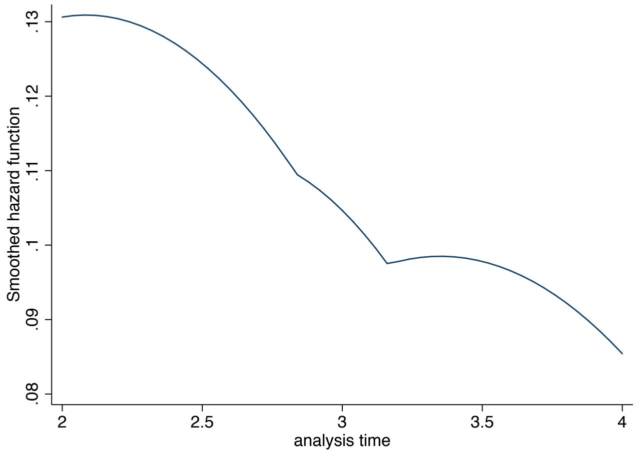

The figure shows the smoothed estimate of the baseline hazard function from specification 3 of the Cox proportional hazard regression.

B Micro-foundations of production function

Following the task-based production function from Acemoglu and Restrepo (2018a), we derive the production function used in Section 4 as in Humlum (2019).

Consider a firm operating with a task-based production function:

$$Y := \left(\int_0^1 y(i)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}, \tag{15}$$

where $y(i)$ denotes the production of a task $i \in [0,1]$. To complete the task i, a firm can employ a factor X_o from occupation $o \in O \subset \mathbb{N}$, with productivity $z_o(i,R)$ which depends on the robotization state $R \in \{0,1\}$.[17] The production function of completing task i with factor X_o is given by

$$y(i) := z_o(i,R) X_o(i). \tag{16}$$

With a slight abuse of notation, we will suppress the robotization argument of task productivity so that $z_o(i,R) \equiv z_o(i)$. Conditional on the robotization state, the optimal task-assignment problem is given by

$$\max_{\{X_o(i)\}_{o=1}^{|O|}} \left(\int_0^1 y(i)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}} - \sum_{o=1}^{|O|} \left(\int_0^1 p_o X_o(i) di \right), \tag{17}$$

where p_o is the price of factors from occupation o. In other words, the firm chooses: (1) the assignment of tasks $A = \{A_o\}_{o \in O}$ and (2) the amount of factors X_o used for production of each task. Once the assignment of tasks is chosen, the first-order condition for a factor X_o and a task $i \in A_o$ is

$$z_o(i)^{\sigma-1} Y = X_o(i) p_o^{\sigma}. \tag{18}$$

Combining first-order conditions for tasks $i,k \in A_o$ and integrating by k over A_o, we obtain that

$$\frac{\int_{A_o} z_o(k)^{\sigma-1} dk}{z_o(i)^{\sigma-1}} = \frac{\int_{A_o} X_o(k) dk}{X_o(i)}. \tag{19}$$

Finally, the first-order condition simplifies to

$$X_o(i) = \mathbf{X}_o \frac{z_o(i)^{\sigma-1}}{\int_{A_o} z_o(k)^{\sigma-1} dk}, \tag{20}$$

where $\mathbf{X}_o \equiv \int_{A_o} X_o(k) dk$ is the total amount of factors from occupation o used in production. Plugging in equation (20) into the definition of the task-based production function,

[17] By assumption, the set of occupations O is finite.

we get:

$$Y \equiv \left\{ \int_0^1 y(i)^{\frac{\sigma-1}{\sigma}} di \right\}^{\frac{\sigma}{\sigma-1}} = \left\{ \sum_{o=1}^{O} \left(z_o \mathbf{X}_o \right)^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{\sigma}{\sigma-1}} \tag{21}$$

$$z_o := \left(\int_{A_o} z_o(i)^{\sigma} di \right) \left(\int_{A_o} z_o(i)^{\sigma-1} di \right)^{-1} \tag{22}$$

The decision to robotize changes the production function along two dimensions. First, it changes the productivity by an absolute magnitude, leading to direct gains in efficiency. Second, and more importantly, it changes the relative productivity of different occupations, which leads to assignments of different tasks across occupations. For our simulations, we rewrite the aggregate robotization-dependent productivity as Humlum does in his paper:

$$z_o = \exp(\varphi_o + R\gamma_o) \tag{23}$$

$$\varphi_o = \log \frac{\int_{A_o} z_o(i,0)^{\sigma} di}{\int_{A_o} z_o(i,0)^{\sigma-1} di} \tag{24}$$

$$\gamma_o = \log \frac{\int_{A_o} z_o(i,1)^{\sigma} di}{\int_{A_o} z_o(i,1)^{\sigma-1} di} - \log \frac{\int_{A_o} z_o(i,0)^{\sigma} di}{\int_{A_o} z_o(i,0)^{\sigma-1} di} \tag{25}$$

C Simulation

C.1 Parametrization and methodology

Table 13 shows the parametrization used for the different models simulated.

Table 13: Parametrization

Parameter	Full Model	Permanent Adoption	Certain Costs
β	0.8	0.8	0.8
σ	1.5	1.5	1.5
γ_H	1.74	1.74	1.74
γ_1	-0.1	-0.1	-0.1
γ_2	0.2	0.2	0.2
w_M	15	15	15
w_1	10	10	10
w_2	20	20	20
C_R	500	500	500
p_R	1260	1010	1160
δ_P	0.03	1×10^{20}	0.03
δ_S	0.17	0.17	0.17
ε	1.5	1.5	1.5
Y_M	10	10	10
P_M	10	10	10
φ_H	[0,0.7,1.3,1.8,2]	[0,0.7,1.3,1.8,2]	[0,0.7,1.3,1.8,2]
φ_1	[0,0.25,0.5,0.75,1]	[0,0.25,0.5,0.75,1]	[0,0.25,0.5,0.75,1]
φ_2	[1,1.25,1.5,1.75,2]	[1,1.25,1.5,1.75,2]	[1,1.25,1.5,1.75,2]
ε_R	[-220,-110,0,110,220]	[-220,-110,0,110,220]	[0,0,0,0,0]
ρ_H	2/3	2/3	2/3
ρ_1	0.3	0.3	0.3
ρ_2	0.5	0.5	0.5

C.1.1 Value function iterations

Our simulations are programmed using MATLAB. The value function iteration relies mostly on the usage of multidimensional arrays, with the dimensions corresponding to the different vectors of states. In a first step, optimal profits and inputs are calculated by exploiting the properties of the CES production function, allowing for independent optimization of each input.

In order to account for the different properties of first-time and repeated-time robotization, the arrays are then extend before specifying initial guesses and iterating the value functions. Due to the different conditional expectations in different robotization states, the expectations are computed separately for a known and unknown ε_R.[18] As soon as some specified precision is reached, policy and value functions are reported.

C.1.2 Time series simulation and event studies.

We start by specifying the initial state profile according to the unconditional distribution of the discretized AR processes, and proceed to generate a time series of state profiles for each firm by randomly generating new exogenous states each period according to the transition probabilities.

Using the policy function arrays from the value function iteration, we obtain a time series of robotization behavior and input demands. On this simulated dataset, we then perform the calculation of summary statistics and generate event studies. The latter is done by identifying each robotization and derobotization event, excluding those which are either too early or late and so do not contain enough data regarding the pre- and post- behavior, normalizing the event period to zero, and plotting the averages of the corresponding samples.

C.2 Counterfactual comparisons

As described in Section 5, derobotization is jointly driven by the productivity effect and the revelation effect. The latter in particular is what enables the model to realistically capture the behavior observed in the data. To illustrate this, Figure 11 shows the differ-

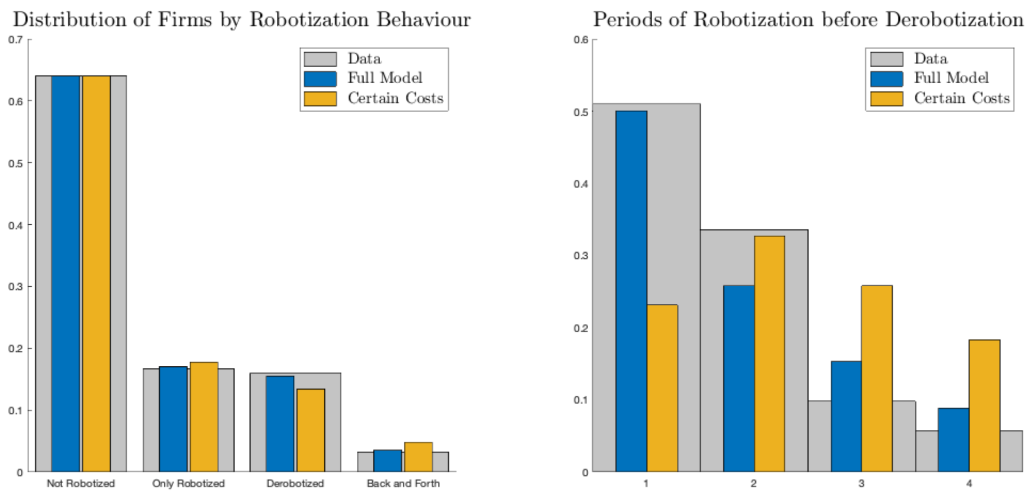

Figure 11: Summary statistics for the full model and the certain costs model

ences when allowing for an initial uncertainty in robotization costs. While being able

[18] All expectations are calculated using matrix multiplication adapted to higher-dimensional arrays.

to account for the overall phenomenon of derobotization fairly well (as apparent in the left histogram), the shortcomings become apparent when considering the behavior among derobotizers. Having only decreasing productivity as a catalyst for deautomation, it relies heavily on its persistence. However, as derobotization occurs at the opposite part of the state space than adoption, it is highly unlikely that a full fast decrease of productivity (from above to below average) takes place in the first period. Thus the modeling of a predominance of first-period reversion requires some form of uncertainty in the success of automation, captured by the production costs using robots.

On another note, we also consider the effect of setting $\delta_S = 0$. While the timing of derobotization remains mostly unchanged (see Figure 12), setup costs are necessary to generate firms which switch back and forth, as in the data. As stated in Section 4, the difference in price between p_R and $(1-\delta_S)p_R$ is the main incentive of rerobotization.

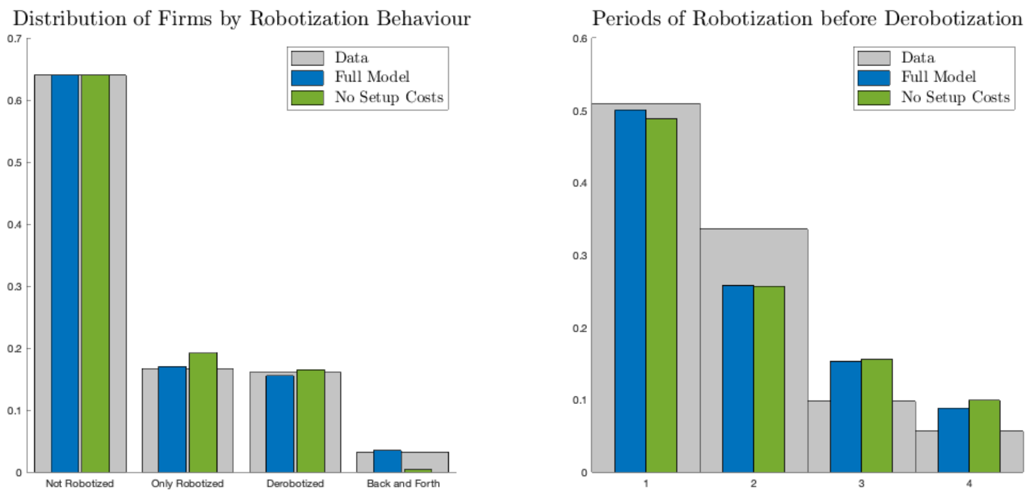

Figure 12: Summary statistics for the full model and the model without setup costs.

C.3 Policy function heatmaps

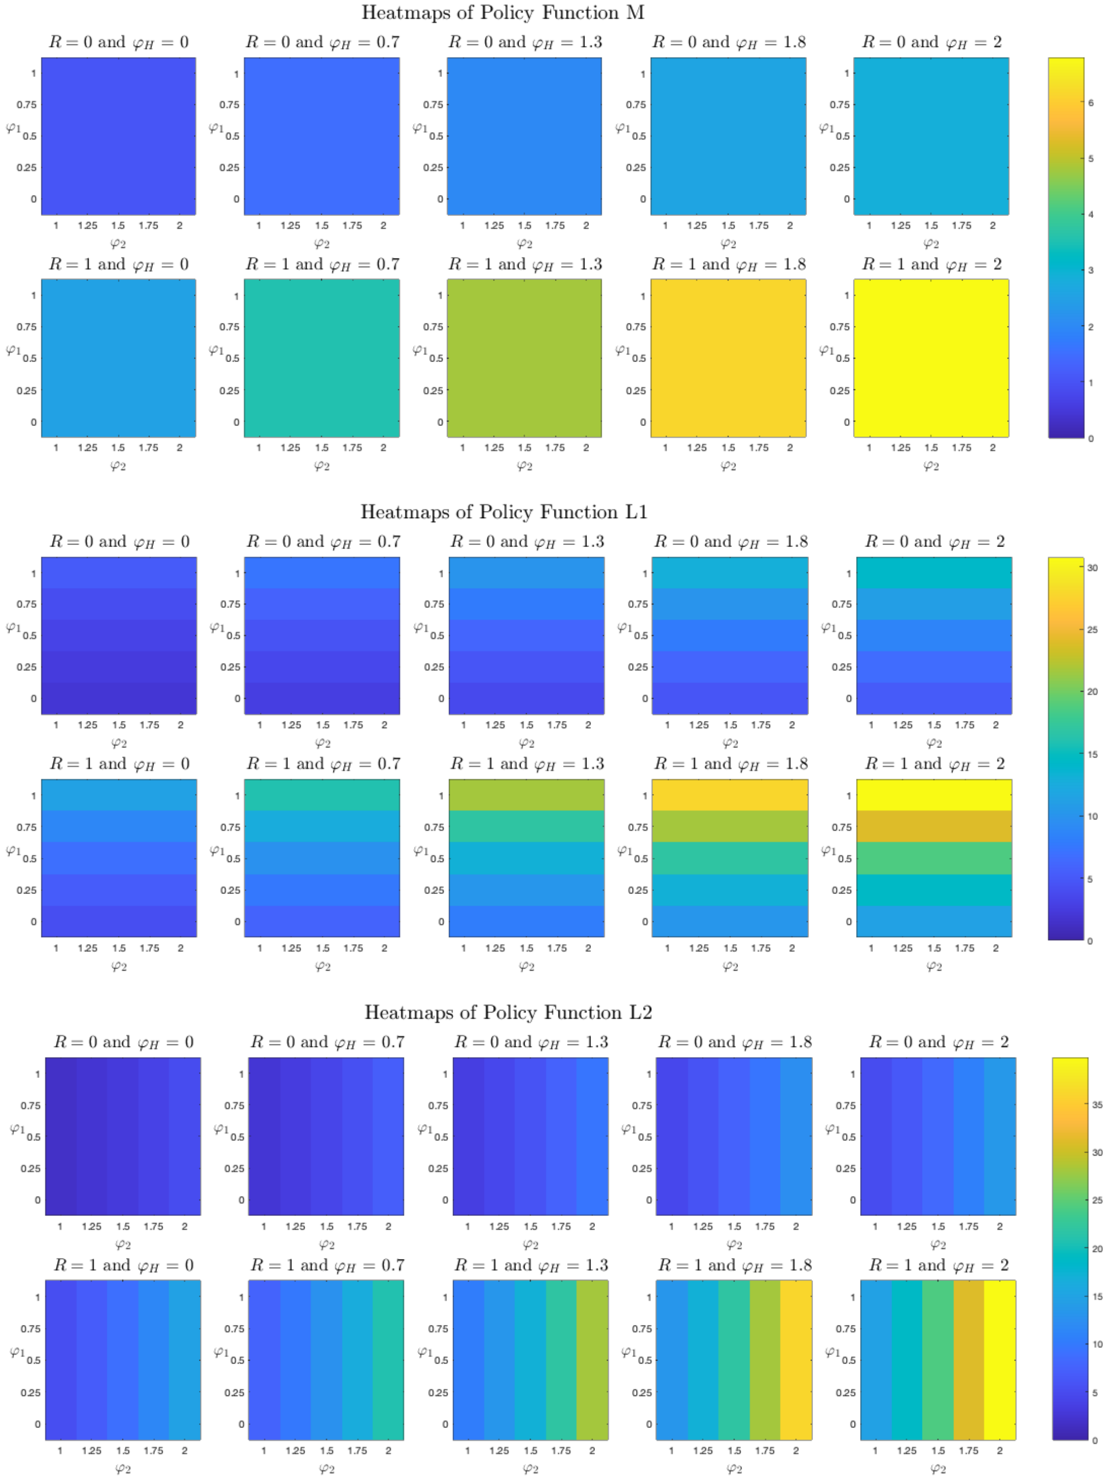

Figure 13: Heatmaps for the policy function M, L_1 and L_2

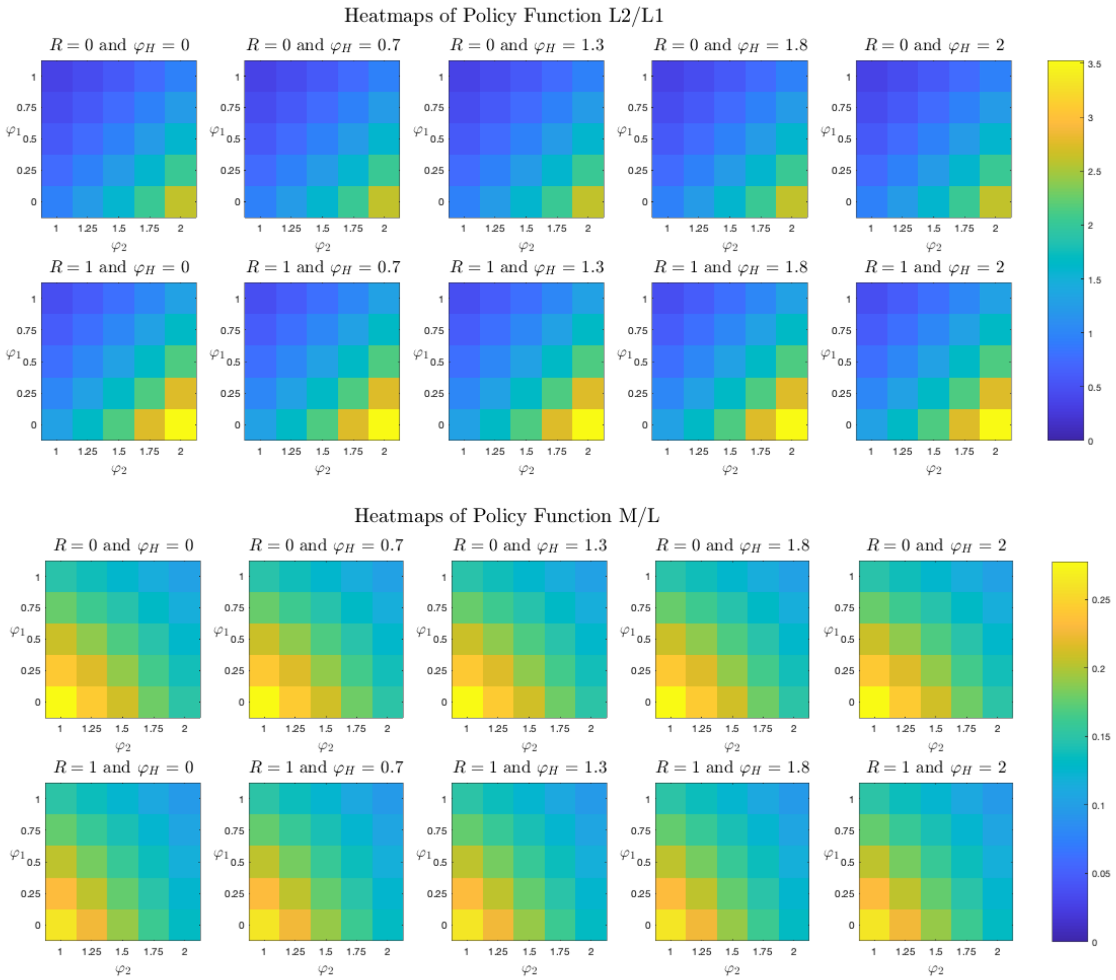

Figure 14: Heatmaps for the policy function L_2/L_1 and $M/(L_1 + L_2)$