Lost Knowledge and New Products: Event-Study Evidence from Worker Deaths *

Anders Akerman and Kerstin Holzheu‡

April 24, 2026

Abstract

Do worker-embodied knowledge stocks shape firms' ability to expand product scope? We exploit unexpected worker deaths as plausibly exogenous shocks to firm knowledge, using Swedish linked employer–employee data to recover workers' prior product experience from their former employers. Event-study estimates show that each year of prior product-specific experience lost to a worker death lowers the on-impact probability of entering the corresponding product by about 0.13 percentage points, implying a roughly 0.7 percentage point decline — about 27 percent of the baseline rate — for a worker with average prior experience. The effect peaks in the first post-event year before attenuating. Static estimates confirm that cumulative lost experience predicts lower subsequent entry and substantially lower product-level sales. A descriptive decomposition indicates that the experience-scaled component accounts for about 92 percent of the sales decline and about 58 percent of the entry decline, with the latter interpreted more cautiously because entry is rare. Effects are roughly twice as large for university-educated workers, markedly attenuated for core products, and suggestive of greater exposure among high-skill firms. A back-of-envelope calculation implies an expected loss of roughly $70,000 per exposed worker for a product line with $10 million in net entry value, providing a benchmark for codification, retention, and redundancy investments.

JEL codes: J24, J41, J63, J64
Keywords: Productivity, Worker Mobility, Innovation

*We thank seminar participants and discussants at IE University, IFS, NOITS, Lisbon Macro Workshop, SED 2024, Stockholm University, Yale University, as well as Johannes Boehm, Nicolas Coeurdacier and Isabelle Mejean for helpful discussions. We thank Statistics Sweden for data access. Akerman gratefully acknowledges support from the Research Council of Norway (grant no. 326000). Holzheu thanks the Banque de France for financial support.

‡Akerman: Stockholm University (anders.akerman@su.se), Holzheu: Sciences Po (CNRS), IZA and CESifo (kerstin.holzheu@sciencespo.fr)

1 Introduction

"In Africa, when an old man dies, a library burns to the ground."
— Attributed to Amadou Hampâté Bâ (Waberi, 2018)

Firms expand product scope by combining physical capital with knowledge that is often embodied in their workers. A central question in strategic management is whether this knowledge constitutes a durable source of competitive advantage or a fragile dependency that exposes the firm to human capital risk (Barney, 1991; Teece et al., 1997). When knowledge walks out the door, do firms become less able to launch new product lines? This paper provides causal evidence on this question by studying how unexpected worker deaths—plausibly exogenous shocks to the firm's stock of embodied knowledge—shape firms' product choices. Using linked employer–employee data with detailed career histories, we measure product-specific knowledge accumulated at prior employers and track whether firms subsequently enter related product lines. Our estimates show that the loss of worker-embodied knowledge has economically meaningful effects on horizontal expansion. On impact, each year of prior product-specific experience lost to an unexpected worker death lowers the probability of entering the corresponding product by about 0.13 percentage points; evaluated at the mean level of prior experience among exposed deceased workers, this implies a roughly 0.7 percentage point decline at impact, with the effect becoming larger in the first post-event year before attenuating thereafter. For a product line with an expected net entry value of $10 million, this implies an expected loss of roughly $70,000 per exposed worker, a benchmark against which retention, cross-training, and knowledge codification investments can be evaluated on expected-value grounds.

Well-known cases illustrate how new product lines can be sparked by specific employees and the know-how they carry. Amazon's cloud business is often traced to an internal group led by Andy Jassy that productized infrastructure originally built to run Amazon's retail operations into AWS.[1] 3M's Post-it Notes are commonly attributed to two employees—a scientist (Spencer Silver) who discovered a low-tack adhesive and a colleague (Art Fry) who identified its application as a repositionable note—illustrating how worker-embodied expertise can seed adjacent product entry.[2] Worker mobility can also be a direct channel for related diversification *across* firms. Zoom's founder Eric Yuan previously helped build WebEx and later left to found a competing video-conferencing product, explicitly drawing on that experience.[3] In the electric-vehicle sector, Lucid Motors' CEO Peter Rawlinson—previously a Tesla engineer—has emphasized how engineering experience from earlier EV

[1] TechCrunch (2016), "Andy Jassy's brief history of the genesis of AWS."
[2] People (2025), "Post-It Notes Were Introduced to the World 45 Years Ago."
[3] Reuters (2018), "Video conferencing company Zoom readies IPO—sources."

programs informed Lucid's approach to its first vehicle.[4] Yet due to selection—workers join firms poised to expand, and firms recruit workers with relevant expertise—causal claims are difficult to establish from observational data on hires or voluntary mobility.

This paper makes two contributions to our understanding of how worker-embodied knowledge shapes firm boundaries. First, we provide causal evidence that such knowledge is a binding constraint on product scope. By exploiting the timing of unexpected deaths and leveraging within-firm, cross-product variation in exposure, we isolate the knowledge channel from confounding labor supply effects and show that the penalty is targeted to products in which the deceased worker had accumulated experience, with economically meaningful magnitudes on both entry and sales. Second, we develop and implement a descriptive decomposition that separates the total knowledge loss into transferable and match-specific components. This decomposition is most informative on the sales margin, where the component of the loss that scales with prior product-specific experience accounts for about 92 percent of the total decline. On the entry margin, the analogous ratio is about 58 percent, though we interpret it more cautiously because product entry is rare and decomposition ratios are sensitive to small changes in the underlying coefficients. We interpret this pattern as evidence that a substantial share of the lost capability is portable across employers rather than purely match-specific.

To study this question, the fundamental empirical challenge is that separations are endogenous to both worker and firm conditions: workers may leave declining firms, and firms may allow departures when expansion prospects are poor. We address this challenge by exploiting sudden worker deaths—identified from administrative cause-of-death records—as plausibly exogenous shocks to the firm's stock of embodied knowledge. Importantly, we recover each worker's product-specific experience from the production portfolios of their *prior* employers, ensuring that the relevant knowledge was accumulated before the current employment relationship and is therefore predetermined with respect to the firm's subsequent product decisions. Our event-study specification is estimated at the firm–product–year level with firm–product and year fixed effects. Identification thus comes from comparing product entries within the same firm across products that do and do not match the deceased worker's experience, before and after the death event.

We document four main findings. First, pre-event coefficients are small and statistically indistinguishable from zero, providing support for the parallel-trends assumption underlying our research design. Second, the probability of entering a new product line falls on impact and reaches its largest decline in the first post-event year before attenuating thereafter.

[4] ABC News (2023), "Lucid CEO has a plan to end Americans' range anxiety."

Third, static estimates show that cumulative lost product-specific experience predicts lower subsequent product entry and substantially lower product-level sales, with the strongest and most robust evidence appearing on the sales margin. Fourth, heterogeneity analyses show that the effects are larger for university-educated workers, substantially attenuated for core products, and suggestive of greater exposure among high-skill firms.

We interpret these patterns as evidence that firms "learn by hiring"—and conversely lose the ability to implement adjacent product expansions when that embodied knowledge is unexpectedly removed. Because our identification leverages deaths rather than hires, the estimates are not confounded by selective recruiting or sorting on unobservables that typically complicate studies of worker inflows. Product entry is a direct, economically salient outcome that connects worker knowledge to firm growth options, complementing more aggregate productivity measures used in prior work.

Our paper relates to several literatures spanning economics and strategic management. A first stream studies the performance effects of losing key individuals. Research on leader or entrepreneur deaths documents sizable effects on firm survival and profitability but does not examine product choice (Becker and Hvide, 2021; Sauvagnat and Schivardi, 2023; Jones and Olken, 2005). Studies of inventor deaths show spillovers on collaborators' innovation output (Jaravel et al., 2018; Azoulay et al., 2010), while Hoey et al. (2023) examine the impact of temporary incapacitations on team productivity. Closer to our setting, Bertheau et al. (2022) quantify the revenue costs to firms from worker deaths, and Jäger et al. (forthcoming) use wage responses to infer complementarity and substitutability among coworkers. Methodologically, we build on Andersen and Nielsen (2011), who pioneered the use of cause-of-death records to identify sudden, unanticipated mortality events for causal inference, applying the design in a household-finance setting; we adapt their classification to the firm–product context.

A second stream examines how worker mobility transmits knowledge across firms. Patault and Lenoir (2024) show that sales managers carry customer capital to new employers, generating revenue losses for origin firms and gains for destinations. Stoyanov and Zubanov (2012) find that hiring workers from more productive firms raises the productivity of destination firms, consistent with knowledge spillovers through labor mobility. Neffke and Henning (2013) use Swedish administrative data on labor flows between firms to construct a measure of skill relatedness and show that firms preferentially diversify into related products; we use a closely related empirical setup to document the converse — that the loss of workers carrying product-specific skills constrains diversification. A key limitation of studying inflows is that workers sort into firms poised to diversify, and firms poised to diversify recruit workers with matching expertise, so the observed correlation between skill inflows and product entry cannot be cleanly interpreted as a causal knowledge channel. Our reverse-direction design, exploiting involuntary outflows through unexpected deaths, bypasses this sorting problem and thus provides a causal counterpart to the skill-relatedness mechanism documented by Neffke and Henning (2013).

In the innovation domain, patent citation patterns suggest that knowledge flows track worker movements (Almeida and Kogut, 1999; Singh and Agrawal, 2011). Closer to our setting on the production margin, the industry-evolution literature has long emphasized that worker mobility from incumbents shapes new entry: Klepper and Sleeper (2005) and Klepper (2007) document that spinoffs from established firms inherit knowledge from their parents and account for a disproportionate share of entry in the laser and automobile industries, while Agarwal et al. (2004) show analogous patterns in the disk-drive industry. Neffke et al. (2011) extend the logic to regional diversification, showing that regions enter products that are skill-related to their existing portfolio, with worker flows operating as the underlying transmission mechanism. Our contribution complements this literature by exploiting an exogenous source of knowledge *loss* rather than transfer, which avoids the selection concerns inherent in studying voluntary spinoff formation, and by focusing on the *production* margin—product scope rather than patents or sales relationships.

A third stream in strategic management theorizes the conditions under which worker-embodied knowledge constitutes a source of competitive advantage. The resource-based view emphasizes that valuable, rare, and inimitable resources underpin sustained performance (Barney, 1991), but the mobility of human capital complicates this logic: if knowledge walks out the door, advantage may be fleeting. Dynamic capabilities theory distinguishes between operational routines embedded in the organization and higher-order capabilities that reside in individuals who sense and seize new opportunities (Teece et al., 1997). A more recent strategic-management literature has confronted directly the tension between human capital's mobility and its role as a source of advantage. Coff (1997) formalized the central dilemma; Argote and Ingram (2000) argued that knowledge transfer through people is the primary mechanism of organizational learning; Hatch and Dyer (2004) showed empirically that firm-specific human capital investments yield sustainable performance differences in semiconductor manufacturing; and Campbell et al. (2012) synthesize these threads into a framework for when human capital generates rents despite mobility. Our experience-graded versus experience-invariant decomposition provides a quantitative complement to this conceptual literature: the finding that the experience-scaled component accounts for about 92 percent of the sales loss suggests that product expansion capabilities are more fragile—and more dependent on retention—than theories emphasizing firm-embedded routines would predict. Our heterogeneity results showing attenuated effects for core products provide direct evidence that firms can buffer against human capital risk by investing in codification and organizational redundancy. In particular, we find that while high-skill workers generate larger absolute losses, a smaller share of their value is portable across firms, suggesting that retention and codification serve as complements rather than substitutes for different parts of the workforce.

The findings have implications at both micro and macro levels. At the micro level, they highlight that product expansion depends not only on financial capacity and demand conditions but also on the retention of specific, worker-embodied know-how. Managers can mitigate vulnerability through three channels: codifying tacit knowledge into organizational routines, distributing expertise across multiple workers rather than concentrating it in single individuals, and designing retention incentives for high-knowledge employees. Our estimates provide a benchmark for evaluating such investments: losing a worker with relevant product-specific experience generates a meaningful short-run reduction in targeted entry probabilities and a substantial decline in product-level sales. At the macro level, our findings suggest that worker mobility events—whether voluntary quits, retirements, or deaths—can influence the reallocation of production across product lines and, in the aggregate, shape growth through the diffusion or destruction of knowledge capital. While we do not attempt a full aggregation, our results provide disciplined inputs for quantifying how shocks to the distribution of worker knowledge propagate to firm dynamics and industry evolution.

The remainder of the paper proceeds as follows. Section 2 presents a short theoretical framework that guides the empirical analysis and generates testable predictions. Section 3 describes the empirical design, including measurement of worker product experience and the event-study specification. Section 4 details the data sources and sample construction. Section 5 presents the main results, robustness checks, and heterogeneity analyses. Section 6 concludes by discussing implications for hiring, retention, and policies that affect the portability of knowledge across firms.

2 Conceptual Framework

Entry Decision and Firm Capabilities. We model the decision of firm i to enter product j in year t as a threshold rule: entry occurs if an underlying capability index C_{ijt} exceeds zero. The capability index is

$$C_{ijt} = \theta_0 + a\,T_{ijt} + b\,S_{ijt} + r\,W_{it} - F_{ij} + u_{ijt}, \tag{1}$$

where T_{ijt} is transferable, worker-embodied, product-specific knowledge that is portable across employers with marginal contribution $a > 0$; S_{ijt} is non-transferable, firm-embedded knowledge such as routines, blueprints, other assets with marginal contribution $b > 0$; and W_{it} is non-transferable, match-specific knowledge with marginal contribution $r > 0$. u_{ijt} is an idiosyncratic entry shock and F_{ij} is a fixed cost of product expansion. While T_{ijt} and S_{ijt} capture codifiable or institutionalized expertise, W_{it} reflects the accumulated understanding between a specific worker and the firm about how to deploy that expertise productively, for instance, which suppliers to approach, how to navigate internal approval processes for a new line, or how to coordinate with colleagues whose cooperation is required for new product entry. Loss of W_{it} therefore reduces the capability index directly, making entry less likely even holding the stock of codifiable knowledge constant. We index match-specific capital W_{it} at the firm–year level rather than the firm–product–year level, reflecting the assumption that a worker's relational capital and organizational knowledge benefits the firm broadly rather than being tied to specific product lines. This contrasts with transferable knowledge T_{ijt}, which is product-specific by construction.[5]

Let $\bar{C}_{ijt} \equiv \theta_0 + aT_{ijt} + bS_{ijt} + rW_{it} - F_{ij}$ denote the deterministic component of the capability index, that is, C_{ijt} evaluated at its systematic part before the realization of the idiosyncratic entry shock u_{ijt}. New product entry occurs if $\bar{C}_{ijt} + u_{ijt} \geq 0$, and therefore

$$\Pr\{u_{ijt} \geq -\bar{C}_{ijt}\} \equiv \Lambda(\bar{C}_{ijt}) = \Pr\{X_{ijt} = 1\}, \tag{2}$$

for an increasing link function $\Lambda(\cdot)$ induced by the distribution of u_{ijt}. Let $X_{ijt} \in \{0, 1\}$ indicate entry of firm i into product j in year t (a $0 \to 1$ transition), and restrict attention to the risk set in which product j has not been produced by firm i prior to t. [6]

[5] We adopt the firm–year formulation for three reasons. First, it preserves a clean separation between the product-specific dimension of capability, which loads entirely on T_{ijt}, and the worker–firm dimension, which loads on W_{it}, making the empirical decomposition in Section 3.3 interpretable. Second, the leading components of match-specific capital we have in mind, such as familiarity with internal approval processes, knowledge of who in the firm to consult, and informal authority within the organization, are plausibly more firm-wide than product-specific. Third, to the extent that a portion of true match capital is in fact product-specific, it would be absorbed into T_{ijt} in our empirical implementation and inflate the estimated experience-graded coefficient β_A. Our reported Share_T therefore represents an upper bound on the contribution of genuinely portable knowledge whenever match capital has a product-specific component, and the qualitative ordering of our heterogeneity results, in particular the attenuation for core products, is unaffected.

[6] A profit-based microfoundation can deliver (1) and (2). Denote V_{ijt} as the expected net value of adding product j at time t. Let the firm enter if the expected value of entry is non-negative:

$$V_{ijt} \;=\; \Pi\big(T_{ijt}, S_{ijt}, W_{it}\big) \;-\; F_{ij} \;+\; \varepsilon_{ijt},$$

where $\Pi(\cdot)$ is the operating value of the product line and ε_{ijt} is an i.i.d. entry shock. A first-order approximation of $\Pi(\cdot)$ around the current state yields $\Pi(T, S, W) \approx \theta_0 + aT + bS + rW$, which delivers (1). If ε_{ijt} is type-I extreme value, (2) follows with a logit link; with normal ε_{ijt}, it follows with a probit link.

Extensive Margin Response to Disappearance Event. Conceptually, let T_{ijt} summarize the stock of portable product-j know-how available in firm i at time t. Let each worker ℓ in the firm's workforce $\mathcal{L}_{it}$ contribute additively to capabilities, such that $T_{ijt} = \sum_{\ell \in \mathcal{L}_{it}} e_{\ell j}$, where $e_{\ell j}$ denotes worker ℓ's prior product-j knowledge. Consider an unexpected worker disappearance of worker d employed by firm i at time τ. Let e_{dj} denote worker d's prior product-j knowledge and w_{di} their contribution to firm-level match capital W_{it}. The event removes two components:

$$\begin{aligned} T_{ijt}^{+} &= T_{ijt}^{-} - e_{dj}, && \text{(loss of transferable knowledge)} \\ W_{it}^{+} &= W_{it}^{-} - w_{di}, && \text{(loss of match-specific knowledge)} \end{aligned}$$

where superscripts $-$ and $+$ denote pre- and post-event values. Under this specification, the death of worker d reduces T_{ijt} by exactly e_{dj}. At the same time, S_{ijt} does not mechanically fall on impact by assumption (it is firm-embedded and non-personal), though it may adjust endogenously with lag.[7] In our empirical application, we measure product knowledge e_{dj} using workers' experience in producing a good j.

We assume that the match-specific component w_{di} is mean-independent of prior product knowledge e_{dj}, conditional on observable match characteristics:

$$\mathbb{E}[\, w_{di} \mid e_{dj}, \mathbf{H}_{di\tau} \,] = \mathbb{E}[\, w_{di} \mid \mathbf{H}_{di\tau} \,], \tag{3}$$

where $\mathbf{H}_{di\tau}$ collects predetermined match characteristics. The economic content of this assumption is that the quality of the worker-firm match does not systematically depend on how many years the worker previously spent producing any particular product at other firms. This is plausible because match-specific capital w_{di} is generated within the current employment relationship, through firm-specific relationship-building, learning internal processes, and adapting to organizational routines, rather than imported from prior employers. Since e_{dj} is accumulated before the current match formed, it is predetermined with respect to w_{di}. One might worry that firms strategically recruit workers with product-j experience precisely because they intend to enter product j, inducing a correlation between e_{dj} and w_{di} through shared organizational intent. However, under our within-firm, cross-

Up to normalization, the latent index in (1) can be written as $C_{ijt} = V_{ijt}$, so u_{ijt} corresponds to ε_{ijt}.

[7] This assumption may be violated when the deceased worker was the primary custodian of firm-embedded knowledge, for instance as the sole operator of a specialized process. In such cases, the effective stock of S_{ijt} available to the firm may erode even though underlying routines nominally remain. Our heterogeneity analysis in Section 5.3 shows substantially attenuated effects for core products relative to non-core ones, consistent with the view that firms with deep embedded knowledge (large, distributed S_{ijt}) are better insulated from individual worker losses.

product comparison, the orthogonality condition requires only that the match-specific component w_{di} does not vary systematically across products in proportion to e_{dj}, conditional on $\mathbf{H}_{di\tau}$. This is substantially weaker than requiring w_{di} to be unconditionally independent of e_{dj}. Instead, the orthogonality condition requires that, conditional on $\mathbf{H}_{di\tau}$, the match-specific component does not generate systematic post-event differences across products that scale with e_{dj}. Section 3 discusses the empirical design features that support this assumption.

Entry Response to a Disappearance Event. A first-order approximation of the entry response around the disappearance event yields

$$\Delta p_{ijt} \;\equiv\; \Pr\{X_{ijt}=1 \mid t \geq \tau\} - \Pr\{X_{ijt}=1 \mid t < \tau\} \;\approx\; -\Lambda'\big(\bar{C}_{ijt}\big)\,\big(a\,e_{dj} + r w_{di}\big). \tag{4}$$

Equation (4) is a first-order approximation evaluated at the pre-event capability level $\bar{C}_{ijt}$, and is therefore local: it characterizes the response for firms near the entry threshold where $\Lambda'(\bar{C}_{ijt})$ is non-negligible. For firms far from the threshold, the same knowledge shock produces little change in observed entry probability, not because knowledge is unimportant, but because entry was never a likely possibility. Our restriction to the risk set ensures that entry is feasible, and the non-trivial baseline entry rate in the estimation sample supports that many firm–product cells are near the entry margin where $\Lambda'(\bar{C}_{ijt})$ is non-negligible.

Equation (4) yields three testable implications. First, the decline in entry probability should be increasing in the deceased worker's prior product experience e_{dj}: larger e_{dj} implies a larger loss of transferable knowledge. Second, the slope of this relationship identifies the marginal contribution of transferable knowledge (proportional to a), while any intercept shift common across products reflects firm-level disruption or match-specific losses (proportional to r). Third, because the orthogonality condition (3) ensures that w_{di} does not covary with e_{dj}, the experience-graded and experience-invariant components can be separately identified. Hence

$$\frac{\partial\,\Delta p_{ijt}}{\partial e_{dj}} \;\approx\; -\Lambda'\big(\bar{C}_{ijt}\big)a, \tag{5}$$

because w_{di} is mean-independent of e_{dj} conditional on $\mathbf{H}_{di\tau}$ by (3).

Intensive-Margin Response. The same capability logic implies effects on product-level sales. Let Y_{ijt} denote sales of product j by firm i at time t, and suppose that, conditional on product j being produced, expected sales are increasing in the firm's operating capability:

$$\mathbb{E}[\,Y_{ijt} \mid X_{ijt}=1\,] = \Psi(T_{ijt}, S_{ijt}, W_{it}), \qquad \frac{\partial\Psi}{\partial T} > 0, \quad \frac{\partial\Psi}{\partial S} > 0, \quad \frac{\partial\Psi}{\partial W} > 0. \tag{6}$$

This reduced-form formulation captures the idea that transferable knowledge, firm-embedded knowledge, and match-specific capital improve the firm's ability to operate a product line efficiently, maintain quality, and serve customers.[8] Following the disappearance of worker d, a first-order approximation implies

$$\Delta Y_{ijt} \; \approx \; -\Psi_T \, e_{dj} - \; \Psi_W \, w_{di}, \tag{7}$$

where Ψ_T and Ψ_W denote the relevant local derivatives of $\Psi(\cdot)$. Thus, as with entry, the decline in sales should be more pronounced for products in which the disappeared worker had greater prior knowledge. Under Equation (3), the experience-graded component of the sales response identifies the marginal contribution of transferable knowledge, while any experience-invariant shift reflects firm-wide disruption or match-specific losses.

Testable Implications. The framework delivers four testable predictions. First, the probability of product entry should decline following an unexpected worker disappearance, with the decline increasing in the worker's prior product-specific knowledge e_{dj}. Second, this experience-graded penalty should be larger for workers who embody more valuable transferable knowledge, such as more educated workers. Third, the penalty should be attenuated in settings with deeper firm-embedded knowledge, such as core product lines. Fourth, analogous experience-graded declines should appear on the intensive margin of product sales. Flat pre-event trends across high- and low-experience products constitute the primary empirical test of the identifying assumption.

3 Empirical Design

This section maps the conceptual framework to observables and outlines our empirical designs. We begin by describing how we measure the core variables in the model (Section 3.1). We then present three complementary empirical approaches. First, we estimate dynamic average marginal effects using an event-study design for product entry at the extensive margin (Section 3.2). Second, we complement these dynamic estimates with a static panel

[8]Equation (6) is deliberately reduced-form: any structural model in which T_{ijt}, S_{ijt}, and W_{it} enter operating profits monotonically delivers the same first-order prediction. For instance, a log-linear specification $\log Y_{ijt} = \psi_0 + \psi_T T_{ijt} + \psi_S S_{ijt} + \psi_W W_{it} + \nu_{ijt}$ implies $\Delta \log Y_{ijt} \approx -\psi_T e_{dj} - \psi_W w_{di}$, which maps directly into our empirical specification on the inverse hyperbolic sine of sales. Equivalently, a CES production function in which worker-embodied knowledge enters as one input bundle yields the same local approximation around the pre-event allocation up to a rescaling of Ψ_T and Ψ_W. The empirical exercise therefore identifies a local elasticity of sales with respect to lost product-specific experience that is robust to a range of structural assumptions on how knowledge enters production, while the precise mapping from Ψ_T to a deep parameter requires committing to a specific functional form.

regression that quantifies average post-event effects on both product entry and product-level sales (Section 3.3). Third, we use heterogeneity analyses to relate the empirical results back to the mechanisms emphasized in the conceptual framework (Section 3.4). We describe each in turn.

3.1 Measurement

Our conceptual framework contains two objects that must be operationalized in the data: an unexpected worker disappearance and the worker's portable, product-specific knowledge. We proxy unexpected disappearances with sudden and unexpected worker deaths identified from administrative cause-of-death records. We measure portable product-specific knowledge using the deceased worker's prior experience in product j, denoted e_{dj}, which we recover from the product portfolios of the worker's former employers.

Concretely, our linked employer–employee data allow us to observe each worker's full employment history together with the products produced by each prior employer. We therefore define worker d's experience in product j as the number of years in which the worker was employed at firms that actively produced product j. For example, if a worker was previously employed at a firm producing a particular type of circular glass bottle, we classify that worker as having experience in that product line and measure e_{dj} as the cumulative years of such exposure. Grounding e_{dj} in *prior* employer product portfolios is a deliberate design choice. It makes the experience measure predetermined with respect to the current worker–firm match: the relevant product knowledge was accumulated before the worker joined the current firm and therefore cannot have been shaped by the firm's subsequent product-entry decisions.

The identification of sudden worker disappearances through worker deaths avoids spurious selection effects. Alternative identification strategies of sudden worker departures, such as using voluntary quits or retirements, would conflate the knowledge channel with selection: workers who quit may do so because the firm's product expansion prospects are poor, and retirements may be anticipated and planned for. Deaths, by contrast, are sudden and unforeseeable, ensuring that pre-event firm behavior is not contaminated by anticipation effects. Additional details on the construction of sudden-death events and prior product experience are provided in Section 4.

3.2 Event-Study for Product Entry

Event Study Design. To test the predictions of the capability-threshold model, we estimate an event-study specification around the unexpected death of workers in firm i at time τ, where $e_{i\tau,j}$ denotes the dying workers' years of prior experience of producing product j. An event is defined as a 9-year event window of a firm i-product j combination where the death occurs in the middle of the window, and where the dying worker has at least some experience of producing product j. Our main outcome is product entry of firm i into product j. The unit of observation is the firm–product–year. Thus, for the baseline specification, X_{ijt} is an indicator equal to one if firm i starts producing product j in year t.

We define event time as $k = t - \tau$ and estimate event-time effects of worker disappearance. First, we consider the overall effect of worker disappearance, scaled by the years of experience of those workers, that is the effect δ_k in

$$X_{ijt} = \sum_{k=-4}^{+4} \delta_k \, I_{\text{period}_k} e_{dj} + \gamma_{ij} + \eta_t + \varepsilon_{ijt}, \tag{8}$$

Coefficient δ_k then represents the effect of the disappearance of a year of experience among the firm's workforce. Note that the event-study coefficients δ_k are reduced-form objects that identify average marginal effects proportional to the structural parameter a from Equation (1); they are not equal to a itself, which is not separately identified without additional distributional assumptions on $\Lambda(\cdot)$.

If multiple workers die unexpectedly in the same firm-year, we aggregate exposure across the deceased workers, defining

$$e_{i\tau,j} = \sum_{d \in \mathcal{D}_{i\tau}} e_{dj},$$

and interpret the estimates as the effect of a one-year increase in cumulative lost product-specific experience.[9] Standard errors are clustered at the firm-product level to account for within-firm-product correlation induced by the death event.[10]

Identification and Interpretation. The identifying variation in equation (8) comes from within-firm, across-product differences in exposure to the deceased worker's prior product-specific experience. Intuitively, we compare, within the same firm and year, products for

[9] Some firms experience sudden deaths in multiple years. In the event-study design, we treat each death-year as a separate event, so the same firm–product may enter the stacked sample more than once with different event times. The resulting coefficients should therefore be interpreted as average effects per death event rather than per firm.

[10] The key identifying regressor, the interaction $I_{\text{period}_k} \times e_{dj}$, varies at the firm-product level, since different products within the same firm receive different treatment intensities depending on the deceased worker's prior product-specific experience. Following Abadie et al. (2023), we cluster at the level at which the effective treatment varies in the regression, which is the firm-product level. This accounts for serial correlation in ε_{ijt} within firm–product cells over time, the primary source of residual dependence given the within-firm, cross-product identifying variation.

which the deceased worker had more prior experience to products for which the worker had less experience. Firm–product fixed effects absorb persistent differences in a firm's propensity to enter particular products, while year fixed effects absorb aggregate shocks common to all firms.

The interacted coefficients δ_k are the core parameters of interest: they trace the experience-graded response to the disappearance of worker experience and therefore isolate the relative penalty associated with the loss of portable, product-specific knowledge.

This mapping follows directly from the framework. Equation (4) implies that the death of worker d reduces the entry probability by an amount proportional to $ae_{dj} + rw_{di}$. Under the orthogonality condition in Equation (3), the match-specific component w_{di} does not vary systematically with e_{dj} conditional on predetermined worker characteristics. As a result, the slope of the response with respect to e_{dj} identifies the marginal contribution of transferable knowledge in probability space. Formally, the coefficients δ_k correspond to average marginal effects proportional to $-\Lambda'(\bar{C}_{ijt})a$.

The identifying assumption is a within-firm, cross-product parallel-trends condition: absent the death event, the gap in entry probabilities between high-experience and low-experience products would have remained stable over time within the firm. This condition is weaker than the cross-firm parallel-trends assumption of a standard difference-in-differences design because all comparisons are made within the same firm facing the same shock. Identification would fail if, even in the absence of the death, the firm were already on differential pre-trends precisely in the product lines in which the worker had prior experience.

Two concerns are particularly relevant. First, firms may have anticipated the worker's disappearance and already begun scaling back entry attempts in the relevant product lines. The use of sudden and unexpected deaths makes this channel implausible by construction. Second, the worker's prior product experience might be correlated with the firm's underlying demand conditions for product j, for instance because workers are hired from firms in the same narrow industry. Our use of experience accumulated at prior employers rather than at the current firm severs the mechanical link between e_{dj} and the current firm's own product demand. Moreover, such confounds would imply differential pre-event trajectories across products, which would appear as non-zero coefficients for $k < 0$. The flat and statistically indistinguishable pre-event estimates reported in Section 5.1 therefore provide the central empirical test of the identifying assumption.[11]

[11] Recent work on staggered difference-in-differences designs has raised concerns about negative weighting when treatment effects vary over time or across cohorts (Goodman-Bacon, 2021; Sun and Abraham, 2021). In our setting, these concerns are mitigated by two features. First, our primary identifying variation is cross-product within firm-year, not cross-firm over time. Second, we estimate fully dynamic specifications

3.3 Static Specification and Descriptive Capability Decomposition

To complement the dynamic event-study evidence on product entry, we estimate a static panel specification that summarizes the average post-event effect of losing targeted worker-embodied knowledge. This specification allows us to examine both the extensive margin of product entry and the intensive margin of product-level sales in a unified framework. The unit of observation is the firm–product–year, and the estimating equation is

$$E_{ijt} = \beta \text{Shock}_{ijt-1} + \gamma_{ij} + \eta_t + \mathbf{Z}'_{it}\Gamma + \epsilon_{ijt}, \tag{9}$$

where γ_{ij} are firm–product fixed effects, η_t are year fixed effects, and $\mathbf{Z}_{it}$ is a vector of predetermined firm-level controls, with our baseline specification including lagged log firm employment. We report specifications varying the choice of scale control in Table 3. As robustness, we further expand the set of control variables Z to include a rich set of worker-level covariates (dummies for five age groups, gender, tenure, income, as well as fixed effects for 3-digit occupation codes) in Appendix Table A.2.

The dependent variable E_{ijt} captures product-level adjustment on either the extensive or intensive margin. In the extensive-margin specifications, E_{ijt} is an indicator for whether firm i starts producing product j in year t.[12] In the intensive-margin specifications, E_{ijt} is the change in the inverse hyperbolic sine of sales of product j from year $t-1$ to t.

The key regressor, Shock_{ijt-1}, measures the loss of targeted capabilities induced by unexpected worker deaths at firm i in year $t-1$. We construct two versions. The first is the cumulative years of product-j experience among workers who died unexpectedly, which yields coefficient β_A. The second is a count of how many such workers had any product-j experience, which yields coefficient β_B. By construction, $\text{Shock}_{ijt} = 0$ for products outside the deceased workers' expertise, so the resulting coefficients identify penalties specific to relevant human-capital losses rather than broad firm-level disruption. Standard errors are clustered at the firm–product level.

The two shock measures serve different purposes. The coefficient β_A captures the marginal

that do not impose homogeneity across event time.

[12] We use a linear probability model as our baseline specification. In principle, one could consider nonlinear link functions such as logit, probit, or Poisson pseudo-maximum likelihood (PPML). In our setting, however, the combination of a rare binary entry outcome and high-dimensional firm–product fixed effects makes such alternatives less attractive. Conditional fixed-effects logit/probit are not feasible with our fixed-effects structure, and PPML drops the large set of firm–product cells with all-zero entry histories because they are perfectly predicted by the fixed effects. This induces a substantial change in the effective sample and estimand, shifting attention to the selected subset of firm–product pairs with within-cell outcome variation. By contrast, the linear probability model retains the full risk set and delivers coefficients that are directly interpretable as average marginal effects in probability units.

effect of an additional year of lost product-specific experience, and therefore corresponds to the portable-knowledge channel emphasized in the conceptual framework. The coefficient β_B captures the average total loss associated with the disappearance of a worker who had any relevant product experience. Evaluating the marginal effect at the sample mean of experience among deceased workers with strictly positive product-j experience, denoted $\bar{e}_{dj}$, yields a descriptive decomposition of the total targeted loss into a transferable and a residual component:

$$\text{Share}_T \equiv \frac{\beta_A \cdot \bar{e}_{dj}}{\beta_B}, \qquad \text{Share}_W \equiv 1 - \text{Share}_T = \frac{\beta_B - \beta_A \cdot \bar{e}_{dj}}{\beta_B}. \tag{10}$$

This decomposition is motivated by the additive structure of the capability index in Equation (1). Under additive separability, the total targeted loss associated with the disappearance of a worker with relevant experience can be written as

$$\beta_B = \underbrace{\beta_A \cdot \bar{e}_{dj}}_{\text{transferable knowledge}} + \underbrace{\left(\beta_B - \beta_A \cdot \bar{e}_{dj}\right)}_{\text{residual match-specific component}}.$$

The first term captures the average loss that scales with prior product-specific experience and therefore maps to transferable knowledge. The second term captures the residual experience-invariant component, which we interpret as match-specific capital or other firm-level losses not captured by experience alone.

This interpretation rests on two maintained assumptions. First, the capability index is additively separable in transferable knowledge T_{ijt} and match-specific capital W_{it}, so the marginal contribution of one component does not depend on the level of the other. Second, the orthogonality condition in Equation (3) implies that the experience-based coefficient β_A isolates the marginal contribution of transferable knowledge rather than a mixture of transferable and match-specific effects. If these assumptions fail—for example, because transferable and match-specific knowledge are strong complements—the decomposition in Equation (10) should be interpreted descriptively rather than structurally. In that case, Share_T measures the portion of the targeted loss that scales with prior experience, rather than a structural share.

3.4 Heterogeneity and Mechanism Tests

We next examine whether the empirical effects vary in ways predicted by the capability framework. The conceptual model highlights three margins along which the impact of a worker disappearance should differ: the amount of transferable knowledge embodied in the worker, the extent of firm-embedded knowledge already available for the product, and the firm's overall reliance on skilled labor. We therefore estimate Equation (9) across three pre-specified sample partitions. These heterogeneity tests are not intended as separate identification strategies; rather, they assess whether the pattern of estimates lines up with the mechanisms emphasized by the framework. Because these tests are intended to summarize average differences in the targeted penalty across environments, we implement them using the static specification in Equation (9).

Worker educational attainment. Our first prediction is that the penalty from a worker disappearance should be larger when the worker embodied more valuable transferable knowledge. We proxy this with educational attainment. Because education is a standard proxy for the breadth and depth of human capital, the framework predicts larger capability losses for workers with a university degree than for workers without one. Formally, this corresponds to allowing the marginal contribution of experience to vary with worker skill, so that $a = a(\text{education})$. If more educated workers convert prior product experience into more valuable knowledge, then for a given level of e_{dj} the loss $a \cdot e_{dj}$ should be larger for high-education workers.

Core versus non-core products. Our second prediction concerns the buffering role of firm-embedded knowledge S_{ijt}. When S_{ijt} is large, the firm's capability index should be less sensitive to the loss of any one worker because routines, blueprints, and organizational knowledge already support the relevant product line. We proxy this dimension by distinguishing between core and non-core products. Core products, defined as products belonging to the firm's primary two-digit product category, are precisely those for which firm-embedded knowledge is likely to be deepest. The framework therefore predicts that the targeted penalty from worker disappearance should be attenuated for core products relative to non-core products.

Firm skill intensity. Our third prediction is that the penalty from a worker disappearance should be larger in firms whose production relies more heavily on worker-embodied knowledge. We proxy this dependence using firm-level skill intensity, measured by the share of skilled workers in the workforce. In firms with greater skill intensity, transferable worker knowledge T_{ijt} should account for a larger share of total capability relative to codified routines S_{ijt}, making product expansion more vulnerable to the loss of any one worker. Conversely, in lower-skill firms, production should rely relatively more on codified routines and physical capital, so the disappearance of a specific worker should have a smaller targeted effect. We test this prediction by splitting the sample at both the within-sector and across-sector medians of firm skill intensity.

4 Data

In the following, we outline our data sources in section 4.1 and discuss our sample in section 4.2.

4.1 Data Sources

We leverage four main Swedish data sources: a firm level data set, a worker level dataset, a product-level data set, and the death registry. The set of variables used per dataset, together with the period covered, is summarized in Appendix Table A.1.

The firm-level data derives from *Företagsdatabasen* (FDB), which contains value added, wage bills, and other production costs drawn from Swedish Tax Authority records of firms' balance sheets.

The worker-level data is the *Longitudinell Integrationsdatabas för Sjukförsäkrings- och Arbetsmarknadsstudier* (LISA), covering all Swedish private-sector workers with information on income, education, age, and occupation (see also Balke and Lamadon, 2022; Saez et al., 2019). It includes information such as income, education, and age. Occupations are reported according to the Swedish Standard Classification of Occupations (SSYK).[13] We can link employers and employees using firm identifiers and therefore track workers' experience based on the firm they work at. The reliability and quality of this data is regarded as very high, since it is based on tax reports by firms and misreporting is punishable by law.

The product-level data is drawn from the dataset *Industrins Varuproduktion (IVP)*, which is based on surveys on the production of Swedish manufacturing firms, and has previously been used by Carlsson and Skans (2012). The dataset includes all manufacturing firms with at least five employees, and contains information on what products they produce up to the 8-digit Combined Nomenclature (CN) level. For each year, firms report both quantity and price for each product. Using the worker-level data, this product-level data allows us to track workers' experience in specific lines of production. For our empirical specification, we use 4-digit product codes.

Our analysis is therefore restricted to the Swedish manufacturing sector. This restriction is

[13] The base for SSYK is the international standard classification of occupations with reference year 2008 (ISCO-08).

a consequence of data availability as granular, firm-level product portfolio data of the kind required to construct e_{dj} is not available for service industries in our sample period, but it is also a substantively well-motivated setting for the research question. Manufacturing is precisely the sector where product-specific, worker-embodied knowledge is most directly tied to physical production processes: experience with a particular product line reflects knowledge of materials, tolerances, equipment configurations, and supply chain relationships that are difficult to codify and genuinely portable across employers. The 4-digit CN classification further ensures that "product experience" is defined at a level of granularity—distinguishing, for example, between specific types of milk depending on whether it has added sweetening or not—where worker expertise is meaningfully product-specific rather than broadly sectoral.

The manufacturing focus does, however, bound the external validity of our estimates. The magnitudes we document should be interpreted as applying to firms whose expansion decisions depend on embodied production knowledge of the kind that is prevalent in manufacturing. In knowledge-intensive service industries, such as consulting, software, finance, analogous mechanisms are likely to operate, but through different channels such as client relationships, algorithmic expertise, or project management know-how, and the relevant experience measure would differ accordingly.

We merge the employer–employee records (LISA) with firm financials (FDB) and product output (IVP) using firm identifiers, and link workers' past employers' product mix to construct product-specific prior experience measures. For each worker in our sample, we recover the complete sequence of prior employers from the LISA employment histories. For each prior employer spell, we retrieve the set of products that employer produced in each year of the spell from the IVP dataset. We then define worker d's prior experience in product j as the total number of years in which the worker was observed employed at a firm that was actively producing product j in that year. Formally,

$$e_{dj} \;=\; \sum_{f\in\mathcal{F}_d}\sum_{t\in\mathcal{T}_{df}} \mathbf{1}[j\in\mathcal{P}_{ft}]\,, \tag{11}$$

where $\mathcal{F}_d$ is the set of prior employers of worker d (excluding the current employer), $\mathcal{T}_{df}$ is the set of years worker d was employed at firm f (restricted to spells of at least one calendar year), and $\mathcal{P}_{ft}$ is the set of products produced by firm f in year t.

A number of construction choices merit explicit discussion. First, regarding breadth of experience assignment: when a prior employer produced multiple products simultaneously, we assign the worker experience in all of them. This choice may introduce measurement error by attributing experience to workers who were not directly involved in every product line their employer maintained. To the extent that this creates classical measurement error in e_{dj}, it would attenuate our estimates toward zero, implying that our main estimates are conservative with respect to the portable-knowledge channel. We address remaining measurement noise in three ways. First, we use continuous years of experience rather than a binary indicator: under continuous coding, a product line that the worker's prior employer produced for only one year contributes one unit to e_{dj}, whereas a product line produced for ten years contributes ten units, so accidental assignment of a single-year product carries proportionally less weight than under a 0/1 coding. Second, we show in Appendix Table A.4 that the results are robust to coarser product classifications, where assignment errors from diversified employers are mechanically less frequent. Third, regarding minimum tenure: we require a worker to have been observed at a prior firm for at least one full calendar year before assigning any product experience from that spell. Spells shorter than one year are coded as zero on the grounds that transitory employment is unlikely to generate the depth of production knowledge that is relevant for firm-level product entry decisions.

Finally, we also leverage information on worker deaths from the registry *Dödsorsaksregistret* which includes information on date and cause of death as classified through the International Classification of Disease ICD, version 10 (see Brooke et al., 2017, for an extensive description of the dataset).

Table 1: ICD-10 codes used to define sudden unexpected death

ICD-10 code	**Description**
Natural causes	
I22–I23	Acute myocardial infarction
I46	Cardiac arrest
I50	Congestive heart failure
I60–I69	Stroke
R95–R97	Sudden death from unknown causes
Unnatural causes	
V00–V89	Traffic accidents
V90–V99, X00–X59, X86–X90	Other accidents and violence

Notes: The table reports the ICD-10 causes of death used to identify sudden unexpected worker deaths in the analysis, following Andersen and Nielsen (2011). Suicides and violent deaths due to relatives are excluded.

In the medical literature, a sudden and unexpected death (SUD) is defined "as a natural, unexpected fatal event that occurs within one hour of the beginning of symptoms in an apparently healthy subject or in one whose disease was not so severe that such an abrupt outcome could have been predicted" (Lim et al., 2010). We identify unexpected worker deaths using information on the cause of deaths from the registry, building on Andersen and Nielsen (2011). Similar to them, we consider the list of ICD-10 causes of death in Table 1 to identify sudden deaths. Among natural causes of death, we thus consider acute myocardial infarction (ICD-10: I22-I23), cardiac arrest (I46), congestive heart failure (I50), stroke (I60-I69) and sudden deaths by unknown causes (R95-R97). Among unnatural deaths, we consider traffic accidents (V00-V89) and deaths caused by other accidents and violence (V90-V99, X00-X59, and X86-X90) which excludes suicides or violent deaths due to relatives. Because the registry reports cause and date of death, our event definition is independent of separations driven by firm conditions (e.g., layoffs) and avoids conflating deaths with other forms of worker exit.

4.2 Sample

Our data covers the years 1997–2019. Because the death registry coverage ends in 2016, our event-study analysis includes deaths occurring through 2016, with post-event outcomes observed through 2019. We restrict the population to workers 15–65 years of age with an observed occupation code, and to firms with at least five employees. Table 2 includes summary statistics for our data sample. Our data spans around 1.2 million unique workers working for around 25,000 firms. Workers are on average 43 years old. Around 31% of workers have a college degree and 1% hold a PhD. A worker is likely to leave the current employer with probability 17% each year. Our firm data shows that firms' labor productivity grows at around 2% per year on average. The median firm employs only 12 workers, implying that individual worker exits are likely to be operationally salient. Around 7% of firms adopt a new product every year. The mean age at death is 51, and is as expected higher than the sample average age. There are around 2,000 sudden death events in our matched sample where a firm experiences at least one death, out of the 242,380 firm-year observations in total, totaling about 1% of all matched observations. This means that death events are relatively rare, but more frequent than for example in the study by Sauvagnat and Schivardi (2023), who look at deaths of Italian CEOs (0.1%) and similar in magnitude to Bertheau et al. (2022) but smaller than Jäger et al. (forthcoming).[14] Among deceased workers with positive prior product experience, the average number of experience years is $\bar{e}_{dj} = 5.33$ at 4-digits.

[14]Note that Jäger et al. (forthcoming) do not observe the cause of death and hence start from a larger sample of disappearing workers. To focus on unexpected deaths, they discard workers who had been on sick leave in the five years prior to the death event.

Table 2: Summary statistics

Variable	**Obs.**	**Mean**	**SD**	**Median**
A. Workers				
Age	8,105,124	42.91	11.62	43.00
Female	8,105,124	0.24	0.43	0.00
Income	8,105,124	3,556.81	2,054.78	3,262
Less than high school	8,105,124	0.69	0.46	1.00
More than high school	8,105,124	0.31	0.46	0.00
PhD	8,105,124	0.01	0.08	0.00
Stayer	7,643,679	0.83	0.37	1.00
Number of workers	1,168,308			
Number of firms	21,508			
Number of occupation groups (4-digit)	366			
B. Firms and products				
Labor productivity growth (Y/L)	212,932	0.02	0.34	0.03
Firm size	242,380	50.67	329.63	12.00
Product adoption	72,454	0.07	0.25	0.00
Products per firm	69,504	2.09	2.49	1.00
Number of firms	25,216			
Number of product codes (4-digit)	236			
C. Death events				
Age at death	2,027	50.81	11.68	54.00
Prior experience (4-digit, worker–product pairs)	12,368	5.33	4.47	4.00
Number of events	2,027			

Notes: The table reports summary statistics for the worker, firm–product, and death-event samples used in the analysis. Panel C reports statistics at two levels. "Age at death" and "Number of events" are at the deceased-worker level and count 2,027 unique workers. "Prior experience (4-digit, worker–product pairs)" is at the worker–product level: each deceased worker contributes one observation per 4-digit product in which they had positive prior experience, yielding 12,368 worker–product pairs across 2,027 workers. The reported mean $\bar{e}_{dj} = 5.33$ is therefore the average years of prior experience per worker–product pair conditional on positive experience, consistent with the definition used in the main text. On average, each deceased worker had positive prior experience in approximately 6.1 distinct 4-digit product lines (12,368/2,027).

5 Results

In the following, we show results for the dynamic event study specifications (Section 5.1) as well as the static effect specification (Section 5.2) and the examination of mechanisms (Section 5.3).

5.1 Event Study Results

Overall effect. Figure 1 reports the baseline event-study estimates based on equation (8). The pre-event coefficients are small in magnitude and statistically indistinguishable from zero, supporting the within-firm parallel-trends assumption. Focusing on impact, the disappearance of a year of experience reduces the probability that the firm enters the relevant product line immediately at the time of the shock. The impact coefficient ($k = 0$) is -0.0013, implying that each additional year of prior product-specific experience held by the disappeared worker lowers the on-impact probability of entering the corresponding product by about 0.13 percentage points. Evaluated at the sample mean of prior experience among unexpectedly deceased workers with positive product-specific experience, $\bar{e}_{dj} = 5.33$, this corresponds to a total on-impact decline of roughly 0.7 percentage points, or about 27 percent of the pre-event baseline entry probability among exposed products, $E[X_{ijt} \mid e_{dj} > 0, k = -1] = 0.0254$. The effect becomes larger in the first post-event year ($k = 1$) before moving back toward zero, indicating that the disruption is immediate but that its full implications for expansion decisions may unfold with some delay.

As shown in Appendix Figure A.1, worker disappearances are also followed by broader declines in firm-level labor productivity (defined as output per worker). One year after the event, log labor productivity falls by roughly 0.03 log points, indicating that the product-level entry response occurs alongside a broader contraction in firm performance. This magnitude is consistent with estimates from the displaced-worker literature showing persistent productivity losses following unexpected separations (e.g., Jacobson et al., 1993), and suggests that the product-level effects we document are not offset by reallocation to other activities.

To benchmark these estimates, we compare them to related findings in adjacent settings. Jaravel et al. (2018) find that the unexpected death of an inventor reduces follow-on innovation by collaborators by about 10%, while Azoulay et al. (2010) show that the death of a "superstar" scientist lowers collaborators' publication rates by roughly 5–8%. These benchmarks are not directly comparable to our estimates: their outcomes are intensive-margin rates of ongoing activity (citations, publications per year), whereas our extensive-margin outcome is a rare binary entry event with a baseline probability of about

Figure 1: Event study of death on product expansion.

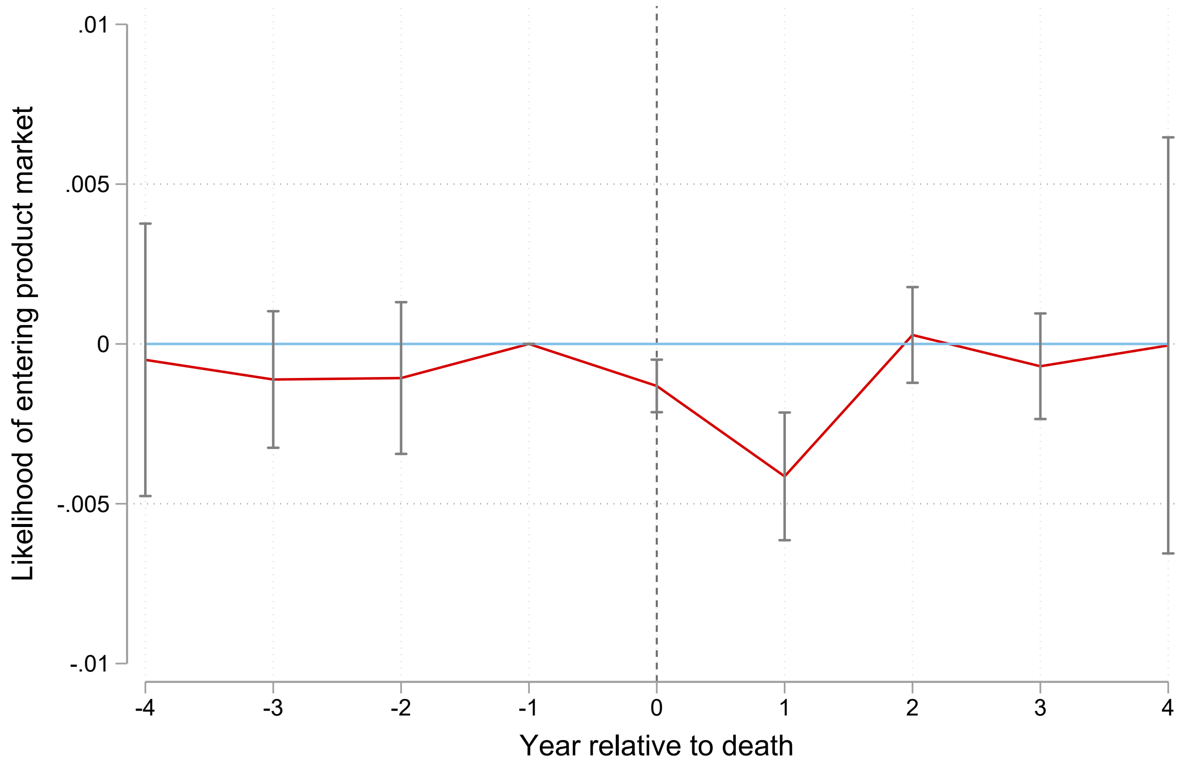

Notes: The figure shows the effect of an unexpected worker death on whether the firm starts producing a target product, scaled by the dying worker's prior years of experience in that product. The corresponding regression specification is equation (8). The figure plots coefficients δ_k, interpretable as the effect per year of prior product-specific experience lost.

2.5%. A 27% relative decline in an entry probability and a 10% relative decline in follow-on innovation therefore capture different objects, and we do not claim the magnitudes are directly rankable. What the three estimates share is a qualitative pattern: the unexpected loss of a knowledgeable individual produces an economically meaningful decline in the activity most closely connected to their expertise. Our contribution is to document this pattern on the production margin—product scope—where prior work has been silent, and to show that the penalty scales with the departed worker's prior product-specific experience. This pattern is consistent with the view that production know-how is often tacit and difficult to diffuse through codified channels. It also complements Bertheau et al. (2022), who show that unexpected worker separations reduce firm revenues on average: our product-level estimates suggest that these aggregate effects mask substantial heterogeneity across products depending on whether the departed worker had relevant experience.

To illustrate the economic magnitude of the impact effect, consider a simple back-of-the-envelope calculation. Evaluating the impact coefficient at the sample mean of prior product-specific experience among exposed deceased workers implies that the unexpected loss of a worker with average experience reduces the probability of entering the relevant product line on impact by roughly 0.7 percentage points. If the expected net value of entering that product is $10 million, the implied expected loss from the worker's disappearance is about $70,000. On this calculation, an expenditure of similar magnitude on retention, cross-training, or knowledge codification would be justified on expected-value grounds. For workers with above-average product experience, or for product lines with higher expected entry value, the implied benefit of preserving the underlying knowledge would be larger still.

We perform a sensitivity exercise in Appendix Figure A.2 where we restrict the sample to firms that do not replace the deceased worker during the event window, as measured by persistently depressed firm size. Results are qualitatively similar, reinforcing the interpretation that the effect reflects lost worker-embodied knowledge rather than transitional disruption costs associated with vacancy filling and onboarding (Jäger et al., forthcoming).

5.2 Static Estimation Results and Capability Decomposition

To formally test our theoretical predictions and summarize the average post-event effects, we turn to the baseline static panel regressions in Table 3. Our framework predicts that the loss of product-specific worker knowledge should reduce subsequent outcomes more strongly in product lines for which the deceased worker had greater prior experience. In the static specification, the coefficient on cumulative lost experience therefore captures the experience-graded component of the loss, while the worker-count specification captures the average effect of losing an exposed worker.

The strongest and most stable results appear on the intensive margin. In columns (1)–(3), each additional year of lost product-specific experience reduces product-level sales growth by about 0.013 log points, with a coefficient of -0.0134 in the lagged-employment specification. The loss of an exposed worker reduces sales growth by about 0.078 log points, with a coefficient of -0.0780 in the lagged-employment specification. These estimates are highly stable across specifications with no employment control, contemporaneous employment, and lagged employment. The sales results therefore provide strong evidence that worker-embodied knowledge matters not only for expansion opportunities, but also for production performance conditional on operating the product line.

Table 3: Production, product entry, and worker experience

	Sales growth in product j			**Product entry in product j**		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Cumulative lost experience						
Lost experience in j	-0.0134***	-0.0134***	-0.0134***	-0.00067***	-0.00067***	-0.00067***
	(0.00153)	(0.00153)	(0.00153)	(0.00009)	(0.00009)	(0.00009)
Observations	56,935,292	56,935,292	56,935,292	56,935,292	56,935,292	56,935,292
Panel B. Number of exposed workers lost						
Workers lost with experience in j	-0.0780***	-0.0780***	-0.0780***	-0.00618***	-0.00618***	-0.00618***
	(0.01107)	(0.01107)	(0.01107)	(0.00060)	(0.00060)	(0.00060)
Observations	56,935,292	56,935,292	56,935,292	56,935,292	56,935,292	56,935,292
Transferable share, Share$_T$	0.92	0.92	0.92	0.58	0.58	0.58
Year fixed effects	✓	✓	✓	✓	✓	✓
Firm–product fixed effects	✓	✓	✓	✓	✓	✓
Current employment (log)		✓			✓	
Lagged employment (log)			✓			✓

Notes: The table reports estimates from the static specification (9). The unit of observation is the firm–product–year. In columns (1)–(3), the dependent variable is the change in the inverse hyperbolic sine of product-j sales from year $t-1$ to t. In columns (4)–(6), it is an indicator for whether firm i starts producing product j in year t. In Panel A, the shock is the cumulative years of lost product-j experience among workers who died unexpectedly. In Panel B, the shock is the number of workers who died unexpectedly and had any experience in product j. Columns (2) and (5) add contemporaneous log employment; columns (3) and (6) add lagged log employment. $Share_T \equiv \beta_A \cdot \bar{e}_{dj}/\beta_B$ is the descriptive share of the targeted loss that scales with prior product-specific experience, evaluated at $\bar{e}_{dj} = 5.33$. Because Share$_T$ is a ratio of two estimated coefficients, it is not informative when β_B is small or imprecisely estimated. We report Share$_T$ only when β_B is statistically distinguishable from zero at the 1% level; cells failing either criterion are reported as "—". Standard errors clustered at the firm–product level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

On the extensive margin, the evidence is stable across specifications. In Panel A of columns (4)–(6), the cumulative-experience measure is negative and statistically significant throughout, with a coefficient of -0.00067 in each specification.[15] Thus, each additional year of relevant experience lost lowers the probability of subsequent entry into the corresponding product line. In Panel B, the worker-count specification is also negative and statistically significant throughout, with a coefficient of -0.00618 in columns (4)–(6). The extensive-margin results therefore no longer hinge exclusively on the cumulative-experience specification, although the latter remains closer to the conceptual framework's emphasis on the amount of product-specific knowledge lost rather than the mere number of exposed workers.

The sales and entry estimates are highly stable across alternative controls for firm size. Moving from specifications without employment controls to models including contemporaneous or lagged employment leaves the sales coefficients essentially unchanged. The same is true on the entry margin: both the cumulative-experience and worker-count coefficients are virtually identical across columns (4)–(6). This stability reinforces the interpretation that the estimated penalties reflect the loss of targeted worker-embodied knowledge rather than mechanical changes in firm scale following a worker death.

Appendix Table A.2 further shows that the results are not explained by observable characteristics of the deceased worker. Adding a rich set of worker-level controls—including age-group dummies, gender, tenure, income, and 3-digit occupation fixed effects—leaves the estimates qualitatively unchanged. The coefficients attenuate modestly in three of the four comparisons, while the count-based entry coefficient becomes slightly more negative; all estimates remain negative and statistically significant. This pattern suggests that the baseline relationship is not driven by observable differences in worker seniority, compensation, or occupational role. Rather, the economically relevant variation appears to come from the specific product knowledge embodied in the worker.

As a robustness check on product aggregation, Appendix Table A.4 re-estimates the baseline static specification using 2-digit product categories. The cumulative-experience coefficient remains negative and statistically significant at the 1% level on both the sales and entry margins. The worker-count coefficient is also negative and significant on the entry margin, but it is small and imprecisely estimated on the sales margin. As a result, the reported sales-margin ratio is mechanically large and not very informative at this level of aggregation,

[15] To interpret these magnitudes, we evaluate the coefficient in column (6) at the sample mean of 5.33 years of prior experience. This implies that the loss of a typical worker with product-relevant knowledge lowers the probability of entering the targeted product in the following year by about 0.0036, that is, 0.36 percentage points. Relative to the baseline entry probability among exposed products, this corresponds to roughly 14%.

while the entry-margin ratio $\text{Share}_T = 0.72$ sits somewhat above the 4-digit estimate of 0.58. We interpret this difference as a feature of the decomposition rather than a contradiction: because the ratio divides two estimated coefficients, both of which change with aggregation, modest changes in β_A or β_B translate into visible shifts in Share_T, especially when entry rates are low. The qualitative ordering is stable across aggregations—the experience-scaled component accounts for the majority of the targeted entry penalty on both margins—but we caution against reading the precise level of the entry-margin Share_T too literally. This is consistent with the framework in Section 2, which predicts the sign and rank ordering of the components but does not pin down a structural ratio.

As a final robustness and mechanism exercise, we test whether worker deaths induce reallocation across products within the same firm. Specifically, we construct an alternative exposure measure that assigns the deceased worker's prior product-specific experience either to all products within the firm or only to products in which the worker had no prior experience, while setting exposure to zero for the directly exposed product lines. This specification asks whether firms offset the loss of targeted knowledge by shifting activity toward unrelated products. The estimates in Appendix Table A.3 do not support a clean broad contraction across unrelated products. When exposure is assigned to all products within the firm, the coefficients are small and negative on both margins. When attention is restricted to non-exposed products, however, the coefficients become small and positive. We interpret this pattern as evidence of limited offsetting adjustment toward unrelated products rather than a large reallocation response.

Following our empirical strategy, we use the static coefficients to decompose the total capability shock into a component that scales with prior product-specific experience and a residual component that does not. The coefficient β_A captures the marginal effect of lost portable, product-specific knowledge accumulated through prior experience, while β_B captures the overall loss associated with the disappearance of an exposed worker. Evaluating β_A at the sample mean of experience among workers with positive product-relevant tenure, $\bar{e}_{dj} = 5.33$, yields the descriptive ratio

$$\text{Share}_T = \frac{\beta_A \cdot \bar{e}_{dj}}{\beta_B}.$$

In the baseline specification, this ratio is 0.92 on the sales margin and 0.58 on the entry margin. The sales-margin decomposition is especially informative because both underlying coefficients are large, stable, and precisely estimated; it suggests that most of the capability loss associated with worker disappearance reflects knowledge that could, in principle, have been transferred across firms. The entry-margin ratio points in the same direction, although we interpret it more cautiously because entry is rare and decomposition ratios are sensitive to small changes in the underlying coefficients. This interpretation resonates with Gathmann and Schönberg (2010) and Poletaev and Robinson (2008), who document substantial portability of task-specific and occupational human capital across employers.

5.3 Mechanisms: Education, Firm Skill, and Embedded Knowledge

Building on the baseline results, we explore heterogeneity across worker, firm, and product characteristics to map the empirical patterns back to the mechanisms emphasized in the conceptual framework.

First, we consider worker educational attainment, hypothesizing that the disappearance of a high-skill worker represents a larger loss of transferable knowledge, T_{ijt}. Table 4 strongly supports this prediction. On the extensive margin, the cumulative-experience coefficient is nearly twice as large for university-educated workers as for workers without a university degree ($\beta_A = -0.00133$ versus -0.00070). On the intensive margin, the same ordering appears even more strongly: the cumulative-experience coefficients are -0.0359 for high-skill workers and -0.0135 for low-skill workers, while the count-based coefficients are -0.2858 and -0.0743, respectively. The count-based entry coefficients point in the same direction (-0.02059 for high-skill workers versus -0.00617 for low-skill workers). Taken together, these magnitudes indicate that a given amount of prior product experience is substantially more consequential when embodied in more educated workers.

On the sales margin, the experience-scaled component accounts for nearly the full decline for low-skill workers, with $\text{Share}_T = 0.98$, but a smaller share for high-skill workers, with $\text{Share}_T = 0.62$. The same ordering appears on the entry margin, where the corresponding ratios are 0.613 and 0.317. We therefore interpret the education heterogeneity as evidence that high-skill workers generate larger absolute knowledge losses, while a larger fraction of their contribution appears to reflect complementary firm-specific value.

This is consistent with the view that the return to experience is higher for more skilled workers, but that a larger fraction of their contribution is embedded in organizational routines, relationships, and tacit coordination that do not travel one-for-one across firms (Becker, 1964; Lazear, 2009; Gibbons and Waldman, 2004). The pattern also has a natural managerial reading: retaining a low-skill worker with deep product experience preserves something close to the full operational value of that worker, since most of what is lost on their departure is portable; retaining a high-skill worker preserves both the portable component and a substantial complementary firm-specific component, so the marginal value of retention investments is greater but the case for cross-training and codification — which substitute primarily for the portable component — is correspondingly weaker.[16]

Second, we consider the buffering role of firm-embedded knowledge, S_{ijt}. The framework predicts that when a firm already possesses deep routines and organizational know-how in a product category, the loss of any one worker should be less consequential. We proxy this margin by focusing on core products, defined as products belonging to the firm's most important two-digit product category. The estimates in Table 4 are strongly consistent with this prediction. On the extensive margin, the cumulative-experience coefficient for core products is -0.00018, far smaller in magnitude than the corresponding baseline estimate, and the count-based coefficient is -0.00139, also much smaller than in the full sample. On the intensive margin, the sales effects are also substantially attenuated: the cumulative-experience coefficient is -0.0041 and the count-based coefficient is -0.0221, both much smaller than the corresponding estimates for the full sample. These estimates remain statistically significant, but their magnitudes indicate that firms are substantially less vulnerable to worker disappearance in product lines where knowledge is already embedded in routines, processes, and organizational capital. Put differently, worker-embodied and firm-embedded knowledge appear to be partial substitutes: where the latter is abundant, the former matters less for both expansion and ongoing performance. Consistent with this attenuation, the implied Share_T for core products is close to one on both margins (1.00 on the sales margin and 0.69 on the entry margin), indicating that essentially all of the small remaining targeted penalty scales with prior experience while the experience-invariant component is close to zero.

Finally, we examine whether reliance on worker-embodied knowledge varies with the overall skill intensity of the firm. Table 5 splits the sample into low- and high-skill firms, defined relative to within-sector and across-sector medians. Here the evidence points toward greater exposure among high-skill firms, although not uniformly across all cells. In the cumulative-experience specification, the entry coefficients are negative in all four splits and larger in magnitude for above-median firms both within sectors (-0.00075 versus -0.00070) and across sectors (-0.00086 versus -0.00065). On the sales margin, the coefficients are also more negative for above-median firms in both comparisons (-0.0156 versus -0.0050 within sectors; -0.0197 versus -0.0125 across sectors), although the below-median within-sector estimate is imprecise.

[16] An observationally equivalent reading is that the additive separability assumption underlying the decomposition in Equation (10) is itself less tenable for high-skill workers: when transferable knowledge T_{ijt} and match-specific capital W_{it} are complements rather than separable inputs, the experience-graded coefficient β_A absorbs part of the interaction, mechanically depressing the estimated Share_T even when the underlying portable component is large. We cannot distinguish these two interpretations with the static decomposition alone.

The count-based specification points in the same direction. On the entry margin, the coefficients are negative in all four splits and larger in magnitude for above-median firms within sectors (-0.00921 versus -0.00235) and across sectors (-0.01224 versus -0.00548). On the sales margin, the count-based effect is close to zero for below-median firms within sectors, but negative in the other three cells, with the largest magnitude appearing for above-median firms in the across-sector split (-0.1964). Taken together, these estimates suggest that high-skill firms may be more exposed to the loss of worker-embodied knowledge, but the evidence is not as uniformly strong as in the worker-education and core-product analyses. We therefore interpret the firm-skill results as suggestive rather than definitive.

Table 4: Heterogeneity by product centrality and worker skill

	Sales growth in product j			**Product entry in product j**		
	Core products	Low-skill	High-skill	Core products	Low-skill	High-skill
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Cumulative lost experience						
Lost experience in j	-0.0041***	-0.0135***	-0.0359***	-0.00018***	-0.00070***	-0.00133***
	(0.00117)	(0.00163)	(0.00740)	(0.00005)	(0.00010)	(0.00037)
Observations	56,935,292	56,903,732	54,487,288	56,935,292	56,903,732	54,487,288
Panel B. Number of exposed workers lost						
Workers lost with experience in j	-0.0221***	-0.0743***	-0.2858***	-0.00139***	-0.00617***	-0.02059***
	(0.00637)	(0.01153)	(0.05120)	(0.00035)	(0.00061)	(0.00257)
Observations	56,935,292	56,903,732	54,487,288	56,935,292	56,903,732	54,487,288
Transferable share, Share_T	1.00	0.98	0.62	0.694	0.613	0.317
Year fixed effects	✓	✓	✓	✓	✓	✓
Firm–product fixed effects	✓	✓	✓	✓	✓	✓
Lagged employment (log)	✓	✓	✓	✓	✓	✓

Notes: The table reports estimates from the static specification (9). The unit of observation is the firm–product–year. In columns (1)–(3), the dependent variable is the change in the inverse hyperbolic sine of product-j sales from year $t-1$ to t. In columns (4)–(6), it is an indicator for whether firm i starts producing product j in year t. In Panel A, the shock is the cumulative years of lost product-j experience among workers who died unexpectedly. In Panel B, the shock is the number of workers who died unexpectedly and had any experience in product j. Core products are products belonging to the firm's most important 2-digit product category. Low-skill and high-skill workers are defined by educational attainment below and at least at university level, respectively. All columns include lagged log employment. $\text{Share}^T \equiv \beta^A \cdot \bar{e}_{dj}/\beta^B$ is the descriptive share of the targeted loss that scales with prior product-specific experience, evaluated at the relevant subsample mean of prior product-j experience among deceased workers with $e_{dj} > 0$. Because Share_T is a ratio of two estimated coefficients, it is not informative when β_B is small or imprecisely estimated. We report Share_T only when β_B is statistically distinguishable from zero at the 1% level; cells failing either criterion are reported as "—". For core products on the sales margin, the displayed $\text{Share}_T = 1.00$ reflects that the experience-scaled component $\beta_A \cdot \bar{e}_{dj}$ is approximately equal in magnitude to the total targeted loss β_B, leaving a residual experience-invariant component close to zero. This is consistent with the framework's prediction that firms with deep firm-embedded knowledge S_{ijt} are largely insulated from match-specific losses on core product lines. Standard errors clustered at the firm–product level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5: Heterogeneity by firm skill intensity

	Sales growth in product j				**Product entry in product j**			
	Within-sector		Across-sector		Within-sector		Across-sector	
	Below	Above	Below	Above	Below	Above	Below	Above
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Cumulative lost experience								
Lost experience in j	-0.0050	-0.0156***	-0.0125***	-0.0197***	-0.00070***	-0.00075***	-0.00065***	-0.00086***
	(0.00623)	(0.00200)	(0.00177)	(0.00275)	(0.00022)	(0.00013)	(0.00012)	(0.00013)
Observations	26,773,400	28,938,416	31,017,168	24,284,368	26,773,400	28,938,416	31,017,168	24,284,368
Panel B. Number of exposed workers lost								
Workers lost with experience in j	0.0119	-0.1252***	-0.0638***	-0.1964***	-0.00235**	-0.00921***	-0.00548***	-0.01224***
	(0.02367)	(0.01381)	(0.01206)	(0.02457)	(0.00104)	(0.00113)	(0.00066)	(0.00120)
Observations	26,773,400	28,938,416	31,017,168	24,284,368	26,773,400	28,938,416	31,017,168	24,284,368
Transferable share, Share_T	—	0.694	0.939	0.563	—	0.456	0.595	0.394
Year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Firm–product fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Lagged employment (log)	✓	✓	✓	✓	✓	✓	✓	✓

Notes: The table reports estimates from the static specification (9). The unit of observation is the firm–product–year. In columns (1)–(4), the dependent variable is the change in the inverse hyperbolic sine of product-j sales from year $t-1$ to t. In columns (5)–(8), it is an indicator for whether firm i starts producing product j in year t. In Panel A, the shock is the cumulative years of lost product-j experience among workers who died unexpectedly. In Panel B, the shock is the number of workers who died unexpectedly and had any experience in product j. Firms are classified as below or above the skill median either within 2-digit sectors or relative to the overall annual sample median across sectors. Standard errors clustered at the firm–product level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

6 Conclusion

This paper provides causal evidence that worker-embodied, product-specific knowledge is a binding constraint on firms' horizontal expansion. Using unexpected worker deaths as plausibly exogenous shocks to a firm's knowledge stock, and combining rich employment histories with a firm–product panel, we trace the consequences of human capital loss for product entry and sales. Our event-study estimates show that firms become less likely to enter new product lines when the deceased worker had prior experience in those products. The decline is immediate and becomes larger in the first post-event year before attenuating thereafter for the total effect, while the experience-graded component is concentrated at impact and attenuates more quickly.

Static estimates complement this dynamic evidence. Cumulative lost product-specific experience predicts lower subsequent product entry and substantially lower product-level sales, with the strongest and most robust effects appearing on the sales margin. A descriptive decomposition indicates that, on the sales margin, a large share of the loss scales with prior experience, consistent with an important portable-knowledge component. On the entry margin, the analogous split is less stable and should be interpreted more cautiously.

The heterogeneity analysis sharpens these conclusions along two particularly clear dimensions. First, the penalty is roughly twice as large for university-educated workers, consistent with more valuable embodied knowledge. Second, the penalty is strongly attenuated for core products, consistent with the buffering role of firm-embedded routines and codified knowledge. Evidence that high-skill firms are more fragile points in the same direction, but is more mixed across specifications.

Our framework connects to several streams in strategic management. The capability index C_{ijt} can be interpreted as a firm-specific resource in the sense of Barney (1991), with the key insight that this resource is partially embodied in workers and therefore mobile. The decomposition into transferable (T_{ijt}) and firm-embedded (S_{ijt}) knowledge parallels the distinction in Teece et al. (1997) between dynamic capabilities that enable reconfiguration and operational capabilities embedded in routines. Our empirical setting allows us to isolate the worker-embodied component and quantify its contribution to the firm's ability to sense and seize new product opportunities. The finding that T_{ijt} accounts for the majority of targeted capability loss suggests that competitive advantage in product expansion is more fragile, and more dependent on worker retention, than theories emphasizing firm-embedded routines would predict.

These findings carry implications for both managers and policymakers. For managers, our estimates identify three concrete levers for mitigating vulnerability to human capital shocks, each anchored to a specific result in the paper. First, codification of tacit worker knowledge into organizational routines pays off: our finding that core products are largely insulated from worker death — with the cumulative-experience coefficient an order of magnitude smaller than for non-core products — shows that firms with deep, distributed firm-embedded knowledge effectively buffer themselves against individual departures. Investments in documentation, process standardization, and cross-functional handover protocols substitute for the portable component of worker-embodied capability and are most valuable in product lines where firm-embedded knowledge is currently shallow. Second, retention incentives matter most where the embodied knowledge is largest. Our finding that the targeted penalty is roughly twice as large for university-educated workers implies that the marginal value of preserving a high-skill worker's knowledge is substantially higher than for a comparable lower-skill worker; deferred compensation, firm-specific skill development, and internal promotion ladders are correspondingly higher-return for this group. Third, our finding that high-skill firms appear more exposed to worker disappearance suggests that knowledge redundancy — distributing product-specific expertise across multiple workers rather than concentrating it in single individuals — is a complementary lever, particularly for firms whose production relies heavily on worker-embodied knowledge.

These margins can be evaluated on expected-value grounds. The on-impact estimates imply that the unexpected loss of a worker with average prior product-specific experience reduces the probability of entering the relevant product line by roughly 0.7 percentage points. For a product line with an expected net entry value of $10 million, this corresponds to an expected loss of approximately $70,000 per exposed worker. Codification, retention, and redundancy investments at or below this threshold are justified on expected-value grounds for an average exposed worker; for high-skill workers and high-value product lines, the benchmark is larger still.

At the macro level, our findings suggest that shocks to the distribution of worker knowledge, whether from mortality, retirement, poaching, or policy-induced labor market disruptions, can affect the pace and direction of product reallocation across firms and, in turn, aggregate productivity growth. Policies that facilitate knowledge transfer, such as apprenticeship programs, cross-training mandates, or documentation requirements, may mitigate these costs. Conversely, policies that increase labor market fluidity without attending to knowledge continuity may generate hidden productivity costs by disrupting firm-level capability accumulation.

Two caveats merit emphasis. First, our identification is based on worker exits rather than hires, so the results speak to the marginal product of retained knowledge, not to selection in recruiting. The symmetric implication, that firms "learn by hiring", is consistent with our findings but not directly tested. A companion paper (Anonymous (2025)) studies these forces. Second, our measure of knowledge, defined as years of prior product experience, is necessarily a proxy for the underlying stock of tacit expertise. To the extent that this measure introduces classical measurement error, our estimates are attenuated and should be interpreted as lower bounds on the true importance of worker-embodied knowledge.

Several directions for future research emerge from our analysis. First, subsequent work could quantify how quickly firms rebuild lost capabilities through replacement hiring, internal training, or complementary investments in documentation and process improvement. Our robustness checks suggest that replacement hiring does not fully offset knowledge losses in the short run, but the medium- and long-run dynamics of capability reconstruction remain an open question. Second, extending the analysis to service industries—where product-specific knowledge may take the form of client relationships, methodological expertise, or project management capabilities—would clarify the generality of our findings beyond manufacturing. Third, linking worker-level knowledge measures to firm-level innovation outcomes (patents, R&D productivity) would illuminate how human capital constraints shape not only horizontal expansion but also the direction of technological change.

In sum, this paper demonstrates that worker mobility, even when involuntary, shapes the boundaries of the firm. The loss of a single worker with product-specific experience reduces the probability of entering that product line by economically meaningful magnitudes. The effects persist beyond the immediate event window and are amplified for highly educated workers and for expansion into peripheral product lines; the evidence on high-skill firms points in the same direction but is weaker. These patterns underscore that firms' growth options are embodied not only in physical capital and organizational routines but also in the knowledge carried by their workers.

References

Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey M. Wooldridge, "When Should You Adjust Standard Errors for Clustering?," *The Quarterly Journal of Economics*, 2023, *138* (1).

Agarwal, Rajshree, Raj Echambadi, April M. Franco, and M. B. Sarkar, "Knowledge Transfer through Inheritance: Spin-out Generation, Development, and Survival," *Academy of Management Journal*, 2004, *47* (4), 501–522.

Almeida, Paul and Bruce Kogut, "Localization of Knowledge and the Mobility of Engineers in Regional Networks," *Management Science*, 1999, *45* (7), 905–917.

Andersen, Steffen and Kasper Meisner Nielsen, "Participation Constraints in the Stock Market: Evidence from Unexpected Inheritance Due to Sudden Death," *The Review of Financial Studies*, 2011, *24* (5), 1667–1697.

Anonymous, "Companion paper, citation withheld for blind review," Technical Report 2025.

Argote, Linda and Paul Ingram, "Knowledge Transfer: A Basis for Competitive Advantage in Firms," *Organizational Behavior and Human Decision Processes*, 2000, *82* (1), 150–169.

Azoulay, Pierre, Joshua S. Graff Zivin, and Jialan Wang, "Superstar Extinction," *Quarterly Journal of Economics*, 2010, *125* (2), 549–589.

Balke, Neele and Thibaut Lamadon, "Productivity Shocks, Long-Term Contracts, and Earnings Dynamics," *American Economic Review*, 2022, *112* (7), 2139–77.

Barney, Jay, "Firm Resources and Sustained Competitive Advantage," *Journal of Management*, 1991, *17* (1), 99–120.

Becker, Gary S., *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, Chicago: University of Chicago Press, 1964.

Becker, Sascha O and Hans K Hvide, "Entrepreneur Death and Startup Performance," *Review of Finance*, 05 2021, *26* (1), 163–185.

Bertheau, Antoine, Pierre Cahuc, Simon Jaeger, and Rune Vejlin, "Turnover Costs: Evidence from Unexpected Worker Separations," *Working Paper*, 2022.

Brooke, Hannah Louise, Mats Talbäck, Jesper Hörnblad, Lars Age Johansson, Jonas Filip Ludvigsson, Henrik Druid, Maria Feychting, and Rickard Ljung, "The Swedish cause of death register," *European Journal of Epidemiology*, 2017, *32* (9), 765–773.

Campbell, Benjamin A., Russell Coff, and David Kryscynski, "Rethinking Sustained Competitive Advantage from Human Capital," *Academy of Management Review*, 2012, *37* (3), 376–395.

Carlsson, Mikael and Oskar Nordstrom Skans, "Evaluating Microfoundations for Aggregate Price Rigidities: Evidence from Matched Firm-Level Data on Product Prices and Unit Labor Cost," *American Economic Review*, 2012, *102* (4), 1571–95.

Coff, Russell W., "Human Assets and Management Dilemmas: Coping with Hazards on the Road to Resource-Based Theory," *Academy of Management Review*, 1997, *22* (2), 374–402.

Gathmann, Christina and Uta Schönberg, "How General Is Human Capital? A Task-Based Approach," *Journal of Labor Economics*, 2010, *28* (1), 1–49.

Gibbons, Robert and Michael Waldman, "Task-Specific Human Capital," *American Economic Review Papers and Proceedings*, 2004, *94* (2), 203–207.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Hatch, Nile W. and Jeffrey H. Dyer, "Human Capital and Learning as a Source of Sustainable Competitive Advantage," *Strategic Management Journal*, 2004, *25* (12), 1155–1178.

Hoey, Sam, Thomas Peeters, and Jan C. van Ours, "The impact of absent co-workers on productivity in teams," *Labour Economics*, 2023, *83*, 102400.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan, "Earnings Losses of Displaced Workers," *American Economic Review*, 1993, *83* (4), 685–709.

Jäger, Simon, Jörg Heining, and Nathan Lazarus, "How Substitutable Are Workers? Evidence from Worker Deaths," *American Economic Review*, forthcoming. forthcoming.

Jaravel, Xavier, Neviana Petkova, and Alex Bell, "Team-Specific Capital and Innovation," *American Economic Review*, 2018, *108* (4-5), 1034–1073.

Jones, Benjamin and Benjamin A. Olken, "Do Leaders Matter? National Leadership and Growth Since World War II," *The Quarterly Journal of Economics*, 2005, *120* (3), 835–864.

Klepper, Steven, "Disagreements, Spinoffs, and the Evolution of Detroit as the Capital of the U.S. Automobile Industry," *Management Science*, 2007, *53* (4), 616–631.

_ **and Sally Sleeper**, "Entry by Spinoffs," *Management Science*, 2005, *51* (8), 1291–1306.

Lazear, Edward P., "Firm-Specific Human Capital: A Skill-Weights Approach," *Journal of Political Economy*, 2009, *117* (5), 914–940.

Lim, Zaneta, Karen Gibbs, James E. Potts, and Shubhayan Sanatani, "A review of sudden unexpected death in the young in British Columbia," *Canadian Journal of Cardiology*, 2010, *26* (1), 22–26.

Neffke, Frank and Martin Henning, "Skill Relatedness and Firm Diversification," *Strategic Management Journal*, 2013, *34* (3), 297–316.

_ , _ , and Ron Boschma, "How Do Regions Diversify over Time? Industry Relatedness and the Development of New Growth Paths in Regions," *Economic Geography*, 2011, *87* (3), 237–265.

Patault, Béréngère and Clémence Lenoir, "Customer Capital Spillovers: Evidence from Sales Managers in International Markets," *American Economic Journal: Applied Economics*, October 2024, *16* (4), 404–443.

Poletaev, Maxim and Chris Robinson, "Human Capital Specificity: Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984–2000," *Journal of Labor Economics*, 2008, *26* (3), 387–420.

Saez, Emmanuel, Benjamin Schoefer, and David Seim, "Payroll Taxes, Firm Behavior, and Rent Sharing: Evidence from a Young Workers' Tax Cut in Sweden," *American Economic Review*, 2019, *109* (5), 1717–63.

Sauvagnat, Julien and Fabiano Schivardi, "Are Executives in Short Supply? Evidence from Death Events," *The Review of Economic Studies*, 02 2023.

Singh, Jasjit and Ajay Agrawal, "Recruiting for Ideas: How Firms Exploit the Prior Inventions of New Hires," *Management Science*, 2011, *57* (1), 129–150.

Stoyanov, Andrey and Nikolay Zubanov, "Productivity Spillovers across Firms through Worker Mobility," *American Economic Journal: Applied Economics*, April 2012, *4* (2), 168–98.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Teece, David J., Gary Pisano, and Amy Shuen, "Dynamic Capabilities and Strategic Management," *Strategic Management Journal*, 1997, *18* (7), 509–533.

Waberi, Abdourahman A., "A Missive for Youth," *The UNESCO Courier*, 2018, (3).

Appendix A Online Appendix

The following online appendix assembles details for the empirical evidence.

A.1 Data Overview

Table A.1: Overview of datasets and key variables

Dataset	**Period**	**Key variables**
LISA (worker data)	1990–2019	Worker ID, firm ID, age, gender, education, graduation year, occupation, labor income
FDB (firm data)	1997–2019	Firm ID, employment, value added, output, capital, industry
IVP (product data)	1997–2019	Firm ID, year, product code, production value, quantity
Cause-of-death registry	1997–2016	Worker ID, year, cause of death

Notes: The table summarizes the four administrative datasets used in the analysis. LISA = *Longitudinell Integrationsdatabas för Sjukförsäkrings och Arbetsmarknadsstudier*; FDB = *Företagsdatabasen*; IVP = *Industrins Varuproduktion*. The table lists the main variables used in the paper rather than the full set of variables available in each source.

A.2 Additional Empirical Evidence

Figure A.1: Event study of worker death on labor productivity.

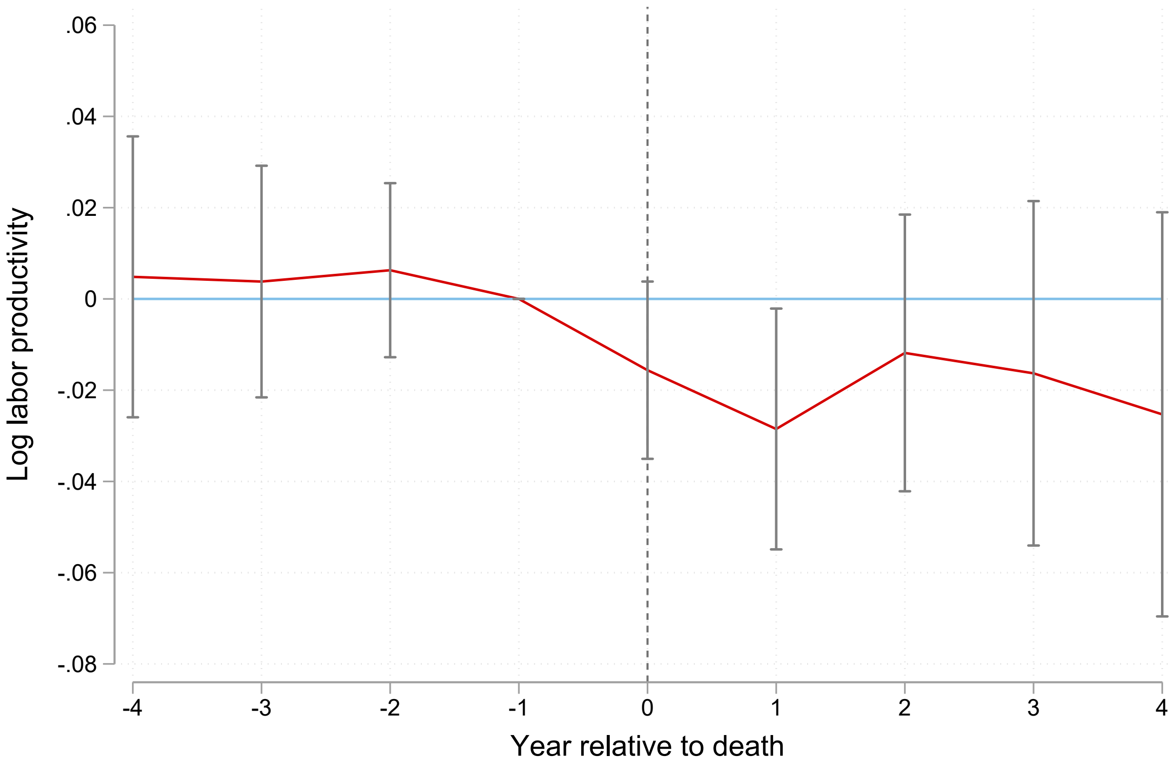

Notes: The figure shows the impact of a worker death on firm log labor productivity around the time of a worker death. The corresponding regression specification is $\log(Y/L)_{jt} = \sum_{\tau=t-4}^{t+4} \alpha_\tau \times I_{\text{period}_\tau} + \gamma_j + \eta_t + \varepsilon_{jt}$.

Figure A.2: Event study of death on product expansion

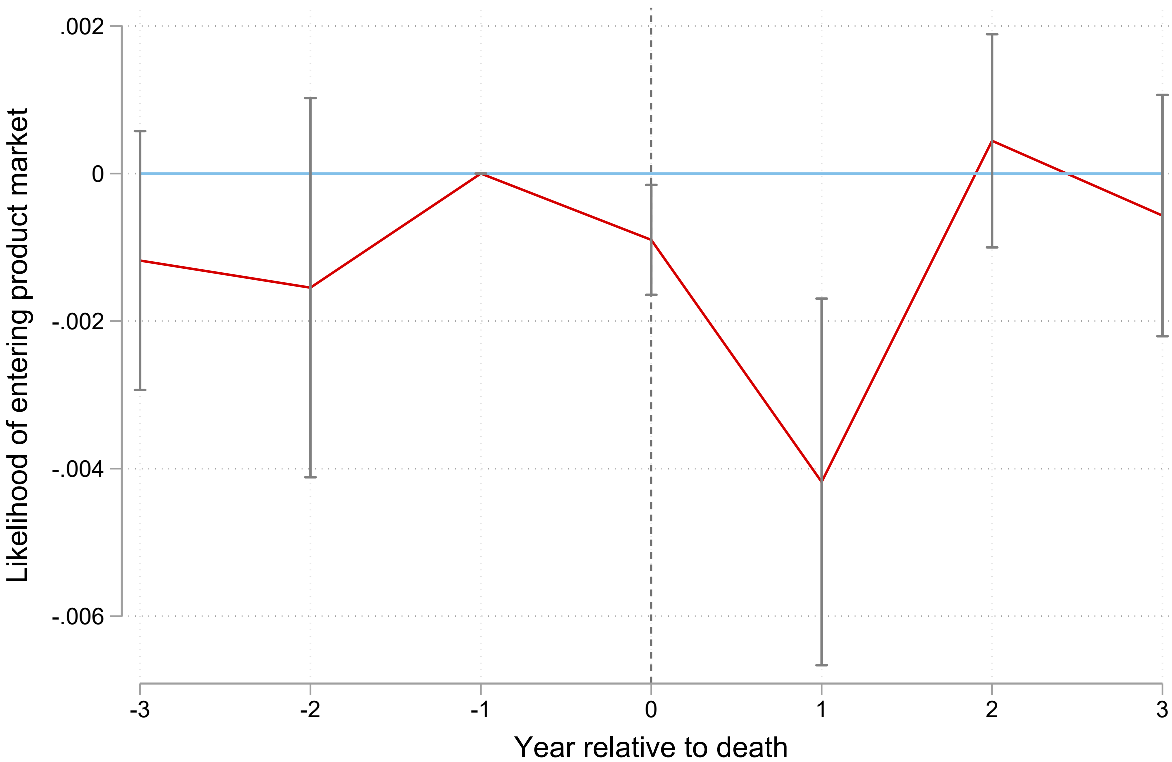

Notes: The figure shows the impact of a worker death on whether a firm starts producing a certain product when the worker that died had experience in producing exactly that product. The corresponding regression specification is equation (8). The figure plots coefficients δ_k. In this figure we include only events in which firms do not replace the disappearing workers, as measured through changes in firm size.

Table A.2: Robustness to controls for dying workers' characteristics

	Sales growth in product j		**Product entry in product j**	
	Without controls (1)	With controls (2)	Without controls (3)	With controls (4)
Panel A. Cumulative lost experience				
Lost experience in j	-0.0134***	-0.0135***	-0.00067***	-0.00062***
	(0.00153)	(0.00178)	(0.00009)	(0.00013)
Observations	56,935,292	56,935,281	56,935,292	56,935,281
Panel B. Number of exposed workers lost				
Workers lost with experience in j	-0.0780***	-0.0848***	-0.00618***	-0.00694***
	(0.01107)	(0.01335)	(0.00060)	(0.00058)
Observations	56,935,292	56,935,281	56,935,292	56,935,281
Year fixed effects	✓	✓	✓	✓
Firm–product fixed effects	✓	✓	✓	✓
Lagged employment (log)	✓	✓	✓	✓
Worker controls		✓		✓

Notes: The table reports estimates from the static specification (9). The unit of observation is the firm–product–year. In columns (1)–(2), the dependent variable is the change in the inverse hyperbolic sine of product-j sales from year $t-1$ to t. In columns (3)–(4), it is an indicator for whether firm i starts producing product j in year t. In Panel A, the shock is the cumulative years of lost product-j experience among workers who died unexpectedly. In Panel B, the shock is the number of workers who died unexpectedly and had any experience in product j. All columns include lagged log employment. Columns (2) and (4) additionally include controls for the dying worker: five age-group dummies, gender, tenure, income, and 3-digit occupation fixed effects. Standard errors clustered at the firm–product level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.3: Potential spillovers across products within the same firm

	Sales growth in product j		Product entry in product j	
	All products in firm	Non-exposed products only	All products in firm	Non-exposed products only
	(1)	(2)	(3)	(4)
Cumulative lost experience				
Lost experience in j	-0.000066**	0.000134***	-0.0000080***	0.0000120***
	(0.000032)	(0.000024)	(0.0000020)	(0.0000020)
Observations	56,935,292	56,935,292	56,935,292	56,935,292
Year fixed effects	✓	✓	✓	✓
Firm–product fixed effects	✓	✓	✓	✓
Lagged employment (log)	✓	✓	✓	✓

Notes: The table reports estimates from the static specification (9). The unit of observation is the firm–product–year. In columns (1)–(2), the dependent variable is the change in the inverse hyperbolic sine of product-j sales from year $t-1$ to t. In columns (3)–(4), it is an indicator for whether firm i starts producing product j in year t. The shock variable is based on the cumulative years of product-j experience among workers who died unexpectedly. In columns (1) and (3), the experience of the deceased worker is assigned to all products within the firm. In columns (2) and (4), the experience of the deceased worker is assigned only to products in which the worker had no prior experience, while exposure is set to zero for the directly exposed product lines. All columns include year fixed effects, firm–product fixed effects, and lagged log employment. Standard errors clustered at the firm–product level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.4: Production, product entry, and worker experience: Two-digit product categories

	Sales growth in product j			**Product entry in product j**		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Cumulative lost experience						
Lost experience in j	-0.0060***	-0.0060***	-0.0060***	-0.00036***	-0.00036***	-0.00036***
	(0.00172)	(0.00172)	(0.00172)	(0.00011)	(0.00011)	(0.00011)
Observations	5,195,616	5,195,616	5,195,616	5,195,616	5,195,616	5,195,616
Panel B. Number of exposed workers lost						
Workers lost with experience in j	-0.0090	-0.0091	-0.0090	-0.00281***	-0.00281***	-0.00281***
	(0.02267)	(0.02267)	(0.02267)	(0.00097)	(0.00097)	(0.00097)
Observations	5,195,616	5,195,616	5,195,616	5,195,616	5,195,616	5,195,616
Transferable share, Share_T	—	—	—	0.72	0.72	0.72
Year fixed effects	✓	✓	✓	✓	✓	✓
Firm–product fixed effects	✓	✓	✓	✓	✓	✓
Current employment (log)		✓			✓	
Lagged employment (log)			✓			✓

Notes: The table reports estimates from the static specification (9) using 2-digit product categories. The unit of observation is the firm–product–year. In columns (1)–(3), the dependent variable is the change in the inverse hyperbolic sine of product-j sales from year $t-1$ to t. In columns (4)–(6), it is an indicator for whether firm i starts producing product j in year t. In Panel A, the shock is the cumulative years of lost product-j experience among workers who died unexpectedly. In Panel B, the shock is the number of workers who died unexpectedly and had any experience in product j. Columns (2) and (5) add contemporaneous log employment; columns (3) and (6) add lagged log employment. $Share_T \equiv \beta_A \cdot \bar{e}_{dj}/\beta_B$ is the descriptive share of the targeted loss that scales with prior product-specific experience, evaluated at $\bar{e}_{dj} = 5.33$. Because Share_T is a ratio of two estimated coefficients, it is not informative when β_B is small or imprecisely estimated. We report Share_T only when β_B is statistically distinguishable from zero at the 1% level; cells failing either criterion are reported as "—". See Section 5.2 for discussion of how the 2-digit entry-margin Share_T compares to the 4-digit baseline. Standard errors clustered at the firm–product level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.