Economic Growth through Worker Reallocation: The Role of Knowledge Spillovers *

Eero Mäkynen†

June 23, 2026

Abstract

This paper studies how knowledge diffusion through worker reallocation contributes to aggregate productivity growth. Using Finnish matched employer–employee data, I estimate a positive productivity response to hires from more productive previous employers. I then use this estimate to discipline an industry-equilibrium model in which hires from more productive previous employers create learning opportunities and producer dynamics determine how often those opportunities arise. The calibrated model implies that knowledge diffusion through worker reallocation makes an important contribution to aggregate growth, both through direct learning and through equilibrium changes in selection, entry, and employment allocation across producers. Policy counterfactuals show that the knowledge-diffusion mechanism amplifies the adverse growth effects of firing costs, while a size-dependent hiring credit shifts employment toward establishments whose workers carry more useful knowledge when reallocated, but the learning gains are too small to offset the misallocation it creates.

Keywords: *knowledge diffusion, firm dynamics, worker reallocation, economic growth*
JEL Classification: *D24, E23, E24, J63, O33, O40*

1 Introduction

Economic growth depends not only on the creation of new ideas, but also on how those ideas reach producers below the frontier.[1] Akcigit and Ates (2023) link declining U.S. business dynamism to weaker idea flows from frontier to laggard firms and point to declining worker flows as one possible contributor to that weakening. This interpretation is natural in light of empirical evidence that hiring producers become more productive after recruiting workers from productive or knowledge-rich previous employers.[2] Taken together, these findings motivate treating producer dynamics as part of knowledge diffusion: worker separations and the exit of producers determine whose knowledge enters worker flows, while hiring choices determine which establishments can learn from that knowledge. Aggregate growth and the effects of labor-market policy may therefore depend not only on the reallocation of labor across producers, but also on how producer dynamics shape the knowledge carried by that reallocation.[3]

*I am indebted to Martin Ellison, William Kerr, Petr Sedláček, Vincent Sterk and Oskari Vähämaa for detailed feedback. I would also like to thank Guido Ascari, Vladimir Asriyan, Vasco Carvalho, Jesús Fernández-Villaverde, Shigeru Fujita, Sampreet Goraya, Axel Gottfries, John Haltiwanger, Heikki Kauppi, Jouko Vilmunen, Ludo Visschers, and seminar and conference participants at the University of Oxford, University of Turku, Helsinki GSE and Nordic Summer Symposium in Macroeconomics 2022 for insightful discussions and suggestions. First version: May 2021. Declarations of interest: none.

†University of Turku, FI-20014 Turun yliopisto, Finland (email: eeanma@utu.fi)

[1]Diffusion models make this transmission channel explicit through meetings, costly search, and imitation from more productive producers (Lucas & Moll, 2014; Perla & Tonetti, 2014), trade exposure (Buera & Oberfield, 2020), and imperfect technology spillovers in innovation (Jo & Kim, 2026).

[2]This evidence comes from empirical work on learning by hiring and knowledge flows through worker reallocation, including hires from more productive firms (Serafinelli, 2019; Stoyanov & Zubanov, 2012) and hires of knowledge carriers (Parrotta & Pozzoli, 2012).

[3]This policy question connects to Hopenhayn and Rogerson (1993), who study firm-level job turnover and reallocation, and to growth models with labor-market frictions, including Poschke (2009), Mukoyama and Osotimehin (2019), and Petit (2023).

This paper combines producer-level evidence on learning by hiring with a calibrated model to show that knowledge diffusion through worker reallocation matters quantitatively for aggregate growth and policy evaluation. Using Finnish matched employer–employee data, I estimate that establishments become more productive after hiring from more productive previous employers. The quantitative model asks how this estimated productivity response aggregates once producer dynamics determine both the supply of worker-carried knowledge and the set of producers exposed to it. Removing this source of diffusion lowers annual growth from 2.56% to 2.01%, with the loss coming not only from fewer direct learning events but also from changes in selection, entry, and the allocation of employment across producers. It also changes policy evaluation: firing costs reduce growth more when they reduce knowledge-transmitting worker flows, while a hiring credit for large establishments shifts employment toward establishments whose workers carry more useful knowledge when reallocated, but the learning gains are too small to offset the misallocation it creates.

In Finnish matched employer–employee data, one additional hire from a more productive previous employer is associated with about 0.56% higher future productivity at the receiving establishment. I estimate this response using an augmented control-function production estimator in which the number of hires from more productive establishments enters the controlled productivity process. The spillover coefficient is therefore identified jointly with the production-function parameters and establishment productivity rankings. The positive response is robust across alternative specifications. It is also largest where learning from previous employers should matter more: among below-mean receivers and among hires from establishments more than 10 percent more productive than the receiver. These patterns motivate the model's learning mechanism, which treats hires from more productive previous employers as opportunities for receiving establishments to learn.

To evaluate the general-equilibrium impact of worker-transmitted knowledge, I introduce the mechanism into a growth version of the industry-equilibrium framework of Hopenhayn and Rogerson (1993). The economy includes the ingredients needed for knowledge diffusion through worker reallocation to operate in equilibrium. Productivity dispersion creates source–receiver gaps, and producer dynamics determine whose knowledge is carried by moving workers. A recruit from a more productive previous employer can then create a productivity-growth opportunity for the receiving establishment. Hiring and separations are therefore not only labor-adjustment decisions. They also determine how often establishments draw knowledge from more productive previous employers and how much useful knowledge those worker flows contain. Growth exists even without this channel through productivity shocks, selection, and entrant imitation, but worker-transmitted knowledge adds a separate growth force whose strength depends on establishment decisions.

Calibrated to the spillover estimate, aggregate growth, and moments of producer dynamics, the model implies that knowledge diffusion through worker reallocation is a sizable source of growth. Shutting down these knowledge transfers lowers annual growth from 2.56% to 2.01%. The associated welfare loss, computed from the full equilibrium allocation of consumption and labor, is equivalent to a permanent consumption decline of about five percent. Direct learning events among continuing establishments account for 0.17 percentage points of productivity growth. The remaining loss comes from equilibrium producer dynamics because the spillover process affects which establishments survive, how large they become, and whose workers later appear in the hiring pool.

The counterfactuals separate two implications of knowledge diffusion through worker reallocation: mobility restrictions become more costly, while subsidizing employment growth at large establishments faces a misallocation tradeoff. These implications are intuitive but not automatic, because the market equilibrium is shaped by two growth-relevant feedbacks that individual establishments do not take into account. Entrant imitation makes today's incumbent distribution matter for future entrants, and knowledge spillovers through worker reallocation make employment and separation decisions affect the knowledge available to future hiring establishments. Halving exogenous worker separations saves replacement costs in an economy without worker-transmitted knowledge, but in the benchmark economy it lowers growth by 0.14 percentage points and requires 1.88% consumption compensation by removing learning opportunities. A one-year-wage firing cost also reduces knowledge-transmitting worker flows, lowering growth by 0.11 percentage points with spillovers compared with 0.02 percentage points in a recalibrated economy without spillovers. The hiring-credit experiment asks whether a size-dependent subsidy can be justified by shifting employment toward establishments whose workers carry more useful knowledge when reallocated. The credit does move employment in that direction and modestly raises aggregate exposure to hires from more productive previous employers, but the resulting learning gains do not overcome the induced misallocation.

The paper builds on models of labor-market policy and producer dynamics. In Hopenhayn and Rogerson (1993), labor-market policy changes job turnover and the allocation of activity across heterogeneous producers. Growth versions of this environment, including Poschke (2009), Mukoyama and Osotimehin (2019), and Petit (2023), connect labor-market frictions to productivity growth through selection, reallocation, or endogenous productivity investment. Closely related on the labor-market side, Bradley and Gottfries (2026) embed technological diffusion in a search-and-matching model of job creation, separations, and growth. A broader mobility-cost mechanism appears in Raurich, Sánchez-Losada, and Vilalta-Bufí (2015), where knowledge is embodied in sector-specific worker types and diffuses across sectors through labor mobility. Separately, Gabler and Poschke (2013) show that distortions can affect aggregate productivity by changing firms' experimentation and productivity evolution. I add measured productivity spillovers from hires with more productive previous employers to this type of environment. Worker flows determine not only where labor is employed, but also which producers' knowledge becomes available to future hiring establishments. This changes policy evaluation because separations and hires affect both current production and future learning opportunities.

The paper also connects to macroeconomic models of knowledge diffusion and idea flows. In modern idea-flow models, aggregate growth depends on how less productive producers gain access to knowledge created elsewhere (Buera & Lucas, 2018). Some work emphasizes knowledge embodied in people and learned through interactions (Lucas, 2009; Lucas & Moll, 2014); other work studies costly search and imitation from more productive producers (Perla &

Tonetti, 2014), trade and the cross-country diffusion of ideas (Buera & Oberfield, 2020), trade-induced technology adoption through export opportunities and foreign competition (Perla, Tonetti, & Waugh, 2021), or imperfect technology spillovers that shape innovation incentives (Akcigit & Ates, 2023; Jo & Kim, 2026). I replace the contact margin in idea-flow models with a measured worker-reallocation margin. A learning opportunity arises when a hire comes from a more productive previous employer. A related labor-mobility mechanism appears in Heggedal, Moen, and Preugschat (2017), who study productivity spillovers through labor mobility in search equilibrium. The distinction here is to discipline the knowledge-diffusion process with matched employer–employee evidence and embed it in a quantitative model of heterogeneous producers.

Empirically, the paper takes evidence on learning by hiring to an aggregate model. Stoyanov and Zubanov (2012) and Serafinelli (2019) show that receiving producers become more productive after hiring workers from more productive or higher-quality previous employers, while Parrotta and Pozzoli (2012) documents productivity gains from hiring knowledge-carrying workers. I estimate a related receiving-establishment productivity response in Finnish data and use it to discipline the model, rather than treating the estimated productivity response itself as the aggregate value of knowledge diffusion through worker reallocation. The estimator builds on the control-function approach of Ackerberg, Caves, and Frazer (2015). Its controlled-productivity-process logic follows work that lets productivity-shifting actions enter the productivity law, including R&D in Doraszelski and Jaumandreu (2013) and export experience in De Loecker (2013). Closest to the spillover setting, Malikov and Zhao (2023) emphasize that productivity-modifying spillovers should enter productivity estimation rather than be tested only in a second-stage regression on recovered productivity.

Throughout, I focus on the productivity of the receiving establishment, not the wage growth of the moving worker, worker-specific human-capital accumulation, or the creation of a new establishment.[4] With this boundary, the remaining sections follow receiving-establishment productivity responses to hires from more productive previous employers from measurement to aggregation. Section 2 estimates how productivity responds to those hires. Section 3 models how producer dynamics determine which previous employers are represented in worker flows. Section 4 quantifies how growth and policy counterfactuals change when worker flows also move productive knowledge.

2 Empirical Evidence on Learning by Hiring

In Finnish matched employer–employee data, one additional hire from a more productive previous employer is associated with 0.56 percent higher future productivity at the receiving establishment. I estimate this response with an augmented control-function production estimator that allows hiring from more productive previous employers to enter the controlled pro-

[4]Related worker-side and entrant margins include coworker learning and wage growth (Jarosch, Oberfield, & Rossi-Hansberg, 2021), firms as learning environments (Gregory, 2026), labor-market fluidity and human-capital accumulation (Engbom, 2022), noncompete contracts and worker mobility (Shi, 2023), spinouts or employee entrepreneurship (Baslandze & Vardishvili, 2026; Sohail, 2021), and ex ante startup heterogeneity in post-entry growth (Sterk, Sedláček, & Pugsley, 2021).

ductivity process, following approaches used for R&D, export experience, and peer productivity (De Loecker, 2013; Doraszelski & Jaumandreu, 2013; Malikov & Zhao, 2023). This setting requires the productivity proxy both to address input simultaneity and to rank establishments, so the measure of hires from more productive previous employers and its coefficient are recovered together. The estimated response is positive across alternative specifications and is larger for below-mean receiving establishments and for hires from previous employers more than 10 percent more productive than the receiver. These patterns motivate the worker-reallocation learning mechanism in the quantitative model.

2.1 Incumbents' Problem with Learning from Hiring

A simple incumbent problem clarifies why the productivity law matters for measuring learning from hires from more productive previous employers. The example removes exit, endogenous separations, and the investment choice so that the hiring first-order condition isolates two benefits of hiring: expanding future employment and, when recruits come from more productive previous employers, shifting future productivity. The exercise motivates why the empirical measure counts hires from more productive previous employers rather than total hiring, and why its coefficient belongs inside the productivity law.

In each period t, an incumbent establishment chooses hiring $h_t \geq 0$ at the end of the period, and employment evolves according to

$$n_{t+1} = (1-\delta)n_t + h_t, \qquad \delta \in (0,1), \tag{1}$$

where δ is an exogenous separation rate. Past investment decisions predetermine capital k_t and govern its evolution according to an exogenous law of motion.

Let z_t denote establishment productivity. The value of an incumbent with state (z_t, n_t, k_t) is

$$V(z_t, n_t, k_t) = \max_{h_t \geq 0} \Big\{ \pi(z_t, n_t, k_t) - C(h_t) + \beta\, \mathbb{E}\big[V(z_{t+1}, n_{t+1}, k_{t+1})\big] \Big\}, \tag{2}$$

where $\pi(\cdot)$ is operating profit and $C(\cdot)$ is the hiring cost. Productivity follows a controlled first-order Markov process,

$$z_{t+1} = \Gamma(z_t, F_t^z; h_t) + \vartheta_{t+1}, \tag{3}$$

where F_t^z denotes the cross-sectional productivity distribution that determines which previous employers are more productive than the hiring establishment.

The learning mechanism enters through the state-contingent composition of hires. For the conceptual problem, write

$$\Gamma(z_t, F_t^z; h_t) = \Upsilon(z_t) + \beta_z \psi^+(z_t; F_t^z)\, h_t, \tag{4}$$

where $\psi^+(z_t; F_t^z)$ is the expected share of hires whose previous employer has productivity above z_t, taken as given by the incumbent when choosing h_t. The term $h_t^+ \equiv \psi^+(z_t; F_t^z) h_t$ is the conceptual measure of hires from more productive previous employers.

For an interior hiring choice, the first-order condition shows the two margins that the empirical design must separate:

$$C'(h_t) = \beta\,\mathbb{E}\big[V_n(z_{t+1}, n_{t+1}, k_{t+1})\big] + \beta\,\mathbb{E}\big[V_z(z_{t+1}, n_{t+1}, k_{t+1})\big]\,\beta_z\,\psi^+(z_t; F_t^z). \tag{5}$$

The first-order condition separates the standard employment-capacity value of hiring from the productivity value of hires from more productive previous employers. The second term appears only because some hires arrive from more productive previous employers and can raise future productivity. Estimating this response therefore requires ranking new hires' previous employers relative to the receiver and estimating the coefficient on hires from more productive previous employers inside the productivity process.

2.2 Estimation Strategy: Augmented Control-Function Approach

The estimation problem is that hiring from more productive previous employers is defined by relative productivity, while its coefficient belongs in the productivity process used to recover productivity. The materials-demand inversion in a control-function estimator produces a productivity proxy that can both address input simultaneity and rank establishments by productivity. I therefore extend the ACF control-function framework to include hiring from more productive previous employers in the productivity process.[5]

For the empirical implementation, I use the value-added Cobb–Douglas production function in Ackerberg et al. (2015):

$$\ln y_{it} = \beta_0 + \beta_n \ln n_{it} + \beta_k \ln k_{it} + z_{it} + \varepsilon_{it}, \tag{6}$$

where n_{it} is predetermined labor input measured by the wage bill, k_{it} is predetermined capital, z_{it} is Hicks-neutral productivity, and ε_{it} is a transitory shock that does not affect input choices.[6]

Productivity evolves as a controlled process in which hiring from more productive previous employers can shift future productivity:

$$z_{it} = \Upsilon(z_{i,t-1}) + \beta_z h^+_{i,t-1} + \vartheta_{it}. \tag{7}$$

Here ϑ_{it} is the productivity innovation realized between $t-1$ and t.[7]

[5] The control-function approach used here builds on Olley and Pakes (1996), Levinsohn and Petrin (2003), and Ackerberg et al. (2015). The extension is closest to Malikov and Zhao (2023), who allow productivity to depend on peer productivity within a controlled process. A difference is that I use the nonparametric ACF proxy rather than a parametrized materials first-order condition to recover productivity. This is computationally more demanding, but it keeps the estimator close to the ACF timing and proxy structure.

[6] The timing of input choices follows the logic of Assumption 3C in Ackerberg et al. (2015), which treats capital as predetermined and allows labor to have dynamic implications. In particular, establishments accumulate capital and labor according to $\ln(k_{it}) = \mathcal{K}(k_{i,t-1}, i_{i,t-1})$ and $\ln(n_{it}) = \iota(n_{i,t-1}, h_{i,t-1}, s_{i,t-1})$, where investment $i_{i,t-1}$, hiring $h_{i,t-1}$, and separations $s_{i,t-1}$ are chosen in period $t-1$. I specialize to this $t-1$ labor-choice timing to align the empirical timing with the quantitative model. The transitory shock assumption follows Assumption 1 in Ackerberg et al. (2015): $\mathbb{E}[\varepsilon_{it} \mid I_{it}] = 0$.

[7] This corresponds to Assumption 2 in Ackerberg et al. (2015), modified to allow a controlled productivity process. The productivity modifier is lagged R&D expenditures in Doraszelski and Jaumandreu (2013), export experience in De Loecker (2013), and average peer productivity in Malikov and Zhao (2023). In the present setting,

The empirical measure is a weighted count of hires from more productive previous employers. Let $H_{i,t-1}$ denote workers hired by establishment i between $t-1$ and t, and let $j(\ell, t-1)$ be worker ℓ's previous employer. The maintained timing is that these hiring choices and previous-employer links are fixed before the period-t productivity innovation is realized. The baseline measure is

$$h^{+}_{i,t-1} = \sum_{\ell \in H_{i,t-1}} v_{\ell,t-1} \mathbf{1}\{z_{j(\ell,t-1),t-1} > z_{i,t-1}\}. \tag{8}$$

The weight $v_{\ell,t-1}$ is a normalized residual-wage correction described in Appendix A. The correction keeps the exposure measure focused on worker-transmitted knowledge rather than on observable differences in the wage-based quality of recruits. Setting $v_{\ell,t-1} = 1$ gives the raw count of hires from more productive previous employers, which I report as a robustness check.

Materials enter the estimator only through their role as a proxy for unobserved productivity. Under the ACF monotonicity condition, intermediate-input demand can be inverted, making the first-stage object the composite output component

$$\Phi_{it} \equiv \beta_0 + \beta_n \ln n_{it} + \beta_k \ln k_{it} + f^{-1}(k_{it}, n_{it}, m_{it}),$$

a flexible function of predetermined inputs and materials.[8] I approximate this first-stage object by

$$\widehat{\Phi}_{it} = g(\ln k_{it}, \ln n_{it}, \ln m_{it}; \gamma),$$

where $g(\cdot;\gamma)$ is a finite-order polynomial and the fitted object $\widehat{\Phi}_{it}$ still includes the input contributions. For a candidate parameter vector $\boldsymbol{\beta} \equiv (\beta_0, \beta_n, \beta_k)$, subtracting those contributions gives the productivity proxy

$$\hat{z}_{it}(\boldsymbol{\beta}) = \widehat{\Phi}_{it} - \beta_0 - \beta_n \ln n_{it} - \beta_k \ln k_{it}. \tag{9}$$

This proxy is used to rank establishments when constructing $h^{+}_{i,t-1}$ and to evaluate lagged productivity in the controlled productivity process.

Substituting the controlled productivity process into the production function gives the

the modifier is hiring from more productive previous employers, so the productivity transition can be written as $p(z_{i,t+1} \mid I_{it}) = p(z_{i,t+1} \mid z_{it}, F^z_t, h^+_{it})$.

[8]Formally, the conditional intermediate-input demand in Assumption 4C of Ackerberg et al. (2015) can be written as $m_{it} = f(k_{it}, n_{it}, z_{it})$, and the monotonicity condition in Assumption 5C requires this demand to be strictly increasing in z_{it}. The demand function can therefore be inverted to write $z_{it} = f^{-1}(k_{it}, n_{it}, m_{it})$. Because materials are nondynamic in the ACF timing, this inversion provides the productivity proxy used to rank establishments for a candidate production function. The following Leontief formulation is one sufficient way to rationalize using value added while materials provide the proxy:

$$y_{it} = \exp(\varepsilon_{it}) \min\left\{\exp(\beta_0 + z_{it}) k_{it}^{\beta_k} n_{it}^{\beta_n},\ \beta_m m_{it}\right\},$$

for some $\beta_m > 0$. See Ackerberg et al. (2015, Sections 4.1 and 5). The control-function argument itself rests on the conditional demand and monotonicity assumptions rather than on this particular example.

second-stage estimating equation

$$
\begin{aligned}
\ln y_{it} = \beta_0 + \beta_n \ln n_{it} + \beta_k \ln k_{it} + \Upsilon(\hat{z}_{i,t-1}(\boldsymbol{\beta})) \\
+ \beta_z\, h^{+}_{i,t-1}(\hat{z}_{i,t-1}(\boldsymbol{\beta}), \hat{\boldsymbol{z}}_{t-1}(\boldsymbol{\beta}), H_{i,t-1}) + \eta_{it},
\end{aligned} \tag{10}
$$

where $\hat{\boldsymbol{z}}_{t-1}(\boldsymbol{\beta})$ is the cross-sectional vector of productivity proxies and $\eta_{it} \equiv \vartheta_{it} + \varepsilon_{it}$. In estimation, I approximate $\Upsilon(\cdot)$ by a third-order polynomial evaluated at $\hat{z}_{i,t-1}(\boldsymbol{\beta})$.

The second-stage moments use predetermined labor and capital stocks and their lags as instruments. Let $I_{i,t-1}$ denote the information set at the end of period $t-1$, after establishments know the lagged states, make the hiring choices used in $h^{+}_{i,t-1}$, and observe the productivity-ranking information used to construct it, but before the period-t productivity innovation and transitory shock are realized. The key exclusion is that, conditional on this information set, the instruments do not predict η_{it}. Under the timing assumptions,

$$
\mathbb{E}\big[\eta_{it}\, B(I_{i,t-1})\big] = 0 \tag{11}
$$

for the instrument vector $B(I_{i,t-1})$, consisting of predetermined labor and capital stocks and their lags.

Recovering productivity under a no-spillover productivity law and then testing for spillovers in a second regression would rank establishments using a productivity process that omits the variable being tested. This concern is familiar in production-function settings with controlled productivity processes: Doraszelski and Jaumandreu (2013) put R&D in the productivity transition, De Loecker (2013) does the same for export experience, and Malikov and Zhao (2023) handle peer-productivity terms that also depend on recovered latent productivity, making their estimator the closest methodological precedent. The distinct feature here is that the productivity-dependent object is rank-based and link-specific: for each candidate production function, the estimator recomputes sender and receiver productivity proxies, classifies which observed hires come from more productive previous employers, constructs $h^{+}_{i,t-1}$, and estimates β_z inside the same productivity law. This keeps the production elasticities, sender–receiver ranking, and spillover coefficient tied to one maintained productivity process.

2.3 Descriptive Evidence on Hires from More Productive Previous Employers

The empirical analysis uses matched employer–employee data for Finnish manufacturing establishments from 1995 to 2012. The estimation sample contains roughly 116 thousand establishment-year observations from 15 thousand unique establishments. Appendix A describes the data construction and capital measurement.

Even under a no-spillover productivity proxy, hiring from more productive previous employers is positively related to establishment productivity. I first compute productivity proxies using a standard control-function approach that imposes $\beta_z = 0$. I then use those proxies to rank establishments, construct h^{+}, and relate hiring from more productive previous employers to current productivity.

Figure 1 displays only a descriptive pattern, as both the productivity proxy and the

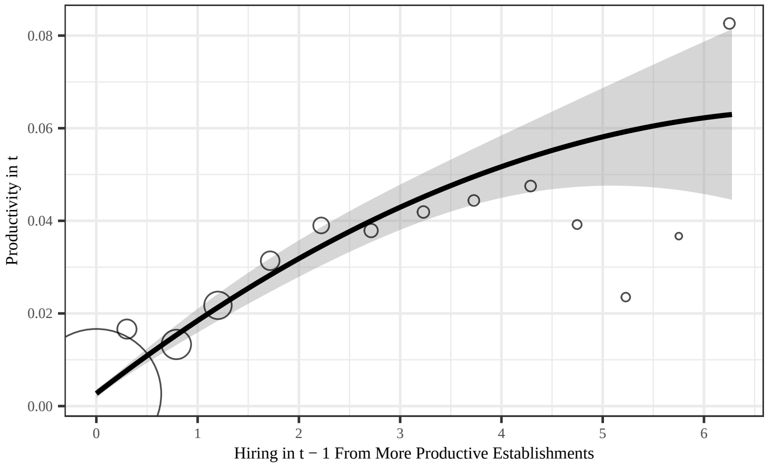

Notes: The figure contains only establishments with positive hiring. The productivity measure is the productivity residual from the ACF estimation. Circles present the number of observations in equispaced intervals on the x-axis.

Figure 1: Suggestive descriptive relationship between productivity and hiring from more productive establishments, using the no-spillover ACF productivity proxy.

previous-employer ranking are constructed under the no-spillover restriction. Among establishments with positive hiring, those with more hires from more productive previous employers have higher contemporaneous productivity residuals. The measure is concentrated near zero, and the fitted relationship is mildly concave, although a linear specification remains broadly consistent with the confidence intervals. The figure therefore motivates the augmented estimator without pinning down the magnitude of the spillover effect.[9]

2.4 Estimation Results of the Augmented Control-Function Approach

The augmented estimator implies a positive productivity response to hiring from more productive previous employers. The baseline estimate is $\beta_z = 0.0056$, with a bootstrap 95 percent confidence interval that excludes zero, so one such hire is associated with about 0.56 percent higher receiving-establishment productivity. At the sample mean of such hires, this corresponds to roughly 0.4 percentage points higher productivity before any general-equilibrium feedbacks.

Adding hiring from more productive previous employers to the productivity law leaves the recovered production function close to the standard ACF benchmark. In Table 1, the labor elasticity is about 0.70 and the capital elasticity about 0.20 in both the standard ACF specification and the specification with spillovers. This stability helps interpretation because the estimated productivity response does not come from a large change in recovered production technology after h^+ is added. The new coefficient instead captures the additional predictive content of hires from more productive previous employers within the enriched productivity law. This comparison is diagnostic rather than a separate identification test, and the interpretation still rests on the timing and information assumptions above.

[9]The figure uses the same residual-wage correction as the baseline estimate. Setting $v_{\ell,t-1} = 1$ leaves the relationship positive and increases its magnitude.

Table 1: Input elasticity and spillover estimates from the augmented control-function estimator.

	Baseline specifications		*Additional specifications*			
Estimate	ACF	ACF-SO	No wage adj.	Scaled by labor	Receiver split	Gap split
Labor elasticity	0.702 [0.684;0.722]	0.700 [0.680;0.720]	0.700 [0.680;0.719]	0.725 [0.701;0.752]	0.700 [0.680;0.719]	0.700 [0.679;0.721]
Capital elasticity	0.206 [0.190;0.220]	0.205 [0.186;0.218]	0.204 [0.188;0.220]	0.187 [0.171;0.216]	0.204 [0.186;0.219]	0.204 [0.186;0.219]
h^+_{t-1}		0.0056 [0.0029;0.0118]	0.0084 [0.0049;0.0143]			
$\frac{h^+_{t-1}}{n_{t-1}}$				0.0355 [0.0180;0.0562]		
$h^+_{t-1}\mathbb{I}(z_{t-1} < \bar{z})$					0.0136 [0.0107;0.0198]	
$h^+_{t-1}\mathbb{I}(z_{t-1} > \bar{z})$					0.0035 [0.0016;0.0101]	
$h^{+\leq 10\%}_{t-1}$						0.0032 [0.0013;0.0124]
$h^{+>10\%}_{t-1}$						0.0119 [0.0086;0.0154]
Residual-wage adjustment	–	Yes	No	Yes	Yes	Yes

Notes: Brackets report 95 percent bootstrap confidence intervals from 100 establishment-level resamples of the full estimator. Mobility links are kept intact within each bootstrap sample, and sender productivities for linked establishments outside the sample are recalculated at the current parameter values. The sample covers 1995–2012. ACF-SO is the augmented specification with hires from more productive previous employers in the productivity process. Labor and capital rows report output elasticities. The moments use predetermined labor and capital stocks and their lags as instruments, and productivity persistence is approximated with a third-order polynomial. The no-wage-adjustment column omits the residual-wage adjustment. The scaled column divides h^+_{t-1} by lagged labor. The receiver-split column allows the coefficient to differ by receiver productivity. The gap-split column allows the coefficient to differ by the sender–receiver productivity gap.

The alternative specifications check whether the baseline result is driven by the wage-based worker-quality adjustment, establishment size, or the productivity distance between a hire's previous employer and the hiring establishment. Removing the adjustment raises the coefficient to 0.0084, so the residual-wage weighting yields a smaller coefficient than the raw-count measure. Scaling hires from more productive previous employers by lagged labor input also preserves a positive coefficient, so the relationship appears when exposure is measured relative to establishment size. The heterogeneous estimates then sharpen the economic interpretation: the coefficient is larger for below-mean receiving establishments, and the gap split shows a larger coefficient when the previous employer is more than 10 percent more productive. This is the empirical pattern behind the gap-closing specification in the quantitative model.

Appendix B also reports a source–receiver gap measure closer to Stoyanov and Zubanov (2012). Instead of counting hires from more productive previous employers, the measure sums positive lagged productivity gaps between each hire's previous employer and the receiving establishment. Establishments with larger gap-based exposure have higher subsequent recovered productivity, so the data point in the same direction when exposure is measured by source–receiver distance rather than by counts.

Worker separations after hiring from more productive previous employers are consistent with an establishment-level interpretation of the productivity response. Appendix C compares establishments where such a hire separates immediately with establishments where the hire remains. Subsequent productivity dynamics are similar across the two groups, which points toward an establishment-level response rather than the productivity contribution of a particular retained worker.

Taken together, the estimates deliver the empirical input needed for the quantitative analysis: an average establishment-level productivity response to hiring from more productive previous employers. The next section embeds the same exposure measure, h^{+}, in a general equilibrium model to evaluate how this estimated productivity response changes aggregate growth and policy counterfactuals.

3 Model with Knowledge Diffusion through Worker Reallocation

The model embeds worker-transmitted knowledge in a growth version of the Hopenhayn–Rogerson industry-equilibrium framework. Establishments are the units of production, and baseline growth is sustained by productivity shocks, selection, and entrant imitation as in Poschke (2009). The additional channel is that hires from more productive previous employers create learning opportunities for receiving establishments. Since separations, exit, and cross-sectional employment determine which previous employers are represented in the worker-weighted sender distribution, establishment decisions shape not only employment dynamics but also how often useful outside knowledge is available to hiring establishments.

3.1 Environment and Timing

Time is discrete.[10] The economy consists of a representative household and a continuum of establishments whose mass evolves endogenously through entry and exit. The household consumes the final good, supplies labor, and owns establishment profits. Establishments differ in actual log productivity z_t and beginning-of-period employment n_t. An establishment with state (z_t, n_t) produces[11]

$$y_t = \exp(z_t) n_t^{\alpha}, \qquad 0 < \alpha < 1. \tag{12}$$

Let μ_t denote the beginning-of-period measure of establishments over states (z, n). Here and below, z and n without time subscripts denote generic state variables used inside distributions and policy functions. The law of motion for μ_t is given below; for now, μ_t is the cross-sectional state faced by establishments.

Within each period, current establishment states and worker flows determine next-period employment as follows. At the beginning of the period, an incumbent observes (z_t, n_t) and chooses whether to continue into the next period. All incumbents produce in the current period with inherited employment n_t. A continuing establishment chooses hires $h_t \geq 0$ and endogenous separations $s_t \geq 0$, and enters the next period with

$$n_{t+1} = (1 - \delta) n_t + h_t - s_t, \tag{13}$$

[10] The model is written from period t decisions to period $t+1$ states: worker flows and learning opportunities in period t affect n_{t+1} and z_{t+1}. Section 2 writes the corresponding transition with lagged objects predicting period-t productivity, which keeps the empirical presentation close to the production-function estimation convention.

[11] To keep the quantitative model focused on the hiring and separation margins through which knowledge diffuses, I abstract from capital, following labor-only firm dynamics models in the literature with decreasing returns.

where $\delta \in (0,1)$ is the exogenous separation rate. Feasible choices satisfy $n_{t+1} \geq 0$. As learning opportunities are attached to hires rather than to net employment growth, hiring and separations must be chosen separately; the model therefore allows, but does not require, establishments to churn workers when the learning benefit justifies the gross reallocation cost.

Gross worker reallocation is costly. The cost is paid in labor units and is convex:

$$d(h_t, s_t; n_t, n_{t+1}) = \frac{f_d}{2}\left(\frac{h_t + s_t}{\bar{n}_t}\right)^2 \bar{n}_t, \qquad \bar{n}_t = \frac{n_t + n_{t+1}}{2}. \tag{14}$$

The scaling by mean employment makes the cost depend on the reallocation rate, so gross reallocation is costly in comparable rate terms across establishment sizes. It also gives a finite adjustment cost for expansion from the small employment state at which entrants are valued.

3.2 Incumbent Problem

For a continuing establishment, the current flow payoff is revenue net of production labor, fixed operating labor, and adjustment labor:

$$\pi_t^c(z_t, n_t; h_t, s_t) = \exp(z_t) n_t^\alpha - w_t \left[n_t + f_f + d(h_t, s_t; n_t, n_{t+1})\right], \tag{15}$$

where w_t is the wage and f_f is a fixed operating cost.

If the establishment exits, it still produces in the current period and then separates from its remaining workers. The associated exit cost uses the same adjustment-cost function as continuation, evaluated at $h_t = 0$, $s_t = (1-\delta)n_t$, and $n_{t+1} = 0$. The time-t exit payoff is

$$V_t^x(z_t, n_t) = \exp(z_t) n_t^\alpha - w_t \left[n_t + f_f + d(0, (1-\delta)n_t; n_t, 0)\right]. \tag{16}$$

At the beginning of period t, an incumbent chooses whether to continue or exit. Given wages and the worker-weighted sender distribution defined below, its value is

$$V_t(z_t, n_t) = \max\{V_t^c(z_t, n_t), V_t^x(z_t, n_t)\}, \tag{17}$$

where

$$V_t^c(z_t, n_t) = \max_{h_t \geq 0,\, s_t \geq 0} \left\{\pi_t^c(z_t, n_t; h_t, s_t) + \mathcal{D}_t\, E\left[V_{t+1}(z_{t+1}, n_{t+1}) \mid z_t, n_t, h_t, s_t, F_t\right]\right\}, \tag{18}$$

subject to (13). The household discount factor $\mathcal{D}_t$ is defined below. Hiring has the standard dynamic role of changing next-period employment, and in this model it also raises the probability of learning from more productive previous employers.

3.3 Worker-Weighted Sender Distribution and Spillovers

Worker flows determine the set of previous employers from which hiring establishments can learn. A high-productivity establishment matters as a source of transferable knowledge only when some of its workers enter the hiring pool. The relevant object is therefore not the unweighted cross section of establishments, but the worker-weighted distribution of previous-employer productivity generated by exogenous separations, chosen separations, and exit.

Formally, let $e_t(z,n) \in \{0,1\}$ denote the exit indicator, with $e_t(z,n) = 1$ for establishments that exit. The flow of workers released from state (z,n) is

$$m_t^F(z,n) = \delta n + [1 - e_t(z,n)]\, s_t(z,n) + e_t(z,n)(1-\delta)n. \tag{19}$$

The three terms are exogenous separations, chosen separations by continuers, and liquidation separations at exit. Only released workers enter the sender distribution defined below. Labor used for fixed, entry, or adjustment costs is a resource cost and is not part of the worker flow from which hiring establishments draw.

The sender distribution is the normalized distribution of these worker flows. For any Borel set B,

$$F_t(B) = \frac{\int \mathbf{1}\{z \in B\} m_t^F(z,n)\, \mu_t(dz,dn)}{\int m_t^F(z,n)\, \mu_t(dz,dn)}. \tag{20}$$

Since F_t is constructed from equilibrium worker flows, the same decisions that determine employment dynamics also determine the distribution of previous employers' knowledge represented in the hiring pool.

A hiring establishment can learn from a worker only if the worker's previous employer is more productive than the receiver. For a receiver with productivity z_t, the relevant part of the sender distribution is therefore the upper tail

$$\psi^+(z_t; F_t) = F_t((z_t, \infty)). \tag{21}$$

Let ρ_H map total model hires into the subset of worker moves comparable to the empirical exposure measure. It captures hires with observed previous employers in the relevant manufacturing-worker flows, allowing for the fact that not every worker movement carries establishment-level knowledge that is useful for the receiver. The model analogue of the empirical exposure in Section 2 is

$$h_t^+ = \rho_H h_t \psi^+(z_t; F_t). \tag{22}$$

This exposure measure combines the establishment's hiring intensity with the availability of more productive potential senders. It is zero for establishments that do not hire and for establishments whose hires are not linked to more productive previous employers.

Exposure to hires from more productive previous employers creates an opportunity to absorb outside knowledge rather than a deterministic productivity increment. I treat h_t^+ as the arrival intensity of an implementable transfer. Let ξ_{t+1} equal one if at least one such transfer occurs between t and $t+1$. Then

$$\Pr(\xi_{t+1} = 1 \mid z_t, h_t, F_t) = 1 - \exp(-h_t^+). \tag{23}$$

This formulation treats the empirical coefficient as an average response across realized and unrealized learning opportunities. Additional such hires increase the probability of learning, while the gain is realized only when transferable knowledge is actually absorbed.[12]

When a transfer occurs, its content depends on where the worker came from. If $\psi^+(z_t; F_t) > 0$, the source productivity is drawn from the upper tail of the worker-weighted sender distribution,

$$z_t^+ \sim F_t(\cdot \mid z > z_t).$$

A successful transfer closes a fraction ϕ of the gap to the source:

$$z_{t+1} = z_t + \phi\, \xi_{t+1}(z_t^+ - z_t) + u_{t+1}, \qquad u_{t+1} \sim N(0, \sigma_u^2). \tag{24}$$

The gap-closing form is disciplined by the robustness results in Section 2: the estimated effect is larger for below-mean receivers and for hires whose previous employers are substantially more productive. If the upper tail is empty, then $\psi^+(z_t; F_t) = 0$ and the spillover term is zero.

3.4 Entry

Entry is governed by a standard free-entry condition: the expected value of a new establishment equals the labor cost of creating it. A potential entrant uses f_e units of labor, starts from employment $n_0 \geq 0$, and draws productivity from G_t. Thus, whenever entry is positive,

$$w_t f_e = \int V_t(z, n_0)\, G_t(dz) \tag{25}$$

The balanced-growth restriction below specifies how the location of G_t moves with the economy.

3.5 Household and Market Clearing

A representative household consumes the final good, supplies labor, and owns the establishments. Preferences are logarithmic in consumption and linear in labor, and aggregate establishment profits are rebated to the household. Normalizing the marginal disutility of labor to one, the static labor-supply condition is $w_t = C_t$. The household's intertemporal marginal rate of substitution prices next-period establishment payoffs, so $\mathcal{D}_t = \beta C_t / C_{t+1}$. Equivalently, the gross risk-free return satisfies $1 + r_{t+1} = 1/\mathcal{D}_t$.

Labor is used in production, fixed operating costs, worker adjustment, and entry. Let Ξ_t denote continuing establishments and Ξ_t^x exiting establishments. The labor requirements of a continuing establishment and an exiting establishment are

$$\ell_t^c(z, n) = n + f_f + d(h_t(z, n), s_t(z, n); n, n_t'(z, n)),$$
$$\ell^x(n) = n + f_f + d(0, (1-\delta)n; n, 0).$$

[12] The conditional productivity histogram in Figure 7 is consistent with, but not diagnostic of, this stochastic-transfer interpretation. After hiring from more productive previous employers, future productivity is bimodal rather than a simple shifted version of the unconditional distribution, as would be expected if only some exposed establishments realize a discrete productivity improvement.

With entry mass $m_{e,t}$, aggregate labor demand is

$$L_t^d = \int_{\Xi_t} \ell_t^c(z,n)\, \mu_t(dz,dn) + \int_{\Xi_t^x} \ell^x(n)\, \mu_t(dz,dn) + f_e m_{e,t}. \tag{26}$$

Labor market clearing requires $L_t = L_t^d$.

Aggregate profits are establishment revenues net of labor payments and entry costs:

$$\begin{aligned}\Pi_t = &\int_{\Xi_t} \left[\exp(z)n^\alpha - w_t \ell_t^c(z,n)\right] \mu_t(dz,dn) \\ &+ \int_{\Xi_t^x} \left[\exp(z)n^\alpha - w_t \ell^x(n)\right] \mu_t(dz,dn) - w_t f_e m_{e,t}.\end{aligned} \tag{27}$$

Goods market clearing is

$$C_t = \int \exp(z) n^\alpha\, \mu_t(dz,dn). \tag{28}$$

Adjustment, operating, and entry costs do not appear as separate goods uses because they are paid in labor units and enter feasibility through (26).

3.6 Distribution

The establishment distribution evolves through continuing incumbent transitions and entry. Let $Q_t((z,n),A)$ denote the probability that a continuing incumbent with current state (z,n) enters next period in a set A of establishment states, under the time-t policies and stochastic productivity law. Then, for any measurable set A of establishment states,

$$\mu_{t+1}(A) = \int_{\Xi_t} Q_t((z,n),A)\, \mu_t(dz,dn) + m_{e,t} \int \mathbf{1}\{(z,n_0) \in A\}\, G_t(dz). \tag{29}$$

Exiting establishments do not transition, and entrants enter at n_0 with productivity drawn from G_t.

3.7 Equilibrium

A competitive equilibrium is a sequence of household allocations, establishment policies, prices, entry masses, distributions, and worker-weighted sender distributions that are mutually consistent. Given an initial measure μ_0, it consists of

$$\{C_t, L_t, w_t, r_{t+1}, \mathcal{D}_t, V_t, h_t, s_t, e_t, n_t', m_{e,t}, \mu_t, F_t, G_t\}_{t \geq 0}.$$

In every period, households satisfy their optimality and budget conditions; incumbent establishments solve the continuation and exit problem; free entry holds when entry is positive; the establishment measure evolves according to (29); the worker-weighted sender distribution is given by (20); and labor and goods markets clear.

3.8 Balanced Growth and the Stationary Problem

To solve the model, I look for a balanced-growth path on which the distribution of establishments is stationary after removing the common productivity trend. Let ζ_t denote the location of the entrant log-productivity distribution and define temporary relative productivity and aggregate growth by

$$\tilde{z}_t = z_t - \zeta_t, \qquad g = \zeta_{t+1} - \zeta_t.$$

Goods-denominated variables are divided by $\exp(\zeta_t)$, while employment, worker flows, entry masses, and establishment measures remain counts or labor quantities. During this derivation, tildes mark relative productivity; after normalization, I drop tildes, so z denotes relative productivity in the stationary problem and quantitative analysis.

Entrants track the incumbent distribution from a fixed distance. In actual units,

$$z_e = \zeta_t + \varepsilon, \qquad \varepsilon \sim G,$$

where G is fixed with mean zero. Entrants are centered at zero in relative units, and the parameter κ sets the incumbent mean relative to the entrant location. Thus, in the stationary distribution over (z, n),

$$\bar{z} = \frac{\int z \, \mu(dz, dn)}{\int \mu(dz, dn)} = \kappa.$$

The equilibrium growth rate g is the value that makes this relative distribution, including its incumbent mean κ, time invariant.

The stationary equations have the same form as the time-indexed equations above, with actual productivity replaced by relative productivity and G_t replaced by G. Dividing values by $\exp(\zeta_t)$ makes the stationary discount factor equal to β: the household discount factor is $\beta C_t / C_{t+1}$, and the growth of next-period values contributes the offsetting factor $\exp(\zeta_{t+1})/\exp(\zeta_t)$. In particular, the stationary exposure measure is $h^+(z, n) = \rho_H h(z, n) F((z, \infty))$. The only substantive notational change is the productivity law. After applying the tilde-drop convention, if $z_t^+ \sim F(\cdot \mid z > z_t)$, subtracting ζ_{t+1} from (24) gives

$$z_{t+1} = z_t - g + \phi \, \xi_{t+1}(z_t^+ - z_t) + u_{t+1}, \qquad u_{t+1} \sim N(0, \sigma_u^2). \tag{30}$$

The shift leaves upper-tail shares and source–receiver gaps unchanged because senders and receivers are measured relative to the same location ζ_t. The household Euler equation implies $1 + r = \exp(g)/\beta$. From this point on, $(C, w, V, h, s, e, n', m_e, \mu, F)$ denotes the stationary representation used in the quantitative analysis.

3.9 Equilibrium Interpretation of Worker-Transmitted Knowledge

Productivity shocks, exit selection, and entrant imitation generate growth even without worker-transmitted knowledge. The shock u_{t+1} spreads the incumbent productivity distribution, exit truncates the lower tail, and entrants imitate the rising incumbent mean from a distance κ. These forces stand in for innovation, experimentation, and imitation margins that are not explicitly modeled here. Worker-transmitted knowledge adds a separate source of productivity growth because it makes the incumbent transition kernel depend on hiring choices and on the endogenous worker-weighted sender distribution.

The empirical regression coefficient summarizes how productivity growth covaries with exposure to hires from more productive previous employers in equilibrium. The structural parameter is narrower: ϕ is the fraction of a source–receiver productivity gap absorbed after a useful transfer occurs. A given count of hires from more productive previous employers can therefore correspond to different economic situations. A low-productivity establishment hiring from a thick upper tail has many potential sources and large gaps to close, while an establishment near the frontier may hire from more productive previous employers and still have little useful distance left.

The empirical slope is therefore an equilibrium statistic, not a primitive technology coefficient. The exposure measure h^+ records how often hiring draws workers from more productive previous employers. The productivity gain from such hiring also depends on where those previous employers sit in the worker-weighted sender distribution and how far the receiver is from them. The regression coefficient therefore combines exposure to more productive previous employers with source–receiver gap sizes, while ϕ governs the fraction of an available gap absorbed after a useful transfer occurs.

The same logic matters for counterfactuals because individual establishments do not take all aggregate feedbacks into account. A hiring or separation decision affects the establishment's own employment and learning opportunities, but in the aggregate these decisions also reshape the worker-weighted sender distribution and the incumbent distribution that entrants imitate. Policies that change worker flows therefore change both the number of hires from more productive previous employers and the economic content of those hires. The reduced-form spillover coefficient can move with the equilibrium environment, even when the structural gap-closing parameter ϕ is held fixed.

4 Quantitative Results

The quantitative exercise shows that worker-transmitted knowledge is not only a direct learning term. Once embedded in equilibrium, it changes which establishments hire, survive, and supply workers to the future hiring pool. In the calibrated economy, removing the channel lowers annual growth from 2.56% to 2.01%, and only part of this decline comes from the direct productivity gains of successful transfers. The rest reflects equilibrium changes in selection, entry, and the allocation of employment across establishments. I discipline the benchmark economy with establishment-level reallocation moments, aggregate growth, and the pooled spillover coefficient, and then use it to evaluate the no-spillover counterfactual and labor-market policy experiments.

4.1 Calibration

The calibration is annual and uses the empirical estimates from Section 2 wherever possible. I set $\beta = 0.95$ and $\alpha = 0.70$, with the latter equal to the estimated employment elasticity. Entrants start from the lowest employment grid point, so $n_0 \simeq 0$ in the quantitative model.

Table 2: Benchmark calibration and model fit

Variable	Value	Explanation	Data	Model
External				
β	0.950	Discount factor		
α	0.700	Labor curvature		
σ_z	0.310	Std. of entrant productivity		
σ_u	0.140	Std. of residual shock		
ρ_H	0.110	Exposure scale parameter		
Internal				
ϕ	0.0287	Inferred from spillover regression	0.00560	0.00557
f_e	2.28	Inferred from mean size	20.0	20.0
f_f	0.178	Inferred from exit rate	0.0670	0.0669
f_d	2.14	Inferred from job turnover	0.120	0.119
δ	0.106	Inferred from worker turnover	0.300	0.296
κ	0.393	Inferred from aggregate growth rate	0.0256	0.0256
Non-targeted Moments				
Mean hires from more productive previous employers			0.0760	0.131
Mean hires			3.00	2.96
Mean separations			3.30	2.96

The productivity process uses $\sigma_z = 0.31$ for entrant productivity dispersion and $\sigma_u = 0.14$ for residual innovations, both taken from the estimated productivity process after removing the spillover term. The exposure scale parameter is $\rho_H = 0.11$, which maps total model hires into the subset of moves with observed and relevant previous-employer productivity.

The remaining parameters are calibrated jointly: the spillover parameter ϕ, the entry cost f_e, the fixed operating cost f_f, the convex adjustment cost f_d, the exogenous separation rate δ, and the productivity-location parameter κ. The last parameter pins down the mean of the stationary incumbent productivity distribution in the balanced-growth equilibrium. The targets are mean establishment size, the exit rate, worker turnover, job turnover, the aggregate growth rate, and the pooled spillover coefficient from the empirical section. The spillover target is constructed exactly as in the empirical regression: all establishments enter the regression, so non-hiring establishments and establishments with no hires from more productive previous employers contribute zeros to exposure.

Table 2 reports the calibrated parameters and model fit. The targeted moments are matched closely, including the moments most important for the aggregate exercise: worker turnover, job turnover, aggregate growth, and the pooled spillover coefficient. The fit therefore disciplines both the amount of worker reallocation in equilibrium and the average productivity response to hires from more productive previous employers.

Non-targeted hiring and separation moments provide additional checks on the scale of worker flows. The model matches average hires and separations closely, but not fully, reflecting the parsimonious one-sector adjustment-cost structure. Existing legal, contractual, and institutional barriers to worker reallocation are therefore absorbed into the calibrated adjustment cost; the policy experiments below add the firing cost and hiring credit on top of this fitted benchmark.

Table 3: Removing worker-transmitted knowledge: aggregate and allocation effects

Model statistic	Benchmark with spillovers	No knowledge spillovers
Aggregate outcomes		
Output (index)	100	101.6
Growth (%)	2.56	2.01
Compensation needed (%)	0.00	5.13
Labor-market allocation		
Labor demand (index)	100	94.8
Establishment mass (index)	100	191
Entry/exit rate (%)	6.69	5.24
Mean size	20.0	9.86
Job turnover (%)	11.9	13.3
Worker turnover (%)	29.6	26.0

Notes: The economy without knowledge spillovers sets $\phi = 0$ and resolves the balanced-growth equilibrium holding the other benchmark primitives fixed. The benchmark economy has $\phi = 0.029$. Output, labor demand, and establishment mass are indexed to 100 in the benchmark economy. Compensation needed is the permanent consumption increase in the economy without knowledge spillovers required to make the household indifferent to the benchmark balanced-growth path.

4.2 The Quantitative Importance of Knowledge Diffusion through Worker Reallocation

Knowledge diffusion through worker reallocation makes a quantitatively large contribution to aggregate growth. In Table 3, setting $\phi = 0$ while holding the other benchmark primitives fixed lowers the balanced growth rate from 2.56% to 2.01%, a decline of 0.55 percentage points. The household would need 5.13% higher consumption in the economy without worker-transmitted knowledge to be indifferent to the benchmark because this permanent growth loss dominates the small change in the stationary output level.

The growth effect reflects equilibrium changes in both worker reallocation and establishment replacement. In the benchmark, worker turnover is 29.6% rather than 26.0%, so more worker movements create opportunities for learning from previous employers. The entry/exit rate is also higher, 6.69% rather than 5.24%, and in the low-employment states where exit occurs, the benchmark continuation cutoff is higher. Exit therefore truncates more of the lower tail, leaves higher-productivity continuing establishments, and raises the productivity trend more quickly because entrant productivity tracks the incumbent distribution from a fixed distance. These responses show that entry and exit amplify the growth effect of spillovers through entrant imitation as well.

Decreasing returns partly offset the growth loss in the stationary output level. Without spillovers, the output index is 101.6 and labor demand falls to 94.8% of the benchmark, while the mass of establishments rises to 191% and mean size falls from 20.0 to 9.86. The economy without worker-transmitted knowledge is therefore more fragmented. Because production has decreasing returns at the establishment level, spreading activity across many smaller establishments can raise contemporaneous output even when the economy grows more slowly. The welfare result shows that this static level gain is dominated by the growth loss.

Figure 2: Policy responses induced by worker-transmitted knowledge

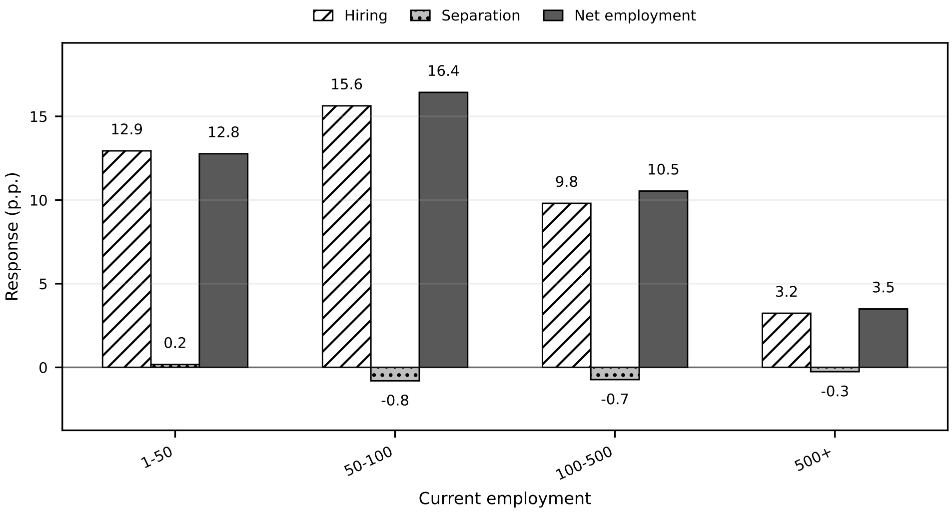

Figure 3: Size distribution with and without worker-transmitted knowledge

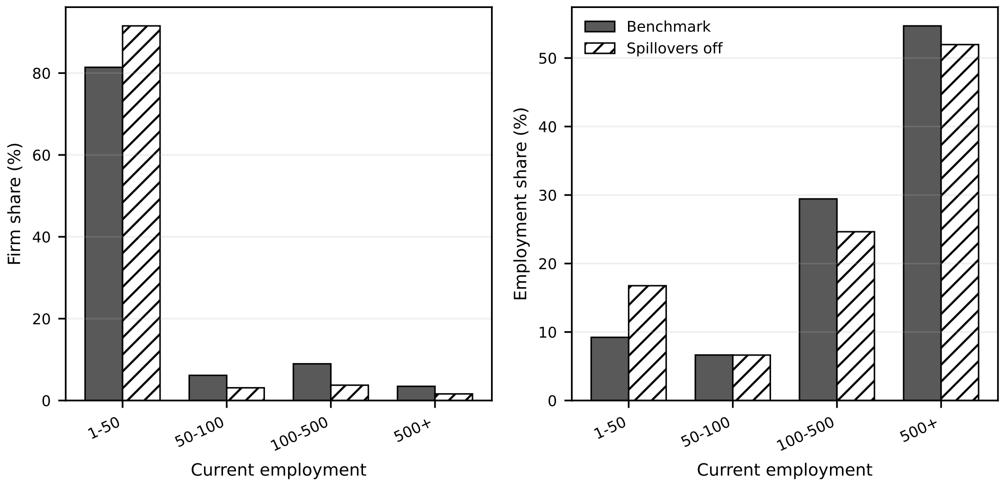

4.3 Equilibrium Responses When Worker-Transmitted Knowledge Is Removed

Worker-transmitted knowledge raises hiring most among smaller incumbents, where useful outside knowledge has the most room to improve productivity. Figures 2–3 compare the benchmark equilibrium with the equilibrium without worker-transmitted knowledge, holding the other primitives fixed. In the benchmark, continuers hire more in every size bin. The increases are largest below 100 workers, with hiring rising by 12.9 percentage points in the 1–50 bin and by 15.6 percentage points in the 50–100 bin.

The extra hiring mainly changes the allocation of employment across establishments. In the benchmark, the 1–50 bin accounts for 81.4% rather than 91.5% of establishments and 9.2% rather than 16.8% of employment, while more employment is allocated to larger establish-

Figure 4: Spillover incidence and gains by establishment size

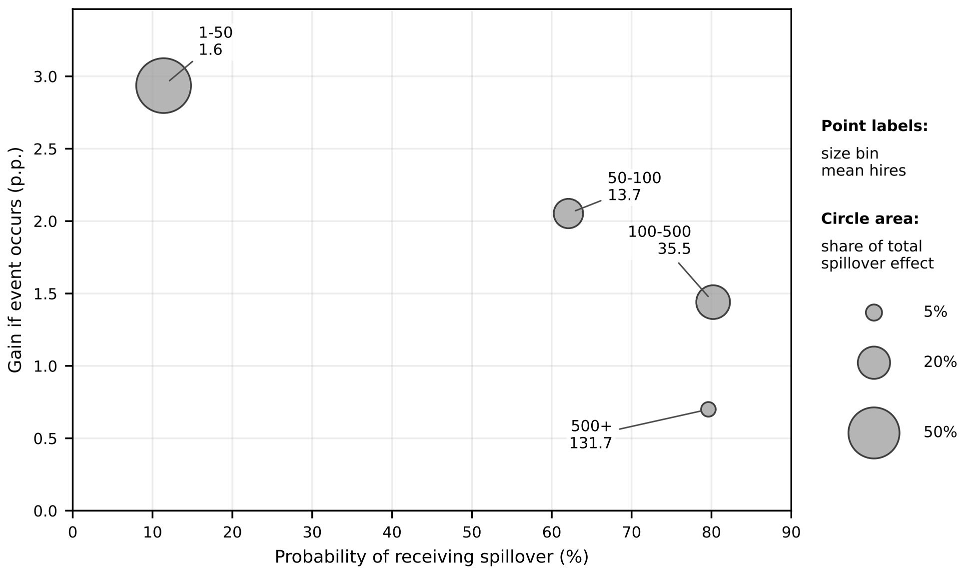

ments. Chosen separations barely move across size bins, so the response operates through hiring and the size distribution rather than substantial churn.

The aggregate spillover contribution depends on the mass of establishments in a size bin, the probability of a successful transfer, and the gain conditional on that transfer. Figure 4 displays this tradeoff by establishment size. Establishments with at least 500 workers have at least one successful learning transfer with probability 79.6%, but their gain conditional on a transfer is only 0.70 percentage points. Establishments with 1–50 workers have a successful transfer with probability 11.4%, but their conditional gain is 2.94 percentage points, so they account for 57.5% of expected productivity gains from useful transfers. A partial-equilibrium use of the empirical coefficient would miss this tradeoff because entry, exit, hiring, and separations jointly determine which previous employers appear in worker flows and how often hiring establishments draw from them.

The fall in growth after removing spillovers can be read from the stationary accounting of relative productivity. Write the establishment state as $x = (z, n)$ and denote the stationary distribution by μ. Continuing establishments form the set Ξ, with mass $M_C = \mu(\Xi)$; total establishment mass is $M = \int \mu(dx)$. Since μ is stationary in relative-productivity units,

$$0 = \underbrace{\frac{M_C}{M}\,(\bar{z}_C - \bar{z})}_{\text{exit selection}} + \underbrace{\frac{M_C}{M}\,\big(\bar{z}'_C - \bar{z}_C\big)}_{\text{continuing transition}} + \underbrace{\frac{M - M_C}{M}\,(\bar{z}_E - \bar{z})}_{\text{below-mean entry}}, \tag{31}$$

where bars denote means, $\bar{z}_C$ is the current mean among continuers, $\bar{z}'_C$ is their expected post-transition mean, and $\bar{z}_E = \int z\, G(dz)$ is the entrant mean. The continuing-transition term is tied to actual productivity growth because actual log productivity equals the common location ζ_t plus relative productivity. If Δ^C denotes expected actual productivity growth among

Table 4: Growth decomposition after removing worker-transmitted knowledge

Component (percentage points)	Benchmark	Spillovers off
Aggregate		
Balanced growth rate	2.56	2.01
Stationary relative-productivity accounting: effect on the mean		
Exit selection: continuers vs. all establishments	4.86	3.97
Continuing transition: post-transition change	-2.24	-1.91
Below-mean entry: entrants vs. all establishments	-2.63	-2.06
Continuing establishments: expected productivity growth		
Growth without successful transfers	-0.01	0.00
Useful-transfer gain	0.17	0.00
Total continuer productivity growth	0.16	0.00

Notes: Entries are annual percentage points. In the stationary accounting block, positive entries raise mean relative productivity and negative entries lower it; the three rows sum to zero by stationarity, up to rounding. The continuing-transition row is the post-transition change for continuers, scaled by their mass share. Continuing-establishment growth rows average over establishments that remain active.

continuers, then

$$\bar{z}'_C - \bar{z}_C = \Delta^C - g,$$

so the continuing-transition term in (31) equals $(M_C/M)(\Delta^C - g)$.

The accounting separates direct learning among continuing establishments from the selection and entrant-imitation terms that operate through the stationary distribution. I split Δ^C into expected productivity growth absent a successful transfer and the expected gain generated by successful transfers, averaging both terms over continuing establishments. This split describes the continuing-establishment term in (31); Table 4 reports it together with the stationary accounting terms in percentage points.

The growth decomposition shows that useful transfers are amplified by the equilibrium composition response. They add 0.17 percentage points to expected productivity growth among continuers, while the full balanced-growth difference is 0.55 percentage points. The gap arises because worker-transmitted knowledge also changes selection and entrant imitation. With spillovers, exit selection raises mean relative productivity by 4.86 percentage points rather than 3.97. Because entry and exit are more frequent and entrants enter below the incumbent mean, the below-mean entry term is also more negative, -2.63 rather than -2.06 percentage points. Worker-transmitted knowledge therefore raises growth through direct learning at continuing establishments and through the entry and exit dynamics that determine which establishments operate and how entrants track the incumbent distribution.

4.4 Interpreting the Spillover Regression Coefficient

The fitted spillover coefficient summarizes three policy-sensitive margins in the stationary economy. Producer decisions determine exposure to hires from more productive previous employers, exposure maps into the probability of at least one successful transfer, and sender–receiver composition determines the productivity gap available to close. For state x, exposure is $h^+(x)$, and the probability of at least one successful transfer is $p_\xi(x) = 1 - \exp[-h^+(x)]$. The

Table 5: Model-implied spillover coefficient decomposition

Contribution to pooled slope	Value
Residual drift and mean reversion	-0.0002
Exposure at average transfer gain	0.0498
Transfer-arrival nonlinearity	-0.0291
Source–receiver gap variation	-0.0150
Total spillover contribution	0.0057
Total model coefficient	0.0056

Notes: Entries decompose the model-implied pooled slope of productivity growth on exposure h^+. Each row applies the covariance formula in the text to one component of (32), using the stationary establishment distribution as weights. The components sum to the total spillover contribution.

conditional gain $D(x) = \mathbb{E}_1[z' \mid x] - \mathbb{E}_0[z' \mid x]$ compares expected next productivity after a successful transfer with expected next productivity without one, holding the current state and stationary policies fixed; $\bar{D}$ is its stationary regression-sample mean. With location growth added back, expected productivity growth in state x can be written as

$$\begin{aligned} m(x) &\equiv E[g + z' - z \mid x] \\ &= \underbrace{g + \mathbb{E}_0[z' \mid x] - z}_{m_0(x)} + \underbrace{h^+(x)\bar{D}}_{m_R(x)} + \underbrace{\big(p_\xi(x) - h^+(x)\big)\bar{D}}_{m_N(x)} + \underbrace{p_\xi(x)\big(D(x) - \bar{D}\big)}_{m_H(x)}. \end{aligned} \tag{32}$$

Table 5 reports

$$\hat{\beta}_j = \frac{\mathrm{Cov}_\mu(h^+(x), m_j(x))}{\mathrm{Var}_\mu(h^+(x))}, \qquad \hat{\beta} = \sum_j \hat{\beta}_j,$$

for the four terms in (32). This is the slope that would be obtained by regressing model-implied productivity growth on exposure to hires from more productive previous employers.

The decomposition gives the economic scale of these margins. If exposure mapped linearly into the average conditional gain, the implied coefficient would be 0.0498 rather than the fitted value of 0.0056. The transfer-arrival nonlinearity subtracts 0.0291 because additional exposure raises the probability of at least one useful transfer less than one-for-one. Gap variation subtracts another 0.0150 because high-exposure states often have less distance left to close. Policies that change hiring, separations, entry, or exit therefore change what the regression coefficient averages over: exposure, the probability of a successful transfer, and source–receiver gaps. The counterfactuals below are chosen around these margins: lower fluidity and firing costs work through worker flows from more productive previous employers, while the large-firm hiring credit changes which previous employers are represented in worker flows.

4.5 Labor-Market Fluidity

Exogenous separations have two opposing roles in the model. They force establishments to rebuild employment through the costly hiring margin before next-period productivity is known, but they also populate the worker-weighted sender distribution that can carry knowledge

Table 6: Lower labor-market fluidity: aggregate, allocation, and sender-distribution effects

Effect of halving δ	With spillovers	Recalibrated without spillovers	Difference
Aggregate outcomes			
Output (%)	1.32	2.08	-0.76
Growth (p.p.)	-0.14	0.01	-0.15
Compensation needed (%)	1.88	-1.35	3.23
Labor-market allocation			
Labor demand (%)	0.62	1.27	-0.65
Establishment mass (%)	10.9	-1.27	12.2
Entry/exit rate (p.p.)	-0.30	0.22	-0.52
Mean size (%)	-9.19	5.66	-14.9
Job turnover (p.p.)	0.91	-0.40	1.31
Worker turnover (p.p.)	-5.94	-10.4	4.51
Sender distribution and learning exposure			
Sender productivity ratio	-0.331	-0.527	0.196
Exposure to hires from more productive previous employers (%)	-31.6	-37.0	5.40
Successful-transfer probability (p.p.)	-1.66	-1.90	0.24
Direct gain from successful transfers (p.p.)	-0.05	0.00	-0.05

Notes: Entries are changes from each economy's calibrated baseline after multiplying the exogenous separation rate δ by 0.5. All other calibrated primitives are held fixed, and each economy resolves wages and growth. Successful-transfer probability is the exposure-implied arrival probability computed with the same mapping in both economies; when $\phi = 0$, it generates no productivity gain. The direct gain from successful transfers is zero in the economy without spillovers by construction. The difference column subtracts the response in the recalibrated economy without spillovers from the benchmark response.

across establishments.[13] The experiment halves the component of worker reallocation that establishments do not choose directly. It asks whether the resource savings from lower exogenous separations are outweighed by the loss of worker-transmitted knowledge.

In the economy without spillovers, lower exogenous separations raise welfare because they mainly force costly rebuilding of employment. Halving the exogenous separation rate raises output by 2.08% with essentially no growth loss, giving a compensation measure of -1.35%. The resource saving therefore dominates when worker reallocation does not transmit productive knowledge across establishments.

The welfare sign reverses when worker reallocation transmits knowledge. Output still rises by 1.32%, but growth falls by 0.14 percentage points and the household would need 1.88% higher consumption to be indifferent. Fewer exogenous separations reduce opportunities for worker-transmitted knowledge: worker turnover falls by 5.94 percentage points, exposure to hires from more productive previous employers falls by 31.6%, and the probability of at least one successful transfer falls by 1.66 percentage points. The worker-weighted sender distribution also becomes less productive, as mean source productivity falls by 0.331 relative to mean establishment productivity. Lower fluidity therefore saves adjustment resources, but in the benchmark economy it also reduces useful-transfer opportunities and weakens the source distribution of worker flows that supports future knowledge spillovers.

[13] The exercise is motivated by the evidence and discussion of declining U.S. labor-market fluidity in Davis and Haltiwanger (2014) and Molloy, Smith, Trezzi, and Wozniak (2016). Related evidence on declining business dynamism is summarized by Decker, Haltiwanger, Jarmin, and Miranda (2014).

4.6 Policy Counterfactuals

Policies that change worker reallocation also change the aggregate knowledge environment. Establishments internalize their own adjustment costs and the learning they receive from hires, but their hiring, separation, expansion, and exit decisions also determine which previous employers are represented in worker flows, how often hiring establishments draw workers from more productive employers, and which incumbent distribution entrants imitate. These aggregate feedbacks mean that a policy can save adjustment resources or reduce static misallocation while also changing growth through learning opportunities, the knowledge carried by worker flows, and imitation.

The counterfactuals use this logic to ask where worker-transmitted knowledge changes policy conclusions. A firing cost mainly compresses separations and worker flows, so it tests the role of hires from more productive previous employers. A large-firm hiring credit shifts expansion toward large establishments, so it tests whether policy conclusions depend on which previous employers are represented in worker flows. For each policy, I solve the change in the benchmark economy and in a separately calibrated economy without worker-transmitted knowledge. The reported entries are policy-minus-baseline changes; the difference column compares the benchmark response with the response in the recalibrated no-spillover economy. Appendix D reports the calibration without spillovers.

4.6.1 Firing Costs

A firing cost isolates the role of knowledge spillovers through worker reallocation by discouraging chosen separations and exit liquidation. Relative to the benchmark incumbent problem in Section 3, the policy sets an incremental firing cost $\tau_f = 1$ on top of the calibrated convex adjustment cost that already disciplines ordinary worker reallocation.[14] A continuing establishment that chooses endogenous separations s pays $w\tau_f s$, and an exiting establishment pays $w\tau_f(1-\delta)n$ on the post-exogenous-separation workforce liquidated at exit. Exogenous separations δn are not taxed. Tax receipts are rebated lump-sum to the household, so the firing cost changes separation incentives but is not a real resource use.

A firing cost is more costly when worker reallocation transmits knowledge. In Table 7, the firing cost lowers growth by 0.11 percentage points and requires compensation equal to 1.40% of consumption in the benchmark economy. In the recalibrated economy without spillovers, growth falls by only 0.02 percentage points and compensation is 0.54%. Tax payments are larger in the benchmark economy, at 1.14% of output rather than 0.30%, and are rebated to households. The welfare gap therefore comes from the induced changes in establishment decisions and worker flows, not from resources absorbed by the tax.

[14] I set $\tau_f = 1$ to match the one-year-wage firing-cost experiment in Poschke (2009). Appendix E reports the same experiment under an exit-exempt convention. Related firing-cost and employment-protection models include Hopenhayn and Rogerson (1993), Moscoso Boedo and Mukoyama (2012), Mukoyama and Osotimehin (2019), and Da-Rocha, Restuccia, and Tavares (2019); suggestive empirical evidence on employment protection and productivity includes Autor, William R. Kerr, and Kugler (2007). This is an incremental employment-protection wedge rather than a measure of all separation frictions. As a legislation-based scale check, weighting Finnish statutory employer notice periods by the OECD tenure distribution gives $\tau_f = 0.269$. Re-solving the experiment at that value changes growth by -0.04 percentage points with spillovers and +0.01 percentage points in the recalibrated economy without spillovers; in both economies, the compensation measure is about 0.3%.

Table 7: Firing costs: aggregate, allocation, and sender-distribution effects

Effect of one-wage firing cost	With spillovers	Recalibrated without spillovers	Difference
Aggregate outcomes			
Output (%)	0.82	-0.23	1.05
Growth (p.p.)	-0.11	-0.02	-0.10
Compensation needed (%)	1.40	0.54	0.86
Labor-market allocation			
Labor demand (%)	0.07	-0.03	0.10
Establishment mass (%)	-2.71	-3.50	0.79
Entry/exit rate (p.p.)	-0.35	0.10	-0.45
Mean size (%)	4.81	4.34	0.47
Job turnover (p.p.)	-2.92	-0.53	-2.39
Worker turnover (p.p.)	-5.33	-1.39	-3.94
Sender distribution and learning exposure			
Sender productivity ratio	0.125	0.034	0.091
Exposure to hires from more productive previous employers (%)	-16.6	-1.01	-15.6
Successful-transfer probability (p.p.)	-1.12	-0.09	-1.04
Direct gain from successful transfers (p.p.)	-0.03	0.00	-0.03
Policy scale			
Firing-tax payments (% output)	1.14	0.30	0.84

Notes: Entries are policy-minus-baseline changes. The firing cost sets $\tau_f = 1$ in (15); it applies to endogenous separations by continuing establishments and to the post-δ workers released at exit. Tax receipts are rebated lump-sum to the household. Successful-transfer probability is the exposure-implied arrival probability computed with the same mapping in both economies; when $\phi = 0$, it generates no productivity gain. The direct gain from successful transfers is zero in the economy without spillovers by construction. The difference column subtracts the response in the recalibrated economy without spillovers from the benchmark response.

Firing costs improve the average productivity of the remaining worker-weighted sender distribution, but they reduce the worker movements through which hiring establishments draw from it. This is the policy analogue of the fluidity experiment: fewer worker movements reduce adjustment activity, but they also reduce learning opportunities from more productive previous employers. Worker turnover falls by 5.33 percentage points, exposure to hires from more productive previous employers falls by 16.6%, and the probability of at least one successful transfer falls by 1.12 percentage points. The direct gain from successful transfers falls by about 0.03 percentage points. A model without spillovers fitted to the standard moments captures part of the labor-market response, but it misses this knowledge-transmission margin.

4.6.2 Large-Firm Hiring Credit and the Sender-Distribution Feedback

A large-firm hiring credit shifts employment toward establishments whose workers carry more useful knowledge when reallocated. Unlike a broad hiring subsidy, it changes which previous employers are represented in worker flows rather than only the overall volume of hiring. The targeted credit asks whether increasing learning opportunities from more productive employers can compensate for the misallocation created by favoring large establishments.[15]

Relative to the benchmark incumbent problem, the only change is a credit for observable

[15] Size-dependent policies are usually studied as sources of misallocation, as in Guner, Ventura, and Xu (2008), Restuccia and Rogerson (2008), and Garicano, Lelarge, and Van Reenen (2016).

Table 8: Large-firm hiring credit: aggregate, allocation, and sender-distribution effects

Effect of large-firm hiring credit	With spillovers	Recalibrated without spillovers	Difference
Aggregate outcomes			
Output (%)	0.10	-0.22	0.32
Growth (p.p.)	0.03	0.02	0.01
Compensation needed (%)	1.05	1.29	-0.24
Labor-market allocation			
Labor demand (%)	2.28	2.07	0.21
Establishment mass (%)	-6.80	-3.00	-3.80
Entry/exit rate (p.p.)	0.07	0.24	-0.17
Mean size (%)	9.72	5.32	4.40
Job turnover (p.p.)	0.24	0.08	0.15
Worker turnover (p.p.)	0.30	0.32	-0.02
Sender distribution and learning exposure			
Sender productivity ratio	0.150	0.111	0.0388
Exposure to hires from more productive previous employers (%)	9.33	6.12	3.21
Successful-transfer probability (p.p.)	0.27	0.06	0.21
Direct gain from successful transfers (p.p.)	0.01	0.00	0.01
Policy scale			
Eligible net-hire volume (%)	44.7	14.2	30.6
Subsidy bill (% output)	1.00	1.00	0.00

Notes: Entries are policy-minus-baseline changes. Establishments with beginning-of-period employment $n \geq 1000$ receive a credit $\tau_{LF} w \max\{h-s, 0\}$, financed by a lump-sum tax and treated as a transfer. The credit rate is chosen separately in each policy economy to target a subsidy bill of 1% of output: $\tau_{LF} = 0.2639$ with spillovers and $\tau_{LF} = 0.1914$ in the recalibrated economy without spillovers. The spillover technology is unchanged. Successful-transfer probability is the exposure-implied arrival probability computed with the same mapping in both economies; when $\phi = 0$, it generates no productivity gain. The direct gain from successful transfers is zero in the economy without spillovers by construction. The difference column subtracts the response in the recalibrated economy without spillovers from the benchmark response.

net job creation among large establishments. Establishments with beginning-of-period employment $n \geq 1000$ receive $\tau_{LF} w \max\{h - s, 0\}$, where h is hiring and s is chosen separations. The credit therefore rewards net new jobs and replacement of exogenous separations, but not simultaneous chosen churn. It is financed by a lump-sum tax on the household and treated as a transfer. The experiment chooses τ_{LF} separately in each policy economy so that the subsidy bill equals 1% of output.

The credit moves employment toward establishments whose workers can carry useful knowledge when they later become available for hire, but this aggregate benefit is quantitatively small. In the benchmark economy, mean establishment size rises by 9.72% as labor demand expands and the mass of establishments contracts, and eligible net-hire volume at large establishments rises by 44.7%. These reallocations change the source distribution in the predicted direction: exposure to hires from more productive previous employers rises by 9.33%, and the probability of at least one successful transfer rises by 0.27 percentage points. The direct gain from successful transfers rises by only 0.01 percentage points, so the intended effect is present but its quantitative pass-through is modest at this subsidy scale.

The aggregate response is correspondingly limited. At a subsidy bill of 1% of output, balanced growth rises by 0.03 percentage points with spillovers, only 0.01 percentage points more than in the recalibrated economy without spillovers. The compensation measure is also only modestly lower with spillovers, 1.05% rather than 1.29%. The signs are consistent with the argument for shifting employment toward establishments whose workers carry more useful knowledge when reallocated, but the induced increase in exposure to hires from more productive previous employers is too small to dominate the policy distortion at this scale.

5 Conclusion

This paper shows that worker reallocation can matter for aggregate growth not only by moving labor across establishments, but also by moving productive knowledge across them. In Finnish matched employer–employee data for manufacturing, establishments that hire from more productive previous employers subsequently become more productive. I embed this estimated response in an industry-equilibrium model where worker flows determine which previous employers are represented in the hiring pool and which establishments hire from those previous employers. The quantitative result is that this channel is large enough to matter: shutting down worker-transmitted knowledge lowers balanced growth by 0.55 percentage points, about one fifth of the benchmark rate. The loss reflects both the disappearance of direct learning events and equilibrium changes in which establishments survive, how large they become, and whose workers are later reallocated.

The policy counterfactuals show that labor-market policy changes growth by reshaping worker flows that generate learning opportunities. Firing costs reduce growth more when worker-transmitted knowledge is active because they reduce worker flows and exposure to hires from more productive previous employers. A hiring credit for large establishments shifts employment toward establishments whose workers carry more useful knowledge when reallocated, but the aggregate gains are modest at the subsidy scale considered. The effects of labor-market policies therefore have to be evaluated together with the producer dynamics that determine whose knowledge enters the hiring pool and which establishments can use it.

References

Ackerberg, D. A., Caves, K., & Frazer, G. (2015). Identification Properties of Recent Production Function Estimators. *Econometrica*, *83*(6), 2411–2451.

Akcigit, U., & Ates, S. T. (2023). What happened to US business dynamism? *Journal of Political Economy*, *131*(8), 2059–2124.

Autor, D. H., William R. Kerr, & Kugler, A. D. (2007). Does Employment Protection Reduce Productivity? Evidence from US States. *The Economic Journal*, *117*(521), F189–F217.

Baslandze, S., & Vardishvili, I. (2026). *Born different: Entrepreneurship through Inventor Mobility, Innovation, and Growth* (CEPR Discussion Paper No. 21016). Centre for Economic Policy Research.

Bradley, J., & Gottfries, A. (2026). Labor Market Dynamics and Growth. *Journal of the European Economic Association*, *24*(1), 242–270.

Buera, F. J., & Lucas, R. E. (2018). Idea Flows and Economic Growth. *Annual Review of Economics*, *10*(1), 315–345.

Buera, F. J., & Oberfield, E. (2020). The Global Diffusion of Ideas. *Econometrica*, *88*(1), 83–114.

Da-Rocha, J.-M., Restuccia, D., & Tavares, M. M. (2019). Firing costs, misallocation, and aggregate productivity. *Journal of Economic Dynamics and Control*, *98*, 60–81.

Davis, S. J., & Haltiwanger, J. (2014). *Labor Market Fluidity and Economic Performance* (NBER Working Paper No. 20479). National Bureau of Economic Research.

Decker, R., Haltiwanger, J., Jarmin, R., & Miranda, J. (2014). The Role of Entrepreneurship in US Job Creation and Economic Dynamism. *Journal of Economic Perspectives*, *28*(3), 3–24.

De Loecker, J. (2013). Detecting learning by exporting. *American Economic Journal: Microeconomics*, *5*(3), 1–21.

Doraszelski, U., & Jaumandreu, J. (2013). R&d and productivity: Estimating endogenous productivity. *The Review of Economic Studies*, *80*(4), 1338–1383.

Engbom, N. (2022). *Labor Market Fluidity and Human Capital Accumulation* (NBER Working Paper No. 29698). National Bureau of Economic Research.

Gabler, A., & Poschke, M. (2013). Experimentation by firms, distortions, and aggregate productivity. *Special issue: Misallocation and Productivity*, *16*(1), 26–38.

Garicano, L., Lelarge, C., & Van Reenen, J. (2016). Firm size distortions and the productivity distribution: Evidence from france. *American Economic Review*, *106*(11), 3439–3479.

Gregory, V. (2026). Firms as Learning Environments: Implications for Earnings Dynamics and Job Search. *Review of Economic Studies*.

Guner, N., Ventura, G., & Xu, Y. (2008). Macroeconomic implications of size-dependent policies. *Review of Economic Dynamics*, *11*(4), 721–744.

Heggedal, T.-R., Moen, E. R., & Preugschat, E. (2017). Productivity spillovers through labor mobility in search equilibrium. *Journal of Economic Theory*, *169*, 551–602.

Hopenhayn, H., & Rogerson, R. (1993). Job Turnover and Policy Evaluation: A General Equilibrium Analysis. *Journal of Political Economy*, *101*(5), 915–938.

Jarosch, G., Oberfield, E., & Rossi-Hansberg, E. (2021). Learning From Coworkers. *Econometrica*, *89*(2), 647–676.

Jo, K., & Kim, S. (2026). Heterogeneous innovations and growth under imperfect technology spillovers. *Journal of Monetary Economics*, *158*, 103899.

Levinsohn, J., & Petrin, A. (2003). Estimating Production Functions Using Inputs to Control for Unobservables. *The Review of Economic Studies*, *70*(2), 317–341.

Lucas, R. E. (2009). Ideas and Growth. *Economica*, *76*(301), 1–19.

Lucas, R. E., & Moll, B. (2014). Knowledge Growth and the Allocation of Time. *Journal of Political Economy*, *122*(1), 1–51.

Malikov, E., & Zhao, S. (2023). On the estimation of cross-firm productivity spillovers with an application to fdi. *Review of Economics and Statistics*, *105*(5), 1207–1223.

Molloy, R., Smith, C. L., Trezzi, R., & Wozniak, A. (2016). Understanding Declining Fluidity in the U.S. Labor Market. *Brookings Papers on Economic Activity*, *47*(1), 183–259.

Moscoso Boedo, H. J., & Mukoyama, T. (2012). Evaluating the effects of entry regulations and firing costs on international income differences. *Journal of Economic Growth*, *17*(2), 143–170.

Mukoyama, T., & Osotimehin, S. (2019). Barriers to Reallocation and Economic Growth: The Effects of Firing Costs. *American Economic Journal: Macroeconomics*, *11*(4), 235–70.

Olley, G. S., & Pakes, A. (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, *64*(6), 1263–1297.

Parrotta, P., & Pozzoli, D. (2012). The effect of learning by hiring on productivity. *The RAND Journal of Economics*, *43*(1), 167–185.

Perla, J., & Tonetti, C. (2014). Equilibrium Imitation and Growth. *Journal of Political Economy*, *122*(1), 52–76.

Perla, J., Tonetti, C., & Waugh, M. E. (2021). Equilibrium Technology Diffusion, Trade, and Growth. *American Economic Review*, *111*(1), 73–128.

Petit, B. (2023). Aggregate effects of firing costs with endogenous firm productivity growth. *Macroeconomic Dynamics*, *27*(4), 945–965.

Poschke, M. (2009). Employment protection, firm selection, and growth. *Journal of Monetary Economics*, *56*(8), 1074–1085.

Raurich, X., Sánchez-Losada, F., & Vilalta-Bufí, M. (2015). Knowledge Misallocation and Growth. *Macroeconomic Dynamics*, *19*(7), 1540–1564.

Restuccia, D., & Rogerson, R. (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics*, *11*(4), 707–720.

Serafinelli, M. (2019). "Good" Firms, Worker Flows, and Local Productivity. *Journal of Labor Economics*, *37*(3), 747–792.

Shi, L. (2023). Optimal Regulation of Noncompete Contracts. *Econometrica*, *91*(2), 425–463.

Sohail, F. (2021). From employee to entrepreneur: Learning, employer size, and spinout dynamics. *Journal of Economic Dynamics and Control*, *133*, 104270.

Sterk, V., Sedláček, P., & Pugsley, B. (2021). The Nature of Firm Growth. *American Economic Review*, *111*(2), 547–79.

Stoyanov, A., & Zubanov, N. (2012). Productivity Spillovers across Firms through Worker Mobility. *American Economic Journal: Applied Economics*, *4*(2), 168–98.

A Finnish Employer–Employee Data

The empirical analysis uses matched employer–employee data from Finnish manufacturing from 1995 to 2012, provided by Statistics Finland. The data track workers across employers and link those mobility histories to establishment-level information on value added, wage bill, materials, employment, and investment. Statistics Finland applies the sampling threshold at the firm level: the employer sample covers manufacturing firms with at least 20 employees, while allowing establishments within those firms to be smaller. I exclude government-owned establishments, special legal forms, and observations with less than one full-time worker. The resulting estimation sample contains roughly 116,000 observations for about 15,000 establishments.

Capital is constructed with a perpetual-inventory method. I winsorize investment at the one percent level, set the depreciation rate to 0.10, and update capital according to $k' = (1 - \delta_k)k + i$. The initial capital stock is $k_{\text{first}} = \max\{i_{\text{first}}/\delta_k, 0\}$.

Labor input is measured with the wage bill rather than headcount. This follows the usual production-function practice of absorbing observable worker-quality differences into labor input. I also quality-adjust the exposure measure for hires from more productive previous employers using a residual-wage weight for each hire. Residual wages come from a log wage regression with gender, age, age squared, education, and year controls. I exponentiate the residual and normalize it so that the average hire has weight one; this normalized weight is $v_{\ell,t-1}$ in (8). The raw-count robustness sets this weight equal to one for all hires.

B Additional Spillover Estimation

This appendix reports a supplemental check based on a source–receiver gap measure close to the productivity-gap approach in Stoyanov and Zubanov (2012). Their firm-level exposure measure uses productivity differences between new hires' sending and receiving firms, averaged over the receiving firm's employment. I construct the establishment-level analogue in the Finnish data by summing lagged productivity gaps for hires whose previous employer is more productive than the receiving establishment, then dividing by lagged employment.

Table 9: Stoyanov-Zubanov spillover measure

	Dependent variable: z_t	
	(1)	(2)
Stoyanov-Zubanov exposure	0.1209	0.1259
	(0.0267)	(0.0268)
Average residual wage of hires	No	Yes
Third-order polynomial in z_{t-1}	Yes	Yes
Observations	100,411	100,411

Notes: Robust standard errors are in parentheses. The Stoyanov-Zubanov exposure is $\sum_{\ell \in H_{i,t-1}} \mathbf{1}\{z^s_{\ell,t-1} > z^r_{i,t-1}\}(z^s_{\ell,t-1} - z^r_{i,t-1})/n_{i,t-1}$, where $z^s_{\ell,t-1}$ is worker ℓ's previous-employer productivity and $z^r_{i,t-1}$ is receiving-establishment productivity. Column (2) adds the average residual wage of new hires as a worker-quality control.

The coefficient is positive in both specifications. This exercise should be read as a supplemental sign check rather than as an alternative augmented control-function estimate. Recovered productivity is positively related to a gap-based exposure measure, matching the qualitative direction in Stoyanov and Zubanov (2012).

C Nature of the Spillovers

The main estimates document a relationship between hiring from more productive establishments and future establishment productivity. This appendix examines whether that relationship is better interpreted as worker-specific human-capital gains or as establishment-level knowledge improvements. The wage-bill measure of labor and the residual-wage weighting address observable worker-quality differences. The remaining concern is that productivity could rise because high-productivity workers replace lower-productivity workers rather than because the establishment absorbs transferable knowledge.

I examine this concern with a descriptive event-study comparison. I take establishments that hire at least one worker from a more productive establishment and compare productivity paths depending on whether that hire leaves immediately in the following period. If the estimated productivity gain were primarily attached to the individual worker, the gain should be weaker when that worker leaves. If the gain reflects establishment-level knowledge, productivity should evolve similarly after the worker leaves.

Figure 5: Productivity after hire from more productive previous employer leaves

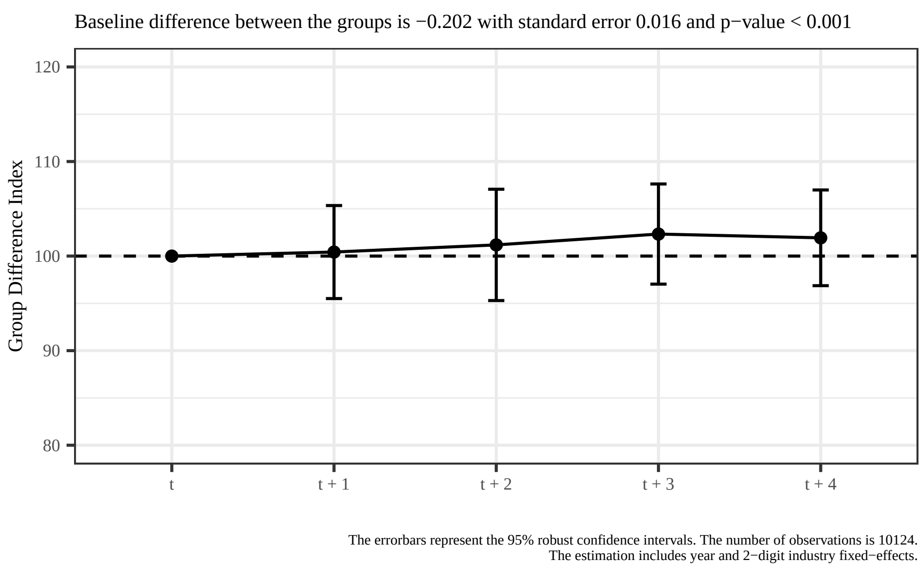

Figure 5 shows no statistically meaningful separation between the two groups. The evidence is descriptive, but it is consistent with the interpretation that the measured spillover is not simply the productivity of an individual worker remaining at the establishment.

Figure 6 repeats the comparison using a broader reference group of establishments that hire from more productive previous employers. The same pattern appears. Productivity paths

Figure 6: Productivity after hiring from a more productive previous employer

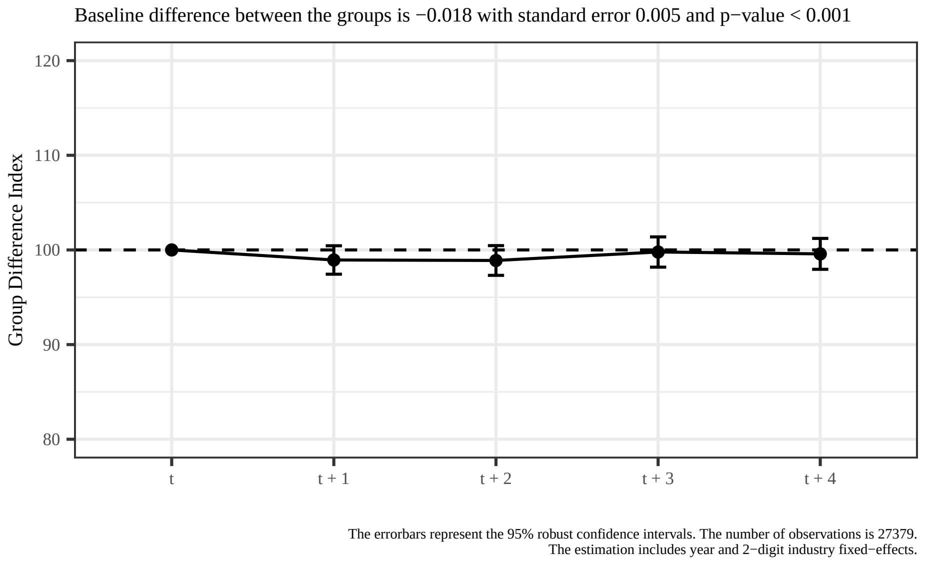

remain similar when that hire leaves.

I also examine whether hiring from more productive previous employers is followed by a shifted or split future-productivity distribution. A stochastic-transfer specification predicts that exposure need not shift every establishment by the same amount. If only some exposed establishments absorb transferable knowledge, the conditional distribution can become bimodal rather than shift uniformly.

Figure 7: Future productivity distribution by hiring from more productive previous employers

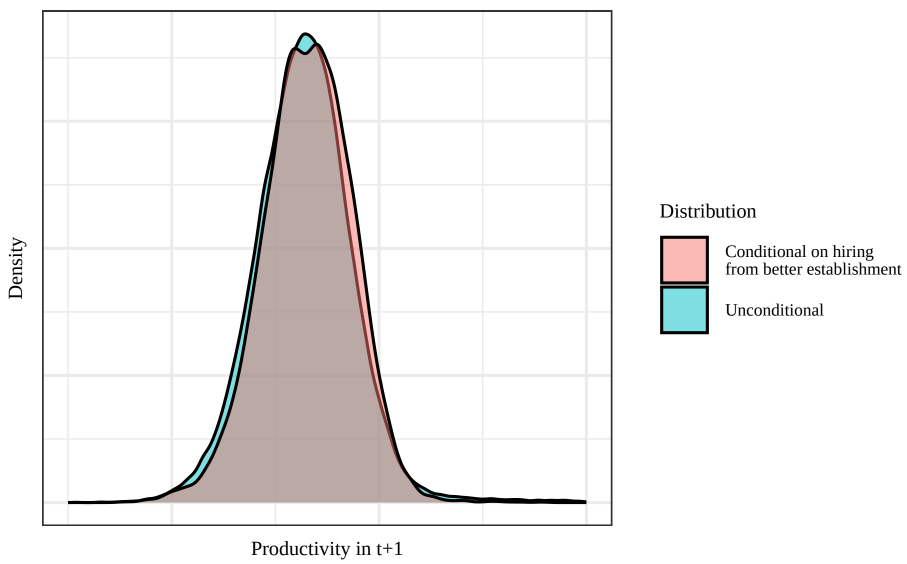

Figure 7 is consistent with a split in the conditional distribution. Conditional on hiring from more productive establishments, next-period productivity is bimodal around the center of the distribution, while the unconditional density is single-peaked. The exercise is not a separate identification strategy, and it is not diagnostic on its own, but it is consistent with a jump-type specification in which hiring from more productive previous employers creates an opportunity for a discrete establishment-level productivity improvement.

D Calibration of the Model without Spillovers

The policy experiments compare the benchmark economy to a separately calibrated economy without worker-transmitted knowledge. This appendix reports that calibration. The calibration sets $\phi = 0$, drops the spillover target, and fits the same core producer-side moments as the benchmark calibration.

Table 10: Calibration and fit without knowledge spillovers

Variable	Value	Explanation	Data	Model
External				
β	0.950	Discount factor		
α	0.700	Labor curvature		
σ_z	0.310	Std. of entrant productivity		
σ_u	0.140	Std. of residual shock		
ρ_H	0.110	Exposure scale parameter		
Internal				
ϕ	0.00000	Set to zero		
f_e	7.28	Inferred from mean size	20.0	20.0
f_f	0.688	Inferred from exit rate	0.0670	0.0660
f_d	4.43	Inferred from job turnover	0.120	0.121
δ	0.138	Inferred from worker turnover	0.300	0.300
κ	0.277	Inferred from aggregate growth rate	0.0256	0.0256
Non-targeted Moments				
Mean hires from more productive previous employers			0.0760	0.136
Mean hires			3.00	3.00
Std. of hires			9.20	40.2
Mean separations			3.30	3.00
Std. of separations			14.0	33.6

Table 10 shows that the economy without knowledge spillovers can match the targeted producer-side moments. The recalibration keeps the policy comparison focused on setting $\phi = 0$ and dropping the spillover target while allowing the remaining parameters to fit the same core moments.

E Exit-Exempt Firing Costs

The firing-cost experiment in Section 4.6.1 treats workers released at exit as a taxed separation. Table 11 reports an alternative convention in which the firing tax applies only to endogenous separations chosen by continuing establishments. The exemption does not remove the exit separation adjustment costs in the baseline problem; it only removes the tax on the post-δ

workforce released by exiting establishments. As in the main experiment, tax receipts are rebated lump-sum and each economy is resolved in general equilibrium.

Table 11: Exit-exempt firing costs: aggregate, allocation, and sender-distribution effects

Effect of exit-exempt firing cost	With spillovers	Recalibrated without spillovers	Difference
Aggregate outcomes			
Output (%)	0.73	1.67	-0.94
Growth (p.p.)	-0.10	-0.08	-0.02
Compensation needed (%)	1.18	-0.22	1.41
Labor-market allocation			
Labor demand (%)	0.12	0.00	0.12
Establishment mass (%)	-0.17	-5.10	4.93
Entry/exit rate (p.p.)	-0.28	-0.41	0.12
Mean size (%)	2.05	6.38	-4.33
Job turnover (p.p.)	-2.73	-0.49	-2.24
Worker turnover (p.p.)	-5.11	-1.65	-3.46
Sender distribution and learning exposure			
Sender productivity ratio	0.133	0.127	0.005
Exposure to hires from more productive previous employers (%)	-18.0	-4.36	-13.6
Successful-transfer probability (p.p.)	-1.21	-0.36	-0.85
Direct gain from successful transfers (p.p.)	-0.03	0.00	-0.03
Policy scale			
Firing-tax payments (% output)	1.03	0.17	0.86

Notes: Entries are policy-minus-baseline changes. The firing cost sets $\tau_f = 1$, as in Table 7, but liquidation separations at exit are exempt from the tax. Continuing establishments pay $w\tau_f s$ on endogenous separations, while exiting establishments still pay the baseline liquidation adjustment cost. Tax receipts are rebated lump-sum to the household. Successful-transfer probability is the exposure-implied arrival probability computed with the same mapping in both economies; when $\phi = 0$, it generates no productivity gain. The direct gain from successful transfers is zero in the economy without spillovers by construction. The difference column subtracts the response in the recalibrated economy without spillovers from the benchmark response.

Under this convention, the firing cost is still costly in the benchmark economy because it reduces knowledge-transmitting worker flows. Growth falls by 0.10 percentage points and the household requires compensation equal to 1.18% of consumption. The entry/exit rate falls by 0.28 percentage points, worker turnover falls by 5.11 percentage points, and exposure to hires from more productive previous employers falls by 18.0%. In the economy without spillovers, the compensation measure is -0.22%, reflecting the resource saving from lower adjustment needs. The exit-liquidation convention therefore matters for the no-spillover welfare comparison, while the benchmark economy continues to lose because the policy reduces hires that can generate knowledge spillovers from more productive previous employers.