Inventor Mobility, Knowledge Diffusion, and Growth*

Yasutaka Koike
UNC Chapel Hill

Toshitaka Maruyama
Bank of Japan

Koki Okumura
Chinese University of Hong Kong

June 16, 2026

Abstract

We study how inventor mobility shapes knowledge diffusion, economic growth, and welfare. Using German patent data linked to employer-employee records, we show that inventor mobility is associated with knowledge diffusion. Motivated by this evidence, we develop and calibrate an endogenous-growth model in which inventors contribute through both internal R&D and mobility across firms. We use the model to evaluate restrictions on non-compete clauses and the macroeconomic consequences of declining inventor mobility. Banning non-compete clauses raises annual growth by 0.036 percentage points, but it also increases vacancy posting and lowers welfare by 0.38%. We also find that the observed decline in inventor mobility reduces growth by 0.25 percentage points, with 0.05 points coming from weaker internal R&D and 0.19 points from weaker knowledge diffusion. The mobility decline also lowers consumption-equivalent welfare by 3.49%.

JEL Classification: O33, O31, O40, J62.

*We thank Andy Atkeson, David Baqaee, Ariel Burstein, Levi Crews, Pablo Fajgelbaum, Gary Hansen, Ryo Horii, Hugo Hopenhayn, Seula Kim, Lee Ohanian, Koki Oikawa, Nik Engbom, Liyan Shi, Jon Vogel, Pierre-Olivier Weill, seminar participants at UCLA and Osaka University, and conference participants at SWET 2023 and the Spring 2026 Midwest Macroeconomics Meeting for their valuable comments. We also thank the Institute for Employment Research (IAB) for providing data access. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Bank of Japan. This paper was awarded the 2023 Moriguchi Prize by the Institute of Social and Economic Research at Osaka University. Corresponding author: Koki Okumura, Department of Economics, The Chinese University of Hong Kong, 9/F & 10/F, Esther Lee Building, Shatin, N.T., Hong Kong; k.okumura1201@gmail.com. Total word count: 15,983.

1 Introduction

Inventors advance technology in two ways: they create knowledge within firms, and they carry knowledge across firms when they change employers. As Arrow (1962) noted, the "mobility of personnel among firms provides a way of spreading information," highlighting that the movement of inventors across firms has long been considered a key factor in inter-firm knowledge diffusion.[1] This dual role places inventor mobility at the center of two debates. The first is about policy. Regulations on non-compete clauses directly shape inventors' ability to move. A classic argument is that California's refusal to enforce non-compete clauses, relative to Massachusetts, facilitated labor mobility and knowledge spillovers and thereby contributed to Silicon Valley's rise relative to Route 128 (Gilson, 1999). Recent U.S. policy debates over non-compete clauses have renewed the question of whether restricting these clauses spurs innovation.[2] The second debate concerns a secular trend: inventor mobility has fallen sharply. In German administrative data, the annual job-to-job transition rate of patent-listed inventors declined from 12.6% in 2001 to 4.8% in 2014, and Akcigit and Goldschlag (2025) document a related decline in U.S. inventor hiring.

Evaluating either the policy question or the secular trend requires a framework in which an inventor's move does two things at once: it reallocates inventive labor across firms and transmits knowledge from the previous employer to the new one. In many endogenous-growth models, knowledge diffuses through random meetings or interactions rather than through employment transitions. Conversely, frictional labor-market models with multi-worker firms generate rich job-to-job flows, but typically do not treat these flows as a direct channel of inter-firm knowledge diffusion. To our knowledge, the effects of policies that shape inventor mobility, such as regulations on non-compete clauses, have not been studied in a structural endogenous-growth framework in which inter-firm inventor mobility itself operates as a direct channel of knowledge diffusion. This paper develops such a framework, disciplines it with German microdata that link patents to matched employer-employee records, and uses it for policy evaluation. Two quantitative findings summarize the analysis. First, banning non-compete clauses raises annual growth by 0.036 percentage points, but lowers consumption-equivalent welfare by 0.38% because the same policy triggers a costly surge in vacancy posting. Second, the observed decline in inventor mobility reduces annual growth by 0.25 percentage points, about 13% of its baseline level,

[1] For evidence from recent studies, see Jaffe et al., 1993; Almeida and Kogut, 1999; Song et al., 2003; Hoisl, 2007; Rosenkopf and Almeida, 2003; Breschi and Lissoni, 2009; Singh and Agrawal, 2011; Kaiser et al., 2015; Rahko, 2017; Braunerhjelm et al., 2020.

[2] See Federal Trade Commission (2024). The Federal Trade Commission's 2024 rulemaking on non-compete clauses is no longer in effect; our quantitative exercise below studies a ban on non-compete clauses rather than an evaluation of a specific legal rule.

and welfare by 3.49%, with most of the growth loss flowing through weaker knowledge diffusion rather than through the allocation of inventors.

We begin with three empirical facts. Our main data set, INV-BIO, links the 1980–2014 labor-market biographies of 152,350 inventors listed on EPO patent filings between 1999 and 2011, recorded in German social security data, to patent-register information. We complement it with SIAB, a 2% random sample of the same administrative universe, in which we construct occupation-based inventor-proxy workers. First, inventor mobility is not predominantly up the ladder. In transition matrices constructed in the spirit of Haltiwanger et al. (2018), a large share of inventor job-to-job moves lands at establishments ranked similarly to or below the origin: 36% of moves go to establishments in lower mean-wage rank bins, there is no clear upward pattern when establishments are instead ranked by citation productivity, and these patterns remain when we restrict attention to moves in which the inventor's own wage rises. Second, workers in inventor-intensive occupations differ from other job-to-job movers. Conditional on worker characteristics and origin-establishment, destination-establishment, and year fixed effects, inventor-proxy workers are significantly more likely than other movers to move to lower-mean-wage establishments, yet they experience larger wage gains when moving.

Third, establishment-level regressions show that inventor inflows carrying higher lagged source inventor productivity are associated with higher incumbent inventor-productivity growth at the destination establishment over the move interval. The first two facts are difficult to reconcile with standard firm-wage job-ladder models, in which job-to-job mobility reallocates workers toward more productive, higher-wage employers; they are instead consistent with hiring firms compensating mobile inventors for the knowledge they bring. The third fact provides establishment-level evidence consistent with the knowledge-diffusion channel and supplies a key empirical moment for the quantitative model. A complementary source-shock exercise, which allocates severe source-employment contractions through lagged inventor hiring networks, points in the same direction and mitigates the concern that the relationship is driven only by contemporaneous destination choices.

Motivated by this evidence, we develop an endogenous-growth model in which inventors contribute to growth both through internal R&D and through the knowledge they transport when they switch employers. As in Klette and Kortum (2004), firms operate a portfolio of product lines, which we identify with their knowledge capital k, and employ inventors n, with internal R&D exhibiting constant returns in (k, n). Inventors and firms meet in a frictional labor market with random on-the-job search. Firms post vacancies anticipating not only the marginal value of an additional inventor but also the knowledge

spillover from the hire's previous employer. When a meeting generates positive joint surplus, the inventor moves, and a share β of the surplus accrues to the coalition of the poached firm and the inventor, with the remainder going to the poaching firm. An inventor's move therefore affects growth through two margins: it reallocates inventors across firms, which changes internal-R&D growth, and it transfers knowledge from source to destination firms, which contributes directly to creative destruction. Both the inter-firm network of inventor flows and the associated diffusion of knowledge are equilibrium objects that respond to economic conditions and policy.

The model clarifies when mobility looks like a job ladder and when it does not. Absent knowledge diffusion, the search margin collapses to standard job-ladder logic: inventors move only toward firms with higher marginal inventor value, equivalently higher inventor productivity $z = k/n$ (Proposition 1). With diffusion, a destination firm may profitably hire an inventor even when the move is not upward, because the hire also transfers the source firm's knowledge. The downward moves we document in the data therefore motivate a model in which inventor mobility carries knowledge across firms.

The model remains tractable because the product-line mass is continuous. The joint value of a firm and its inventors is homogeneous of degree one in (k, n), so a single state variable, inventors per product line $x = n/k$, summarizes firm heterogeneity even though inventor employment is costly to adjust. The same homogeneity delivers a strong version of Gibrat's law and lets us focus on job-to-job transitions among incumbent firms without modeling entry and exit.

Comparing the competitive equilibrium with the planner allocation reveals how search frictions, product-market competition, internal R&D, and knowledge diffusion interact to drive wedges between private and social incentives. Innovation incentives reflect both appropriability and business-stealing forces, while vacancy posting is shaped by surplus sharing and congestion. Because firms are heterogeneous, the Hosios-like benchmark for the search margin is state-dependent, so no constant bargaining weight decentralizes the planner allocation. We therefore treat the bargaining weight β, which maps naturally into the enforceability of non-compete clauses as in Shi (2023), as a second-best policy instrument.

We calibrate the model to the German data. The model implies a law of motion for establishment-level inventor productivity in which incumbent productivity growth responds to inventor inflows weighted by source inventor productivity and to internal R&D at the destination. The inflow term is positive across increasingly demanding specifications and supplies the empirical moment used to discipline the scale of knowledge diffusion, while the internal-R&D term disciplines the scale of in-house innovation. Matching

efficiency then targets the 12.6% inventor job-to-job transition rate in 2001, and the innovation step size targets a 2% annual growth rate. In the calibrated economy, knowledge diffusion accounts for 16.2% of aggregate growth.

Our first quantitative exercise evaluates a ban on non-compete clauses, modeled as a permanent, unanticipated fall in β from its baseline value of 0.5 to 0. The ban raises long-run annual growth by 0.036 percentage points, mainly through stronger knowledge diffusion. Welfare nonetheless falls because the same increase in the poacher's surplus share that fuels mobility also induces costly vacancy posting. The initial consumption loss outweighs the later gains from faster growth, leaving consumption-equivalent welfare 0.38% below the no-policy-change path. Faster growth and higher welfare thus need not coincide.

Second, although no fixed bargaining split implements the planner allocation, a near-symmetric split comes close to it. The welfare gain from implementing the planner allocation is nearly zero around $\beta = 0.52$, consistent with a state-dependent Hosios-like benchmark that stays close to one half over much of the inventor-weighted distribution.

Third, we quantify the macroeconomic consequences of the decline in inventor mobility. We choose a path of matching efficiency that replicates the German fall in the inventor transition rate from 12.6% in 2001 to 4.8% in 2014. This transition lowers long-run annual growth by 0.25 percentage points, about 13% of the initial 2% growth rate. Most of the growth loss comes from weaker inter-firm knowledge diffusion, with the remainder reflecting slower internal-R&D growth from a worse allocation of inventors. Consumption-equivalent welfare falls by 3.49% relative to a balanced growth path in which matching efficiency does not deteriorate.

Taken together, the results show that inventor mobility is not merely a labor-market outcome but a channel through which knowledge is reallocated across firms. Policies that affect mobility must therefore account for knowledge diffusion, the allocation of inventive labor, and recruiting resource costs, which can move growth and welfare in different directions.

Related Literature

Our paper is, to our knowledge, the first to model inventor mobility between firms as a direct channel of inter-firm knowledge diffusion and to embed that mechanism in an endogenous-growth framework. This contribution places the paper closest to the endogenous-growth literature on technology and knowledge diffusion, including Kortum (1997), Lucas and Moll (2014), Perla and Tonetti (2014), Akcigit et al. (2018), Buera and Oberfield (2020),

Benhabib et al. (2021), Prato (2025), Crews (2023), Hopenhayn et al. (2026), and Babalievsky (2023).[3] Similar to our study, Benhabib et al. (2021) and Hopenhayn et al. (2026) analyze the interaction between internal R&D activity and knowledge diffusion between agents. A contemporaneous paper by Babalievsky (2023) develops a Schumpeterian growth model with a frictional market for inventors and studies how search frictions and misallocation in inventor-firm matching affect growth and innovation policy. Within this literature, our work is also closely aligned with Akcigit et al. (2018), Prato (2025), and Crews (2023), which model knowledge diffusion among inventors or high-skilled workers. These papers do not explicitly account for inventor mobility between firms and the associated job-to-job transitions, and they do not use such flows to evaluate non-compete policy or the consequences of declining inventor mobility.

The idea that inventor mobility between firms facilitates knowledge diffusion is supported by numerous empirical studies. One of the early studies in this literature, Almeida and Kogut (1999), demonstrates that in regions with high levels of inventor mobility between firms, knowledge flows tend to be localized. Song et al. (2003) show that inventors who switch firms are more likely to base their work on ideas from their previous employer compared to other inventors at the hiring firm. Rosenkopf and Almeida (2003) analyze pairs of firms and find that firms with higher inventor mobility also experience greater subsequent knowledge flows. These pioneering studies have spurred further research into the relationship between inventor mobility and knowledge spillovers (Hoisl, 2007; Breschi and Lissoni, 2009; Singh and Agrawal, 2011; Kaiser et al., 2015; Rahko, 2017; Braunerhjelm et al., 2020).[4] We also make an empirical contribution to this literature. By comparing the patterns of job transitions and the accompanying changes in wages between inventors and other workers, we provide evidence suggestive of inter-firm knowledge diffusion and compensation for it.

Our paper is complementary to recent studies that analyze learning among workers within firms. Jarosch et al. (2021) and Herkenhoff et al. (2024) develop frameworks for learning from colleagues within firms and, using matched employer-employee data, find significant knowledge spillovers among coworkers. We differ by focusing on inventors engaged in R&D, linking their mobility to aggregate growth through an endogenous-growth

[3] Buera and Lucas (2018) provide a survey of this literature.

[4] Mawdsley and Somaya (2016) provide a review of this literature. A related area of study is the relationship between geography and knowledge diffusion. Early research by Jaffe et al. (1993) suggested a higher probability of cited patents originating from the same location as the citing ones. Breschi and Lissoni (2009) refined this approach by introducing inventor mobility as a control, revealing that the effect of spatial proximity on knowledge diffusion is reduced by more than half. This finding suggests that the critical role of geography in knowledge diffusion primarily stems from the infrequency of inventors relocating across regions.

model, and using data that connect patents to matched employer-employee records.

Finally, our paper is related to the literature on frictional labor markets with multi-worker firms and on-the-job search, including Schaal (2017), Elsby and Gottfries (2021), Lentz and Mortensen (2022), Bilal et al. (2022), and McCrary (2022). We build on this approach to study how job-to-job mobility reallocates inventors and transmits knowledge across firms. Closely related, Bilal et al. (2023) embed on-the-job search in an endogenous-growth model, but their framework does not model knowledge diffusion through worker or inventor mobility.

The rest of the paper proceeds as follows. Section 2 introduces the data and presents empirical findings. Section 3 presents the theoretical framework: the competitive equilibrium, the social planner's problem, and the associated externalities and policy instruments. Section 4 estimates the model using the data and discusses its fit. Section 5 uses the estimated model to examine the impact of the decline in inventor mobility and to evaluate the quantitative policy counterfactuals. Section 6 concludes.

2 Data and Empirical Findings

In this section, we document inventor mobility and related evidence on knowledge diffusion using two German administrative data sets. We first use "Linked Inventor Biography Data 1980–2014" (INV-BIO) to describe job flows of patent-listed inventors between establishments. We then use the "Sample of Integrated Labor Market Biographies" (Stichprobe der Integrierten Arbeitsmarktbiografien—SIAB) to compare occupation-based inventor-proxy workers with other workers.[5]

The INV-BIO data set combines labor market biographies recorded in the German social security data (Integrated Employment Biographies — IEB) with patent-register data from PATSTAT at the European Patent Office (EPO), and supplements these patent records with DPMAregister. This data set tracks labor market information from 1980 to 2014 for 152,350 inventors identified from EPO patent filings between 1999 and 2011. Patent-based measures are constructed from patent filings/applications rather than only from granted patents. For inventors, the data include a unique ID, age, gender, education, wage, and the number of citations received by patents associated with each inventor. The data also contain information about the establishments employing the inventors, such as the establishment ID, the total number of their employees, and the mean wage of their full-time employees. An important advantage over patent-based data sets used in previous studies, such as the

[5]More detailed information is presented in Appendix A.

EPO patent data in Akcigit et al. (2018), is that we can use social security information to track inventors' flows even in years when they do not appear on new patent filings.

The SIAB data set is a 2% random sample from IEB (SIAB 7519, v1; Frodermann et al., 2021; see also Dauth and Eppelsheimer, 2020). It contains the same information about individuals and their employing establishments as INV-BIO, except for patent-related information. In the absence of patent data, we construct inventor-proxy workers in SIAB using broad KldB 2010 parent occupation groups, as described in Section 2.2. The underlying SIAB version covers 1975–2019; our analysis sample covers 3,322,316 individuals from 1980 to 2019.

Because the two data sets cannot be linked at the individual level, we use them separately in the analyses below. When comparing the movement patterns of inventor-proxy workers and other workers in Section 2.2, we rely on SIAB, which includes occupation-based inventor proxies and non-inventor workers. Otherwise, we use INV-BIO, a data set focused exclusively on inventor information.

2.1 Inventor Flows in INV-BIO

We characterize inventor flows in INV-BIO using transition matrices constructed in the spirit of Haltiwanger et al. (2018). We assign each establishment to a percentile rank based on one of two establishment-level measures. We then compute the transition probabilities of inventor flows between these ranks and present them as row-normalized transition matrices, following the format in Albagli et al. (2025).[6]

We use two establishment-level measures related to knowledge quality and productivity: (i) patent citations per inventor and (ii) the mean wage of full-time employees.[7] The first measure is based on forward citations to patents associated with inventors employed at each establishment. Measuring patent quality through forward citations is widely employed in the literature on patent creation (e.g., Pakes, 1986; Hall et al., 2001; Akcigit et al., 2018). To construct this measure, we first define inventor-level productivity. Let I_{jt} denote the set of patents associated with inventor j in year t. We define inventor j's productivity in year t as:

$$p_{jt} = \sum_{i \in I_{jt}} \frac{\text{citations}_{it}}{n_i},$$

[6]Establishments can fall into different percentiles each year depending on the measure used. The ranks of the origin and destination establishments are determined by the measure from the previous year, preceding the movement of inventors.

[7]On-the-job search models with heterogeneous-productivity firms (e.g., Postel-Vinay and Robin, 2002) predict that more productive firms offer higher wages and attract more workers.

where citations$_{it}$ denotes the three-year forward citations received by patent i, measured using EPO DOCDB citation variables and summed over the DE, EP, and US jurisdiction-specific citation variables, and n_i is the total number of inventors on patent i. Thus, each patent's citations are divided equally among its co-inventors. Let J_{et} denote the set of inventors employed at establishment e in year t. We then define establishment-level inventor productivity as the average inventor productivity within the establishment:

$$z_{et} = \frac{1}{|J_{et}|} \sum_{j \in J_{et}} p_{jt}.$$

Because citations are first divided equally among co-inventors, patents coauthored across establishments contribute to z_{et} only through the inventors employed at establishment e. This empirical object is the data counterpart of firm-level inventor productivity $z = k/n$ in the model below.

For the transition matrices, we rank establishments using the three-year backward average of annual citation productivity:

$$\bar{z}_{et}^{(3)} = \frac{1}{3}\left(z_{et} + z_{e,t-1} + z_{e,t-2}\right).$$

We define this average only when it is positive. For inventor moves between $t-1$ and t, origin and destination ranks use previous-year establishment measures, including $\bar{z}_{e,t-1}^{(3)}$ for citation productivity.

Table 1 shows row-normalized transition matrices of inventor flows from origin to destination rank bins. The matrices show substantial mass on and below the diagonal, indicating that many inventors move to establishments that are similar to or lower than their origin establishment in the relevant rank distribution. In the mean-wage ranking, 36% of patent-listed inventor job-to-job flows go to lower destination rank bins. Appendix A.2 shows that this pattern persists even when we restrict the sample to job-to-job flows accompanied by wage increases.

This pattern is difficult to reconcile with standard firm-wage job-ladder models, in which job-to-job mobility tends to reallocate workers toward more productive and higher-wage employers. The lack of a clear upward pattern in both the citation-productivity and mean-wage rankings suggests that inventor moves are shaped by additional forces, consistent with the possibility that hiring establishments value the knowledge carried by mobile inventors.

Table 1: Transition Matrices of Inventor Job-to-Job Flows (Percent)

(A) Rank by Citation Productivity

	Destination rank bin									
Origin rank bin	1	2	3	4	5	6	7	8	9	10
1	12	12	11	10	8	9	9	9	10	10
2	12	12	12	9	9	9	9	9	10	10
3	10	12	12	11	9	10	9	9	9	9
4	11	10	11	11	11	10	11	9	9	8
5	9	9	10	9	10	11	15	10	9	10
6	8	8	10	7	12	9	9	12	11	13
7	6	8	6	14	12	9	12	11	11	12
8	5	6	7	7	8	12	12	14	14	14
9	4	5	10	6	12	8	9	16	13	17
10	5	7	6	6	6	10	10	11	15	25

(B) Rank by Mean Wage

	Destination rank bin									
Origin rank bin	1	2	3	4	5	6	7	8	9	10
1	18	16	13	12	10	9	7	6	5	4
2	11	16	14	12	10	9	8	8	7	6
3	7	13	14	12	10	9	9	9	8	8
4	6	10	11	12	11	11	11	10	10	8
5	5	8	9	11	11	11	10	12	11	12
6	4	7	9	9	10	12	13	12	12	11
7	4	6	7	8	9	11	14	13	14	15
8	3	5	6	7	10	11	14	13	14	17
9	2	5	5	6	9	9	14	14	16	21
10	2	3	4	4	6	7	10	12	19	32

Notes: Each panel reports row-normalized inventor job-to-job flows from origin rank bins to destination rank bins, in percent. Rows sum to 100 up to rounding. Job-to-job moves require an establishment change with no intervening unemployment spell; inventors who remain with the same establishment are excluded. Establishments are ranked within year by the panel variable, and rank bins are formed using cumulative inventor counts. Panel (A) ranks establishments by the positive three-year backward average of citation productivity; panel (B) uses establishment mean wage. Origin and destination ranks use previous-year establishment measures. The sample covers 1980–2014.

2.2 Inventor-Proxy Worker Flows Compared with Other Workers

In this section, we compare the job-to-job flows of inventor-proxy workers with those of other workers. We use SIAB for this comparison because INV-BIO lacks information on non-inventor workers.

To construct inventor-proxy workers within SIAB, we use broad KldB 2010 parent occupation groups. The five largest parent occupation groups among inventors in INV-BIO are groups 25, 27, 41, 26, and 71, which together account for 80.5% of inventor-years with observed occupations. We define inventor-proxy workers as SIAB workers in these groups who also have a university-entrance qualification.

Table 2 compares summary statistics across the two data sets. The mean wage of the inventor-proxy workers in SIAB lies between the average wage of all SIAB workers and that of inventors in INV-BIO. Because the proxy is occupation-based, we interpret the

Table 2: Inventor-Proxy Workers in SIAB and Inventors in INV-BIO

Summary statistics		SIAB		INV-BIO
(1980–2014)		Workers	Inventor-proxy workers	Inventors
Wage, Euro	Mean	63.8	99.5	123.8
	S.D.	45.5	55.6	52.1
Age	Mean	38.5	39.0	38.9
	S.D.	12.6	11.2	10.6
Females, %		44.2	33.7	6.4
University-entrance qualification, %		15.3	100.0	70.3
N of observations, thousand		17,474	1,030	3,517

Notes: This table compares workers in SIAB, inventor-proxy workers in SIAB, and inventors in INV-BIO. Inventor-proxy workers in SIAB are workers whose observed KldB 2010 parent occupation group is in groups 25, 26, 27, 41, or 71 and who have a university-entrance qualification. The workers column includes inventor-proxy workers. The SIAB columns use pooled BeH employment person-year observations with wage, age, gender, occupation, and education observed. The INV-BIO column uses pooled inventor-year observations with wage, age, and gender observed. The education row additionally requires the schooling variable to be observed.

comparison between inventor-proxy workers and other SIAB workers as conservative: measurement error in the proxy would tend to attenuate the estimated differences.

We estimate the following fixed-effects linear regression for SIAB job-to-job movers with no intervening unemployment spell:

$$Y_{it} = \beta I_{i,t-1} + X'_{i,t-1}\gamma + \alpha_{o(i)} + \alpha_{d(i)} + \alpha_t + \varepsilon_{it}, \tag{1}$$

where Y_{it} is either an indicator for moving to a destination establishment with a lower previous-year mean wage than the origin establishment or the worker's wage growth. Specifically, $D_{it}^{\text{wage}} = 1\{\bar{w}_{d(i),t-1} < \bar{w}_{o(i),t-1}\}$, and $\Delta \log w_{it} = \log w_{it} - \log w_{i,t-1}$. The indicator $I_{i,t-1}$ equals one if individual i is an inventor-proxy worker in the previous year. The vector $X_{i,t-1}$ includes age, age squared, gender, and log wage, all measured before the move. We absorb origin establishment, destination establishment, and year fixed effects.

Our coefficient of interest is β. A positive value in the lower-mean-wage-destination regression indicates that inventor-proxy workers are more likely than other workers to move to lower-rung establishments in the firm-wage distribution.

Table 3 reports the estimation results. Column (1) uses establishment mean wage to construct the downgrade outcome, while column (2) studies wage growth directly. The positive coefficient in column (1) shows that, conditional on the controls and fixed effects, inventor-proxy workers are significantly more likely than other SIAB job-to-job movers to move to lower-mean-wage establishments. This result is difficult to reconcile with the standard firm-wage job-ladder prediction that job-to-job mobility reallocates

Table 3: Estimation Results for Inventor-Proxy Worker Flows

	Full job-to-job sample	
	(1)	(2)
Dependent variable	D_{it}^{wage}	$\Delta \log w_{it}$
$I_{i,t-1}$	0.0080***	0.0581***
	(0.0021)	(0.0033)
Controls	Yes	Yes
Origin estab. FE	Yes	Yes
Dest. estab. FE	Yes	Yes
Year FE	Yes	Yes
Observations	651,554	651,554
Adjusted R^2	0.5603	0.5288

Notes: Column (1) reports a fixed-effects linear probability model. Column (2) reports fixed-effects OLS estimates of wage changes. D_{it}^{wage} equals one if worker i's destination establishment has lower previous-year mean wage than the origin establishment; ties are coded as zero. $I_{i,t-1}$ equals one if worker i was an inventor-proxy worker in the previous year. Controls are previous-year age, age squared, gender, and log wage. All columns absorb origin establishment, destination establishment, and year fixed effects. The sample consists of job-to-job moves between establishments with no intervening unemployment spell. Observations are reported after singleton drops implied by the absorbed fixed effects. Standard errors are two-way clustered by origin and destination establishment and reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

workers toward higher-wage employers. The positive wage-growth coefficient in column (2) indicates that inventor-proxy workers experience larger wage gains when changing jobs, so moves to lower-mean-wage establishments need not represent downward individual mobility; they are consistent with compensation for knowledge transfer. Table A.2 shows that the lower-mean-wage-destination result also holds in the subsample of moves accompanied by an increase in the worker's own wage. In that subsample, the coefficient on inventor-proxy status is 0.0094 for moving to a lower-mean-wage establishment and 0.0535 for wage growth.

Knowledge transfer provides a natural interpretation of these patterns. Inventors who leave high-inventor-productivity establishments bring ideas that can raise inventor productivity elsewhere, making lower-inventor-productivity establishments willing to recruit them and compensate them accordingly. Taken together, downward moves in establishment rankings and wage gains for mobile inventor-proxy workers motivate a search model in which inventors carry transferable knowledge across firms.

In Section 4 and Appendix C.1, we use the linked patent information in INV-BIO to discipline this inventor-productivity channel empirically by relating inventor inflows to incumbent inventor-productivity growth at the destination establishment.

3 Model

In this section, we first introduce an endogenous-growth model that highlights the role of inventor mobility and the resulting knowledge diffusion, and we characterize the competitive equilibrium. A central feature of the model is that an inventor move changes both the allocation of inventor labor and the distribution of knowledge capital across firms. Accordingly, mobility affects growth through two margins: it reallocates inventors across firms, which changes internal-R&D growth, and it transfers knowledge across firms, which contributes directly to creative destruction. Next, we formulate the social planner's problem and examine the externalities present in this economy.

3.1 Environment and Technologies

Household and Production Technology

Time is continuous. The representative household is composed of a unit measure of individuals who supply L units of production labor and N units of inventor labor inelastically. The representative household consumes a single final good, and the utility at time 0 is given by

$$\int_0^\infty e^{-\rho t} \log C(t) dt,$$

where $\rho > 0$ is the discount rate. The representative household can borrow and lend at interest rate $r(t)$, and $C(t)$ denotes aggregate consumption at time t. We choose the numeraire so that $P(t)Y(t) = 1$ for all t, where $P(t)$ denotes the final good price and $Y(t)$ denotes the aggregate output.[8] Along a BGP under this numeraire, $r(t) = \rho$.[9]

There is a unit continuum of products used to produce the final good, and the aggregate output is determined by the production function

$$\log Y(t) = \int_0^1 \log \left(a(\omega, t) l(\omega, t) \right) d\omega,$$

where $l(\omega, t)$ is the quantity of product $\omega \in [0, 1]$ at time t, and $a(\omega, t)$ is the productivity of product ω at time t.

Production labor is the only factor in the production of each product. There are no frictions in the production-labor market. Labor productivity is the same for all products and is set equal to 1. This means that $l(\omega, t)$ represents both the output amount of product

[8] The choice of the numeraire implies that $P(t)$ declines over time at the rate of economic growth.
[9] The derivation is in Appendix B.1.

ω and the labor demand for the production of product ω. Each product is supplied by a monopolist. We assume that firms pay an arbitrarily small operating cost $\epsilon \to 0$ before producing a product each period, as in Acemoglu et al. (2018), to prevent price competition between the technology leader of the product and firms that previously produced that product. Let $p(\omega, t)$ denote the price of product ω at t. Because the nominal output $P(t)Y(t)$ is normalized to 1 and the final good technology is Cobb–Douglas, the revenue for each product line $p(\omega, t)l(\omega, t)$ is also equal to 1.[10]

We assume that other firms can copy and use a product-line technology by paying a cost $0 < \pi < 1$. This copying option bounds the profit that an incumbent firm can obtain from a single product line by π.[11][12] A monopoly firm sets the price $p(\omega, t)$ to solve the following problem:

$$\begin{aligned} \max_{p(\omega,t)} \ & \pi(\omega, t) = p(\omega, t)l(\omega, t) - w(t)l(\omega, t) \\ \text{s.t.} \quad & \pi(\omega, t) \leq \pi, \\ & p(\omega, t)l(\omega, t) = 1. \end{aligned}$$

The firm's optimal price is $p(\omega, t) = w(t)/(1 - \pi)$, and labor demand is $l(\omega, t) = (1 - \pi)/w(t)$. From the labor demand and labor market clearing condition $\int_0^1 l(\omega, t)d\omega = L$, the production wage is given by $w(t) = (1 - \pi)/L$. Therefore, we have $p(\omega, t) = 1/L$ and $l(\omega, t) = L$.

Innovation Technology

The quality of each product is determined by the total number of innovations that have been implemented for that product in the past:

$$a(\omega, t) = \lambda^{m(\omega,t)},$$

where $m(\omega, t)$ is the number of innovations made to product ω up to time t, and $\lambda > 1$ is the step size of innovation, assumed to be the same across all products. All products face the same rate of creative destruction, denoted by $\delta(t)$, which is endogenously determined

[10] The derivation is in Appendix B.2.

[11] We consider the size of π to be exogenously determined by factors such as the strength of patent enforcement or taxation on sales.

[12] In a quality ladder model with a Cobb–Douglas aggregator, it is common to assume Bertrand competition between the firm currently producing the product and the firm that previously produced it. In this case, when the step size of innovation is constant, the instantaneous profit for each product line is equal across all product lines, as in the result derived here.

by the aggregation of innovation activities. By the law of large numbers,

$$\log Y(t) = \log \lambda \int_0^t \delta(\tau) d\tau + \log L. \tag{2}$$

Therefore, the economic growth rate is proportional to the creative destruction rate:

$$\frac{d}{dt} \log Y(t) = \delta(t) \log \lambda. \tag{3}$$

Each firm supplies a continuum of products, denoted by k, and employs a continuum of inventors, denoted by n. The mass of a firm's product lines k changes for three reasons:

(i) Internal R&D. The mass of a firm's product lines k changes as a result of internal R&D activities:

$$\left(\mathrm{d} \log k\right)_{R\&D} = \left\{\mu\left(n/k\right) - \frac{\sigma^2}{2}\right\} dt + \sigma dW(t),$$

where $dW(t)$ is a Wiener process, $\mu(\cdot)$ is increasing and concave with $\mu(0) \geq 0$, and $\sigma > 0$. Under this assumption on $\mu(\cdot)$, the drift of R&D outcome per inventor $\mu\left(n/k\right) k/n$ is increasing in k/n. We therefore define $z \equiv k/n$ as firm-level inventor productivity and use its reciprocal, $x \equiv n/k = z^{-1}$, as the firm state below. The stochastic process of k due to internal R&D assumed here is closely related to Klette and Kortum (2004), because the expected R&D outcome $\mu\left(n/k\right) k$ exhibits constant returns to scale in the knowledge capital k and the mass of inventors n. Klette and Kortum (2004) assume that a firm's knowledge capital facilitates innovation, and that knowledge capital can be summarized by the number of a firm's product lines. We adopt the same assumption and refer to k interchangeably as knowledge capital or the mass of product lines. The key difference from Klette and Kortum (2004) is that in our model, k takes a continuous value, whereas in their model, k takes an integer value. Furthermore, as we will discuss later, unlike the model of Klette and Kortum (2004), there is a cost to adjusting the mass of inventors n due to labor market frictions.

(ii) Knowledge diffusion due to inventor job-to-job flows. When inventors flow into a firm with inventor productivity $z = k/n$ from a source firm with inventor productivity $z' = k'/n'$, the knowledge capital of the destination firm increases by $\hat{\alpha}(z, z')$ per inventor inflow:

$$\frac{\left(\mathrm{d}k\right)_{\text{from } (k',n')\text{-firm}}}{\left(\mathrm{d}n\right)_{\text{from } (k',n')\text{-firm}}} = \hat{\alpha}(z, z'),$$

where $\hat{\alpha}(z, z')$ is an exogenous function that maps source and destination inventor productivities into the amount of knowledge capital transferred per inventor inflow. This

formulation allows the amount of transferable knowledge to vary with the characteristics of both the poaching firm and the poached firm.

(iii) Creative destruction by other firms. Each product line of a firm faces the possibility that other firms will innovate on it. If a firm has k product lines, it will lose product lines at a rate of $\delta(t)k$:

$$\left(\mathrm{d}\,\log k\right)_{\text{Destruction}} = -\delta(t)dt.$$

Matching Technology and Bargaining

Firms and inventors meet in a single frictional labor market, and search is random. We focus on on-the-job search by assuming that unemployment benefits are sufficiently low that inventors do not voluntarily choose unemployment. As in Kaas and Kircher (2015), we assume a vacancy cost that is homogeneous of degree one in the number of vacancies and the firm's knowledge capital: a firm pays a cost $c\left(\hat{v}/k\right)k$ to post $\hat{v}$ vacancies. The vacancy cost is denominated in the numeraire. The cost function $c\left(\cdot\right)$ is increasing and convex, and satisfies $c(0) = 0$ and $c'(0) = 0$.[13]

An aggregate matching function $m(N, V)$ determines the flow of meetings between firms and inventors as a function of the measure of inventors, N, and the total measure of vacancies, V. Each firm meets an inventor at a Poisson rate proportional to the number of vacancies it posts. Each inventor meets a firm at a Poisson rate identical across all inventors. We assume that m exhibits constant returns to scale in (N, V) and decreasing returns to scale in either N or V separately. Let $\theta = V/N$ denote labor market tightness. Then, the Poisson rate at which a unit of vacancy meets an inventor is $q(\theta) \equiv m(N, V)/V$, and the rate at which an inventor meets a firm is $\theta q(\theta) = m(N, V)/N$.

Our assumptions on bargaining and vacancy posting follow Lentz and Mortensen (2022). When a firm and an inventor meet, a match is formed if the match surplus among the three parties (the poaching firm, poached firm, and inventor) is positive. In this case, the three parties bargain over the match surplus. As a result of bargaining, the share $\beta \in (0, 1)$ of the surplus is allocated to the coalition of the poached firm and inventor, and the residual is allocated to the poaching firm. We do not specify how the surplus is allocated between the poached firm and inventor. In addition, we assume that each firm posts the privately efficient number of vacancies, which maximizes the sum of the values of the firm and its inventors.[14]

[13]We consider vacancy costs not just as the cost of recruiting but more broadly as labor adjustment costs. The management literature on labor turnover often emphasizes the significant transaction costs associated with labor mobility. Hiring and training new employees are costly, and employee turnover can disrupt relationships and informal communication structures within the firm.

[14]Hawkins (2015) discusses this result with wage-contract commitment; Bilal et al. (2022) provide a

Under these assumptions, firms' and inventors' decisions are privately efficient, as if the firm and incumbent inventors maximize their total value. This joint-value formulation avoids tracking the distribution of individual wage contracts within a multi-inventor firm, which would otherwise become an additional state variable in an on-the-job-search environment. Because inventors are ex ante identical, the total value function has only two state variables: inventor employment n and knowledge capital k.

Distribution of Firms and Product Lines

To formally characterize the equilibrium, we need to introduce the distribution of firms and product lines. Let $g(k, n, t)$ denote the measure of firms in the (k, n)-space. Then, the measure of products in the (k, n)-space is given by $kg(k, n, t)$. As we will confirm later, the state variable of a firm is reduced to the number of inventors per product line, $x \equiv n/k = z^{-1}$. Anticipating this, let $f(x, t)$ denote the induced measure of products in the x-space, which is defined as:

$$f(x,t) = \int_0^\infty \int_0^\infty kg(k,n,t)\delta\Big(\frac{n}{k} - x\Big)\, dk\, dn,$$

where $\delta(x)$ is the Dirac delta function. Since the total mass of products is one, we have:

$$1 = \int_0^\infty f(x,t)dx. \tag{4}$$

Also, since all inventors are employed, we have:

$$N = \int_0^\infty x f(x,t)dx. \tag{5}$$

Let v denote the amount of vacancy per product line, where $v = \hat{v}/k$. Since labor market tightness is $\theta(t) \equiv V(t)/N$, we have:

$$N\theta(t) = V(t) = \int_0^\infty v(x,t) f(x,t)dx. \tag{6}$$

To characterize the equilibrium under random search, it is useful to define the inventor-weighted and vacancy-weighted sampling measures. In the (k, n)-space, define

$$dG_n(k,n,t) = \frac{ng(k,n,t)}{N}\, dk\, dn,$$

communication-protocol foundation without such commitment.

$$dG_{\tilde{v}}(k,n,t) = \frac{\hat{v}(k,n,t)g(k,n,t)}{V(t)}\,dk\,dn.$$

The corresponding weighted product densities in the x-space are

$$f_x(x,t) = \frac{xf(x,t)}{N},$$
$$f_v(x,t) = \frac{v(x,t)f(x,t)}{V(t)}.$$

Let $F_x(x,t)$ and $F_v(x,t)$ denote the cumulative distribution functions associated with these weighted densities. Here, $g(k,n,t)$ is introduced only to define firm-level quantities and map them into product-line-level objects. Because homogeneity implies that firms' dynamics and the equilibrium allocation can be summarized by the one-dimensional state variable $x \equiv n/k = z^{-1}$, the central distributional objects in what follows are $f(x,t)$ and the inventor-weighted and vacancy-weighted distributions derived from it.

3.2 Competitive Equilibrium

HJB Equation

The firm's vacancy-posting decision and the realization of inventor mobility after meetings are characterized by the HJB equation. Let $\Omega(k,n,t)$ denote the joint value of a firm with product mass k and n inventors at time t. This value function satisfies

$$\begin{aligned}
&\rho\Omega(k,n,t) - \frac{\partial}{\partial t}\Omega(k,n,t) \\
&= \max_{\hat{v}\geq 0}\Bigg\{\pi k - c\left(\hat{v}/k\right)k + \left[\mu\left(n/k\right) - \delta(t)\right]k\Omega_k(k,n,t) + \frac{\sigma^2}{2}k^2\Omega_{kk}(k,n,t) \\
&\qquad + (1-\beta)q(\theta(t))\hat{v}\int \hat{M}\left(k/n,k'/n',t\right)dG_n(k',n',t) \\
&\qquad + \beta\theta(t)q(\theta(t))n\int \hat{M}\left(k'/n',k/n,t\right)dG_{\tilde{v}}(k',n',t)\Bigg\},
\end{aligned}$$

where the match surplus $\hat{M}\left(k/n,k'/n',t\right)$ is given by

$$\hat{M}\left(k/n,k'/n',t\right) = \max\Big[\Omega_n(k,n,t) + \hat{\alpha}\left(k/n,k'/n'\right)\Omega_k(k,n,t) - \Omega_n\left(k',n',t\right),0\Big].$$

The terms on the right-hand side have the following interpretation. The term πk is current profit, and $-c\left(\hat{v}/k\right)k$ is the vacancy-posting cost. The drift terms involving Ω_k and

Ω_{kk} capture changes in value from internal R&D, creative destruction, and idiosyncratic shocks. The two remaining terms account for inventor mobility. A vacancy meets an inventor sampled from the inventor-weighted distribution G_n at rate $q(\theta(t))$; if the match surplus is positive, the poaching firm obtains a $(1-\beta)$ share of the surplus. Conversely, each incumbent inventor meets a vacancy sampled from the vacancy-weighted distribution $G_{\tilde{v}}$ at rate $\theta(t)q(\theta(t))$; when the inventor is poached, the coalition of the poached firm and inventor receives a β share.

The match surplus equals the poaching firm's marginal gain from hiring an inventor, including the associated knowledge inflow, net of the poached firm's marginal loss from losing that inventor.

The HJB equation implies that the value function is homogeneous of degree one in (k, n). It can therefore be written in product-line units, with the number of inventors per product line, $x = z^{-1}$, as the single state variable. Let $S(x,t)$ denote the value per product line, where $x \equiv n/k$ denotes inventors per product line and is the reciprocal of firm-level inventor productivity $z = k/n$. By homogeneity,

$$\Omega(k, n, t) = kS(x, t).$$

The firm-level marginal values of inventors and knowledge capital can be written as

$$\begin{aligned}\Omega_n(k, n, t) &= S_x(x, t),\\ \Omega_k(k, n, t) &= S(x, t) - xS_x(x, t),\end{aligned}$$

and the second derivative of firm-level joint value with respect to knowledge capital is expressed as

$$\Omega_{kk}(k, n, t) = \frac{1}{k}x^2 S_{xx}(x, t).$$

These expressions follow from differentiating $\Omega(k, n, t) = kS(n/k, t)$.

Using these expressions and defining the vacancy rate per product line as $v(x,t) \equiv \hat{v}(k, n, t)/k$, the firm-level HJB can be written in product-line units as

$$\begin{aligned}\rho S(x,t) - \frac{\partial}{\partial t}S(x,t) = \max_{v \geq 0}\Big\{&\pi - c\,(v) + \left[\mu(x) - \delta(t)\right]\left[S(x,t) - xS_x(x,t)\right]\\ &+ \frac{\sigma^2}{2}x^2 S_{xx}(x,t) + (1-\beta)q(\theta(t))v\int M\,(x, x', t)\,dF_x(x', t)\\ &+ \beta\theta(t)q(\theta(t))x\int M\,(x', x, t)\,dF_v(x', t)\Big\},\end{aligned} \tag{7}$$

where the match surplus $M(x, x', t)$ is given by

$$M(x, x', t) = \max\Big[S_x(x, t) + \alpha(x, x')\{S(x, t) - xS_x(x, t)\} - S_x(x', t), 0\Big]. \tag{8}$$

Here $\alpha(x, x') \equiv \hat{\alpha}(1/x, 1/x')$ maps the knowledge-diffusion function from inventor-productivity notation into the inventor-per-product state variable.

The match-surplus expression also yields a simple benchmark that clarifies the role of knowledge diffusion.

Proposition 1 (No-spillover benchmark and the job ladder)**.** *Fix any date t and suppose that $\alpha(x, x') = 0$ for all (x, x'). If $S(\cdot, t)$ is concave, then*

$$M(x, x', t) = [S_x(x, t) - S_x(x', t)]_+ .$$

Hence a successful poaching move from a source firm with state x' to a destination firm with state x can occur only if $x < x'$. Since $x = n/k = z^{-1}$, this means that the destination firm has fewer inventors per product line and, equivalently, higher inventor productivity $z = k/n$. In this no-spillover benchmark, inventors therefore move only toward firms with higher inventor productivity.

Proof. If $\alpha(x, x') = 0$, (8) implies

$$M(x, x', t) = [S_x(x, t) - S_x(x', t)]_+ .$$

Concavity of $S(\cdot, t)$ implies that $S_x(\cdot, t)$ is weakly decreasing. Therefore, $M(x, x', t) > 0$ requires $S_x(x, t) > S_x(x', t)$, which in turn requires $x < x'$. Since $x = z^{-1}$, this is equivalent to the destination firm having higher inventor productivity $z = k/n$. □

This benchmark clarifies the role of knowledge diffusion in the model. In the absence of spillovers, the search margin collapses to a standard job ladder: inventors move only toward firms with higher inventor productivity. Therefore, moves toward firms with lower inventor productivity require the knowledge term in (8) to be sufficiently large. In particular, for a move from a source firm with state x' to a destination firm with $x > x'$ to occur, it must be that

$$\alpha(x, x')\{S(x, t) - xS_x(x, t)\} > S_x(x', t) - S_x(x, t).$$

This condition provides the model counterpart to the empirical finding in Sections 2.1 and 2.2 that many inventors move to establishments with lower citation-productivity or wage ranks.

From (7), the first-order condition for the firm's vacancy decision gives

$$c'(v(x,t)) = (1-\beta)q(\theta(t)) \int M(x,x',t)\, dF_x(x',t). \tag{9}$$

Therefore, the vacancy policy is increasing in the share of the match surplus accruing to the poaching firm $(1-\beta)$, the matching probability $q(\theta(t))$, and the expected match surplus $\int M(x,x',t)\, dF_x(x',t)$.

Kolmogorov Forward Equation

To complete the description of the environment, we specify the evolution of the product-line distribution $f(x,t)$ over the state $x = n/k$. First, define the poaching indicator $\mathbb{1}_P$, which equals one when the match surplus is positive and zero otherwise:

$$\mathbb{1}_P(x,x',t) = \begin{cases} 1 & \text{if} \quad M(x,x',t) > 0 \\ 0 & \text{otherwise} \end{cases}. \tag{10}$$

The drift component of x induced by mobility, internal R&D, and creative destruction is

$$\begin{aligned} b_x(x,t) = q(\theta(t))v(x,t) \int_0^\infty \mathbb{1}_P(x,x',t) \left\{x^{-1} - \alpha(x,x')\right\} f_x(x',t)dx' \\ - \theta(t)q(\theta(t)) \int_0^\infty \mathbb{1}_P(x',x,t) f_v(x',t)dx' - \mu(x) + \delta(t). \end{aligned}$$

Applying Ito's lemma to $x = n/k$, we obtain the following Kolmogorov forward (KF) equation governing the evolution of $f(x,t)$:

$$\begin{aligned} \frac{\partial}{\partial t} f(x,t) = \left[q(\theta(t))v(x,t) \int_0^\infty \mathbb{1}_P(x,x',t)\alpha(x,x') f_x(x',t)dx' + \mu(x) - \delta(t) \right] f(x,t) \\ - \frac{\partial}{\partial x}\left[b_x(x,t)x f(x,t)\right] + \frac{\partial^2}{\partial x^2}\left[\frac{\sigma^2}{2}x^2 f(x,t)\right]. \end{aligned} \tag{11}$$

The first term captures net changes in the mass of product lines at state x. The second term is the transport term for the state x, reflecting changes in inventors per product line due to mobility, internal R&D, and creative destruction. The third term captures diffusion in x induced by idiosyncratic shocks.

Creative Destruction Rate

The common creative-destruction rate faced by each product is

$$\delta(t) = \underbrace{\int \mu(x) dF(x,t)}_{\text{Internal R\&D}} + \underbrace{N\theta(t)q(\theta(t)) \int \int \mathbb{1}_P(x,x',t)\alpha(x,x')\, dF_x(x',t) dF_v(x,t)}_{\text{Knowledge Diffusion}}. \quad (12)$$

The first term is the aggregate internal-R&D intensity across product lines. The second term is the aggregate contribution of knowledge diffusion generated by successful inventor flows. For future reference, define these two components by

$$\begin{aligned}
\delta_{RD}(t) &\equiv \int_0^\infty \mu(x)\, dF(x,t), \\
\delta_{KD}(t) &\equiv N\theta(t)q(\theta(t)) \int_0^\infty \int_0^\infty \mathbb{1}_P(x,x',t)\alpha(x,x')\, dF_x(x',t)\, dF_v(x,t),
\end{aligned}$$

so that $\delta(t) = \delta_{RD}(t) + \delta_{KD}(t)$. Because the economic growth rate is proportional to the creative-destruction rate, growth can be decomposed into contributions from internal R&D and knowledge diffusion.

To characterize the internal-R&D component, note that since $F(\cdot,t)$ is the product-line-weighted distribution, (4) implies $\int_0^\infty dF(x,t) = 1$, and (5) implies $\int_0^\infty x\, dF(x,t) = N$. Jensen's inequality therefore yields the following result.

Proposition 2 (Internal R&D growth and misallocation)**.** *Fix any date t. If $\mu(\cdot)$ is concave, then*

$$\delta_{RD}(t) \leq \mu(N).$$

If $\mu(\cdot)$ is strictly concave, equality holds if and only if $x = N$ $F(\cdot,t)$-almost surely.

Proof. Let X_t be distributed according to $F(\cdot,t)$. Then $\delta_{RD}(t) = \mathbb{E}[\mu(X_t)]$ and $\mathbb{E}[X_t] = N$. Jensen's inequality implies

$$\mathbb{E}[\mu(X_t)] \leq \mu(\mathbb{E}[X_t]) = \mu(N).$$

If $\mu(\cdot)$ is strictly concave, equality holds if and only if $X_t = N$ almost surely. □

This proposition shows that dispersion in inventor allocation reduces internal-R&D growth. Holding fixed the total mass of inventors N, the internal-R&D component is maximized when the number of inventors per product line, $x = n/k$, is equalized across product lines. Equivalently, internal-R&D growth is maximized when inventor productivity

$z = x^{-1} = k/n$ is equalized across firms. Thus, if lower inventor mobility widens the dispersion of x, it depresses growth not only by weakening knowledge diffusion but also by worsening the allocation of inventors for internal R&D.

Resource Constraint

Final output is either consumed by the representative household or used for vacancy posting. Vacancy-posting expenditures, measured in units of the numeraire, are $\int_0^\infty c(v(x,t))dF(x,t)$. Because nominal output is normalized to $P(t)Y(t) = 1$, the quantity of final output absorbed by vacancy posting is $Y(t)\int_0^\infty c(v(x,t))dF(x,t)$. The resource constraint is therefore

$$C(t) + Y(t)\int_0^\infty c(v(x,t))dF(x,t) = Y(t). \tag{13}$$

Definition of Equilibrium

We now define the equilibrium.

Definition 1. *A competitive equilibrium* is a collection of paths

$$\left[Y(t), C(t), \left\{S(x,t), v(x,t), \{\mathbb{1}_P(x,x',t)\}_{x'\in(0,\infty)}, f(x,t)\right\}_{x\in(0,\infty)}, \delta(t), \theta(t)\right]_{t\in[0,\infty)},$$

such that, given the initial output $Y(0)$ and the initial product-line state distribution $f(x,0)$ satisfying the $t = 0$ versions of (4) and (5),

1. the joint value $S(x,t)$ satisfies the HJB equation (7),
2. the vacancy policy is optimal, so that $v(x,t)$ satisfies (9),
3. given a match, the inventor moves if the match surplus is positive, so that the poaching indicator function $\mathbb{1}_P(x,x',t)$ is given by (10),
4. the product-line state distribution $f(x,t)$ follows the KF equation (11),
5. the creative destruction rate $\delta(t)$ satisfies (12),
6. aggregate output $Y(t)$ evolves according to (3),
7. the labor market tightness $\theta(t)$ satisfies (6), and
8. the resource constraint (13) holds.

The steady-state equilibrium is one in which all variables are constant over time.

3.3 Firm Growth and Model Properties

Firm Growth and Inventor Productivity

We next derive the firm-level growth equations that link the model to the inventor-productivity growth regressions in Section 4. These equations decompose growth in knowledge capital, inventor employment, and inventor productivity. The growth rate of knowledge capital in a firm with x inventors per product line is

$$d\log k(x,t) = \underbrace{\left(\mu(x) - \delta(t) - \frac{\sigma^2}{2}\right)dt + \sigma dW(t)}_{\text{Internal R\&D−Creative Destruction}} + \underbrace{q(\theta(t))v(x,t)\int_0^\infty \mathbb{1}_P(x,x',t)\alpha(x,x')\, f_x(x',t)dx'dt}_{\text{Knowledge Gain from Poaching}}. \tag{14}$$

The first line captures internal R&D and creative destruction, while the second captures knowledge diffusion through inventor inflows.

The growth rate of inventor employment of a firm with x inventors per product line is given by

$$d\log n(x,t) = \underbrace{\frac{q(\theta(t))v(x,t)}{x}\int \mathbb{1}_P(x,x',t)dF_x(x',t)dt}_{\text{inventor inflow rate}} - \underbrace{\theta(t)q(\theta(t))\int_0^\infty \mathbb{1}_P(x',x,t)f_v(x',t)dx'dt}_{\text{inventor outflow rate}}. \tag{15}$$

This expression is the difference between the inventor inflow rate and the inventor outflow rate.

Since inventor productivity is $z = x^{-1}$, its growth rate is given by $d\log z = d\log x^{-1} = d\log k - d\log n$. We use this relationship in Section 4 to estimate the parameter values.

Advantages of Modeling Product Lines as a Continuum

Our model treats the mass of product lines as a continuous state variable, rather than as an integer. This formulation differs from Klette and Kortum (2004) and related models, in which firms own an integer number of product lines. This continuous-state formulation is useful for two reasons.

First, it preserves homogeneity in the firm's product-line mass and inventor employment. As a result, the firm's value can be written as $kS(n/k)$, and the ratio $x = n/k$ is sufficient to describe the firm's dynamics. This homogeneity also implies a strong version of Gibrat's law. Firm size, whether measured by production worker employment or sales, is proportional to the mass of product lines k. As shown in (14), the growth rate of k is independent of k conditional on inventor productivity. Therefore, not only is the mean growth rate of a firm independent of firm size (a weak version of Gibrat's law), but the entire distribution of growth rates is also independent of firm size (a strong version of Gibrat's law).

By contrast, integer-product-line models such as Klette and Kortum (2004) and Atkeson and Burstein (2019) do not generally deliver this strong version. In these models, a firm acquires and loses product lines at Poisson arrival rates proportional to its size. In the absence of other types of shocks, the mean growth rate of a firm remains independent of firm size, satisfying a weak version of Gibrat's law. However, averaging within the firm implies that the variance of firm growth is inversely proportional to firm size. Consequently, the entire distribution of growth rates depends on firm size, violating the strong version of Gibrat's law.[15]

This reduction to a one-dimensional firm state is useful because inventor employment is costly to adjust. In models where firms can freely adjust R&D resources, those resources need not be state variables. When adjustment is costly, however, firm size and R&D resources typically both become state variables. The continuous-state formulation and homogeneity let us avoid this additional state dimension: the ratio of R&D resources to firm size, $x = n/k$ in our model, is sufficient to describe the firm's dynamics.

Second, the continuous-state formulation allows us to focus on job-to-job mobility among incumbent firms without modeling firm entry, firm exit, or off-the-job search. With continuous product-line masses, one can construct an equilibrium in which firms do not enter or exit. By contrast, if the number of product lines is integer-valued and firms stochastically lose product lines, firms eventually lose all product lines and exit, making entry and exit unavoidable. We therefore abstract from firm entry and exit, inventor unemployment, and off-the-job search, and focus instead on job-to-job transitions among inventors and the resulting knowledge diffusion.

[15]Stanley et al. (1996) show that independent growth across equal-sized units implies a size-volatility relationship that is too steep relative to the data, pointing instead to strong positive correlations across units.

3.4 Planner Problem and Wedges

Planner Allocation

To characterize the efficient allocation, we first state the planner's problem. The social planner chooses the path

$$\left[Y(t), \delta(t), \theta(t), \left\{v(x,t), f(x,t), \{\mathbb{1}_P(x,x',t)\}_{x'\in(0,\infty)}\right\}_{x\in(0,\infty)}\right]_{t\in[0,\infty)}$$

to maximize

$$\int_0^\infty e^{-\rho t}\left[\log Y(t) + \log\left\{1 - \int_0^\infty c(v(x,t))f(x,t)dx\right\}\right]dt, \tag{16}$$

subject to (3), (12), (6), and (11), given the initial output $Y(0)$ and initial density $f(x,0)$.

We state the planner allocation in a form parallel to the competitive equilibrium (Definition 1). Appendix B.3 records the planner formulation and the value representation used for this characterization.

Definition 2. *A planner allocation* is a collection of paths

$$\left[\left\{S(x,t), v(x,t), \{\mathbb{1}_P(x,x',t)\}_{x'\in(0,\infty)}, f(x,t)\right\}_{x\in(0,\infty)}, \delta(t), \theta(t), \varpi(t)\right]_{t\in[0,\infty)}$$

such that, given the initial product-line state distribution $f(x,0)$ satisfying the $t=0$ versions of (4) and (5), the share of output available for consumption is

$$\Psi(t) \equiv 1 - \int_0^\infty c(v(x,t))f(x,t)dx,$$

and the following conditions hold:

1. the joint value $S(x,t)$ satisfies the HJB equation

$$\begin{aligned}
&\left(\rho + \frac{\dot{\Psi}(t)}{\Psi(t)}\right)S(x,t) - S_t(x,t) \\
&\quad = \max_{v\geq 0}\Big\{\varpi(t) - c(v) + \left[\mu(x) - \delta(t)\right]\left[S(x,t) - xS_x(x,t)\right] + \frac{\sigma^2}{2}x^2 S_{xx}(x,t) \\
&\qquad + q(\theta(t))v\int M(x,x',t)\,dF_x(x',t) + \theta(t)q(\theta(t))x\int M(x',x,t)\,dF_v(x',t) \\
&\qquad + v\theta(t)q'(\theta(t))\int\int M(x'',x',t)\,dF_x(x',t)\,dF_v(x'',t)\Big\}.
\end{aligned} \tag{17}$$

2. the vacancy posting policy $v(x,t)$ satisfies

$$c'(v(x,t)) = q(\theta(t)) \int M(x,x',t)\, dF_x(x',t) + \theta(t) q'(\theta(t)) \int \int M(x'',x',t) dF_x(x',t) dF_v(x'',t). \tag{18}$$

3. the poaching indicator function $\mathbb{1}_P(x,x',t)$ satisfies (10),

4. the product-line state distribution $f(x,t)$ follows the KF equation (11),

5. the creative destruction rate $\delta(t)$ satisfies (12),

6. the labor market tightness $\theta(t)$ satisfies (6), and

7. $\varpi(t)$ is the normalized instantaneous social gain from a product line entering the planner HJB, and the associated social valuation of a one-step innovation satisfies

$$\underbrace{\frac{\log \lambda}{\rho}\left[1 - \int_0^\infty c(v(x,t)) f(x,t) dx\right]}_{\text{consumer-surplus gain from innovation}} = \underbrace{\int_0^\infty \left[S(x,t) - S_x(x,t)x\right] f(x,t) dx}_{\text{expected marginal social value of innovation}}. \tag{19}$$

The stationary version of the planner allocation is one in which all variables are constant over time. In that case, $\dot{\Psi}(t) = 0$ and the left-hand side of (17) reduces to the standard stationary expression.

When comparing the competitive equilibrium and the planner allocation, we use the same notation S and M for readability. These objects should be understood as being evaluated under the corresponding allocation. The comparison reveals two wedges: an innovation wedge and a search/matching wedge.

The innovation wedge arises because the competitive equilibrium and the planner allocation assign different values to additional knowledge capital. In the competitive equilibrium, the value of a new product line is disciplined by the private profit flow π. In contrast, the planner values innovation through aggregate consumption and the reallocation of knowledge capital across the economy. The planner's marginal social value of innovation is characterized by (19). The left-hand side of that equation is the consumption-normalized continuation-value gain from a one-step quality improvement, while the right-hand side is the expected marginal social value of the additional knowledge capital. The condition therefore provides the social valuation benchmark for innovation.

Comparing this social valuation with the private valuation based on π reveals two standard externalities.[16] First, innovators do not appropriate the full gain in consumer surplus generated by a quality improvement. This appropriability effect tends to make the private value of innovation too low relative to its social value. Second, innovation replaces incumbent producers, and the innovator does not internalize the value lost by the displaced incumbent. This business stealing effect tends to make the private value of innovation too high relative to its social value. Hence the gap between the private valuation based on π and the social valuation characterized above is the net innovation wedge, whose sign is not pinned down a priori.

The vacancy margin makes the search/matching wedge transparent. The competitive vacancy first-order condition is given by (9), whereas the planner vacancy first-order condition is given by (18). Holding fixed the value functions and distributions, the difference between the planner's and the poaching firm's marginal benefit from a vacancy is

$$\underbrace{\beta q(\theta(t)) \int M(x, x', t) dF_x(x', t)}_{\text{surplus-sharing / hold-up wedge}} + \underbrace{\theta(t) q'(\theta(t)) \int \int M(x'', x', t) dF_x(x', t) dF_v(x'', t)}_{\text{congestion externality}}.$$

The first term is positive and reflects that the poaching firm captures only a $1-\beta$ share of the surplus generated by its vacancy. Holding the congestion term fixed, this surplus-sharing wedge pushes private vacancy posting below the planner's level. The second term is negative because $q'(\theta(t)) < 0$: an additional vacancy lowers the meeting rate of other vacancies. Holding the surplus-sharing term fixed, this congestion externality pushes private vacancy posting above the planner's level. Therefore, the net search/matching wedge is generally ambiguous and depends on the state x, the distribution of potential source firms, and the vacancy-weighted distribution of destination firms.

In our model, these wedges operate through the value of knowledge capital, the match surplus $M(x, x', t)$, vacancy posting, and inventor mobility. Internal R&D intensity is not a separate control: given the allocation of inventors across product lines, the internal-R&D component is $\int \mu(x) dF(x, t)$. Thus the innovation wedge does not appear as an independent R&D effort first-order condition. Instead, it changes the value functions, and therefore the match surplus that governs vacancy posting and the mobility network. From (8), this surplus captures both the reallocation value of moving an inventor across firms and the knowledge diffusion generated by the move. Table 4 summarizes the components of the

[16] See Chapter 12 of Acemoglu (2008) for a more detailed discussion of the typical externalities that arise in Schumpeterian growth models.

Table 4: Wedges and Partial Effects

Broad wedge	Component	Partial effect on private incentives
Innovation	Appropriability	Private value of innovation is too low.
Innovation	Business stealing	Private value of innovation is too high.
Search/matching	Surplus sharing / hold-up	Vacancy posting is too low.
Search/matching	Congestion	Vacancy posting is too high.

Notes: The directions in the last column are partial effects. The net innovation wedge and the net search/matching wedge are not signed a priori. Because internal R&D intensity is not chosen separately, the innovation wedge affects the allocation through the value functions and match surplus rather than through an independent R&D effort margin.

wedges and their associated partial effects.

Bargaining Power as a Policy Instrument

We consider the bargaining power β as a policy instrument. Implementing the planner allocation would generally require policies that vary with the firm state, including inventor productivity k/n, but such policies are rare in practice. As a second-best approach, we focus on the best fixed bargaining power within this class of policy instruments. Bargaining power can be influenced by the regulation of non-compete clauses. For example, Shi (2023) models a contractual environment involving non-compete clauses. Her model, like ours, maximizes the joint value of the firm and the incumbent worker. In that environment, stronger enforceability of non-compete clauses increases the share of surplus captured by the firm and its incumbent worker. Furthermore, in the standard DMP model, optimal bargaining power can fully eliminate inefficiencies caused by congestion externalities and hold-up problems. In our heterogeneous environment, by contrast, the analogous Hosios-like benchmark is state-dependent, so a constant bargaining power is necessarily a second-best instrument. Even so, it can still mitigate a substantial portion of the search inefficiency in the quantitative model. For these reasons, we treat bargaining power as a natural policy instrument within our framework.

In the model, stronger non-compete enforceability is represented by a higher β, which assigns a larger share of match surplus to the poached firm and the inventor and a smaller share to the poaching firm. The direct effect is to reduce the poaching firm's return to vacancy posting, as shown by (9). A higher β also changes continuation values: incumbent firms receive a larger surplus share when their inventors are poached, and the resulting decline in aggregate vacancy posting lowers future poaching risk in equilibrium. These continuation-value effects enter through S and M, but the direct vacancy-posting incentive operates through the poacher share $1 - \beta$.

3.5 Relationship to Existing Endogenous Growth Models with Knowledge Diffusion

We relate our framework to endogenous-growth models in which random meetings generate knowledge diffusion. A first distinction concerns the allocation of match surplus. Early models, such as Lucas and Moll (2014) and Perla and Tonetti (2014), as well as inventor-based models including Akcigit et al. (2018) and Prato (2025), share with our model the feature that random matches transmit knowledge across agents. In these models, however, the knowledge recipient captures the match surplus. As a result, they do not internalize the effects of other agents learning from their own knowledge or technology.

In Hopenhayn et al. (2026) and Benhabib et al. (2021), as in our model, the surplus generated from a match is allocated between the two agents in fixed proportions based on bargaining power. Relatedly, Jovanovic and Wang (2025) study innovation, diffusion, and intellectual property protection in an industry-evolution model and characterize the socially optimal licensing rate in the presence of matching externalities. However, the relevant sharing rule and its relationship to policy differ between these frameworks and ours. While these papers connect the sharing rule to patent enforcement or intellectual property protection, our model—where knowledge is transmitted through inventor mobility—links bargaining power to labor-market institutions that govern inventor mobility, such as the enforceability of non-compete clauses.

A second distinction is that our model makes inter-firm inventor mobility the mechanism of knowledge diffusion. Models with micro-foundations for inventors, such as Akcigit et al. (2018) and Prato (2025), do not account for inventor job-to-job flows across firms. By introducing this mechanism, our framework can be disciplined by microdata that jointly observe inventor job flows, patents, and employer-employee links.

4 Estimation

In this section, we bring the model to the German matched employer-employee-patent data. The quantitative exercise combines external parameter choices and normalizations with parameters chosen by matching model-implied moments to their empirical counterparts. The calibration proceeds in three blocks: moments of inventor-productivity growth discipline the diffusion, internal-R&D, and shock parameters; baseline inventor mobility determines matching efficiency; and aggregate growth pins down the innovation step size. The exposition first specifies the functional forms and summarizes the resulting parameterization, then discusses each parameter group in turn.

Table 5: Functional Forms

Function	Interpretation
$c(v) = \frac{\overline{c}}{\phi+1} v^{\phi+1}$	Vacancy cost function
$q(\theta) = \overline{q}\theta^{\eta-1}$	Rate at which a vacancy meets an inventor
$\mu(x) = \overline{\mu} x^{\epsilon}$	Internal R&D function
$\alpha(x, x') = \overline{\alpha}/x'$	Knowledge-diffusion function

4.1 Functional Forms

We begin by specifying the functional forms summarized in Table 5. The vacancy cost function is $c(v) = \frac{\overline{c}}{\phi+1} v^{\phi+1}$, as in Kaas and Kircher (2015). The matching function is Cobb–Douglas, so a vacancy meets an inventor at rate $q(\theta) = \overline{q}\theta^{\eta-1}$. The internal R&D function is $\mu(x) = \overline{\mu} x^{\epsilon}$, as in Lentz and Mortensen (2008). Finally, we assume that the knowledge-diffusion function is given by $\alpha(x, x') = \overline{\alpha}/x'$. This specification makes the knowledge transferred per inventor inflow proportional to source-firm inventor productivity, $z' = x'^{-1}$. The resulting parameterization contains 13 parameters. We set or estimate them in three groups: eight are set externally or normalized, three are disciplined by inventor-productivity growth regressions, and two are calibrated to balanced-growth-path targets.

4.2 Calibration

Externally Set or Normalized Parameters

As summarized in Table 6, eight parameters are either set externally or normalized. The discount rate, $\rho = 0.004$, reflects a 5% annual real interest rate. Since we cannot identify the scale parameter of the vacancy cost function $\overline{c}$ and the matching efficiency $\overline{q}$ separately, we normalize $\overline{c}$. The vacancy cost elasticity, $\phi = 3.45$, is based on estimates from Bilal et al. (2022). The elasticity of matches with respect to vacancies, $\eta = 0.50$, follows standard values in the matching literature.[17] We set the internal-R&D elasticity to $\epsilon = 0.50$, following related quantitative work (e.g., Acemoglu et al., 2018). The instantaneous profit per product line, $\pi = 0.23$, is set so that the markup equals 1.3, following recent markup estimates (see Cavenaile et al., 2023). The total mass of inventors, $N = 1.95$, is chosen to match average German inventor productivity, measured by the count of three-year forward citations.[18]

We set the baseline bargaining weight to $\beta = 0.5$. This value is not intended to estimate

[17] See Petrongolo and Pissarides (2001) for a review.

[18] In our model, the total knowledge capital is 1. Therefore, the average inventor productivity in the model is $1/N$. Since the average number of three-year forward citations for German inventors is 0.512, we set the total mass of inventors such that $1/N = 0.512$.

Table 6: Parameter Values Used in Quantitative Exercises

Parameter	Value	Description
A. Externally Set or Normalized		
ρ	0.004	Monthly discount rate
$\overline{c}$	1	Scale parameter of the vacancy cost function, normalized
ϕ	3.45	Curvature of the vacancy cost function
η	0.50	Elasticity of matches with respect to vacancies
ϵ	0.50	Internal-R&D elasticity
π	0.23	Instantaneous profit per product line
N	1.95	Total mass of inventors
β	0.50	Externally set baseline bargaining power
B. Inventor-Productivity Growth Estimates		
$\overline{\alpha}$	0.526	Scale parameter of the knowledge-diffusion function
$\overline{\mu}$	0.0212	Scale parameter of the internal-R&D function
σ	0.203	Standard deviation of internal-R&D shocks
C. Baseline BGP Targets		
$\overline{q}$	0.0198	Matching efficiency
λ	1.053	Step size on the quality ladder

Notes: Panel B reports the parameter values disciplined by the inventor-productivity growth estimates in column (1) of Table A.5. The annual internal-R&D coefficient and root mean squared error are converted to monthly model units by dividing by 12 and $\sqrt{12}$, respectively. Panel C reports parameters chosen to match the baseline BGP targets: the monthly log equivalents of the 2001 annual inventor job-to-job transition rate and a 2% annual growth rate.

the enforceability of non-compete clauses; instead, it provides a symmetric Nash-bargaining benchmark for splitting match surplus. Symmetric bargaining weights are standard in search-and-matching models and, with Cobb–Douglas matching and $\eta = 0.5$, coincide with the Hosios condition in the canonical environment (Hosios, 1990; Petrongolo and Pissarides, 2001). In our economy, knowledge diffusion and state-dependent match surplus make the corresponding Hosios-like benchmark state-dependent, so no single value of β implements efficiency. We therefore use $\beta = 0.5$ as a neutral baseline and vary β in the counterfactual analysis to study its allocation and welfare effects.

Inventor-Productivity Growth Estimates

The parameters governing knowledge diffusion, internal R&D, and internal-R&D shocks, $\overline{\alpha}$, $\overline{\mu}$, and σ, are disciplined by establishment-level regressions of inventor-productivity growth. The construction proceeds in three steps. We first derive the law of motion for inventor productivity, $z = x^{-1} = k/n$, then map the diffusion and internal-R&D terms into establishment-year data, and finally estimate the implied regression.

Since $z = x^{-1}$, its growth rate is $d \log z = d \log k - d \log n$. By subtracting (15) from (14),

we obtain

$$d\log z(x,t) = q(\theta(t))v(x,t)\int_0^\infty \mathbb{1}_P(x,x',t)\alpha(x,x')\, f_x(x',t)dx'dt + \left(\mu(x) - \delta(t) - \frac{\sigma^2}{2}\right)dt - d\log n(x,t) + \sigma dW(t).$$

Under $\alpha(x,x') = \overline{\alpha}/x'$ and $z(x') = x'^{-1}$, this becomes

$$d\log z(x,t) = \overline{\alpha}q(\theta(t))v(x,t)\int_0^\infty \mathbb{1}_P(x,x',t)z(x')f_x(x',t)dx'dt + \left(\mu(x) - \delta(t) - \frac{\sigma^2}{2}\right)dt - d\log n(x,t) + \sigma dW(t).$$

To map the diffusion term into establishment-year data, consider inventor flows between $t-1$ and t. An inflow of $m_{e'\to e,t}$ inventors from establishment e' to establishment e contributes to log knowledge capital by

$$\begin{aligned} d\log k_{e,t}\big|_{e'\to e} &= \alpha(x_{e,t-1}, x_{e',t-1})\frac{m_{e'\to e,t}}{k_{e,t-1}} \\ &= \overline{\alpha}\frac{m_{e'\to e,t}}{n_{e,t-1}}\frac{z_{e',t-1}}{z_{e,t-1}}. \end{aligned}$$

Thus, when inflows are measured per incumbent inventor at the destination, the model-implied diffusion exposure is the inflow rate weighted by source inventor productivity relative to destination inventor productivity. The internal-R&D term maps into $\overline{\mu}\Delta z_{e,t-1}^{-0.5}$, because $\mu(x) = \overline{\mu}x^{0.5} = \overline{\mu}z^{-0.5}$ and Δ is the length of the observation interval in monthly model units. The employment-growth term, $-d\log n(x,t)$, is handled by including the net inventor inflow rate in the controls.

In the data, inventor productivity is measured using three-year forward citations. Let $z_{e,t}^{\text{inc}}$ denote the average inventor productivity of inventors who remain at establishment e over the interval from $t-1$ to t, and define $\Delta\log z_{e,t}^{\text{inc}} \equiv \log z_{e,t}^{\text{inc}} - \log z_{e,t-1}^{\text{inc}}$. The model then implies the following regression equation:

$$\Delta\log z_{e,t}^{\text{inc}} = \overline{\alpha}\sum_{e'}\frac{m_{e'\to e,t}}{n_{e,t-1}}\frac{z_{e',t-1}}{z_{e,t-1}} + \overline{\mu}\,\Delta\, z_{e,t-1}^{-0.5} + \text{Controls}_{e,t} + \varepsilon_{e,t},$$

The first regressor is the diffusion exposure used to discipline $\overline{\alpha}$. The inverse-root lagged destination-inventor-productivity term is the internal-R&D exposure used to discipline $\overline{\mu}$. The baseline specification includes establishment and year fixed effects; the more saturated specifications replace the year fixed effects with industry-by-region-by-year fixed effects

and, in the richest specification, add establishment-specific linear trends. The error term $\varepsilon_{e,t}$ has variance $\sigma^2\Delta$.[19] The inflow term uses $m_{e'\to e,t}$, the number of inventors newly hired by establishment e from establishment e' between $t-1$ and t, divided by the lagged number of inventors at the destination, $n_{e,t-1}$. Source and destination inventor productivity in the inflow term are measured at $t-1$, and the internal-R&D term uses lagged destination inventor productivity, $z_{e,t-1}$. We infer the standard deviation of internal R&D shocks, σ, from the regression root mean squared error. Standard errors are clustered by establishment. We set $\Delta = 12$ because the regressions use annual panel data, whereas one model period corresponds to one month. The baseline calibration uses column (1) of Table A.5, yielding $\overline{\alpha} = 0.526$, $\overline{\mu} = 0.0212$, and $\sigma = 0.203$; the latter two values convert the annual internal-R&D coefficient and root mean squared error to monthly model units. A detailed description of the regression analysis, its results, robustness checks, and auxiliary exercises is provided in Appendix C.1. That appendix also reports a source-shock exercise that allocates severe source-employment contractions through lagged inventor hiring networks. Because the exercise does not use contemporaneous realized destination choices, it mitigates a central selection concern and supports a causal interpretation of the knowledge-diffusion moment, although it is not used as the calibrated estimate of $\overline{\alpha}$.

Baseline BGP Targets

Conditional on the externally set parameters and the inventor-productivity-growth estimates, the baseline BGP leaves two parameters to be chosen internally: matching efficiency $\overline{q}$ and the innovation step size λ. We choose these parameters sequentially. First, $\overline{q}$ is chosen so that the baseline BGP matches the 12.6% annual job-to-job transition rate of German inventors in 2001. Because the model period is one month, the target is the monthly log equivalent,

$$\iota = \frac{\log(1 + 0.126)}{12}.$$

In the stationary equilibrium, the model-implied inventor job-to-job transition rate is

$$\iota = \overline{q}\theta^{\eta} \int\int \mathbb{1}_P(x, x')\, dF_x(x')\, dF_v(x),$$

so that

$$\overline{q} = \frac{\iota}{\theta^{\eta} \iint \mathbb{1}_P(x, x')\, dF_x(x')\, dF_v(x)}.$$

[19]The dependent variable is constructed using only inventors who remain at establishment e over the interval. This restriction keeps inventors who arrive between $t-1$ and t out of the dependent-variable measurement, so the inflow term does not mechanically add the entrants' own patent output to the outcome.

Table 7: Untargeted Productivity-Dispersion Moments

Moment	Data	Model
SD of log positive inventor productivity	0.98	0.91
P90/P50 of positive inventor productivity	2.94	3.33
P50/P10 of positive inventor productivity	4.81	3.24

Notes: Data moments use establishment-year observations with positive three-year backward average citation productivity per inventor, pooled over rank years 2001–2011 and weighted by inventor counts. Model moments use inventor productivity k/n in the baseline steady state and the inventor-weighted stationary distribution. None of the moments in this table is targeted in the calibration.

The innovation step size λ is then chosen to match a 2% annual aggregate growth rate. Since λ scales aggregate growth but does not affect the stationary allocation used to recover $\overline{q}$, we use the BGP creative-destruction rate δ from the calibrated stationary equilibrium and impose

$$g = \delta \log \lambda, \qquad \lambda = \exp\left(\frac{g}{\delta}\right),$$

where $g = \log(1.02)/12$ is the monthly log growth target. This procedure gives $\overline{q} = 0.0198$ and $\lambda = 1.053$.

4.3 Model Fit

Table 7 compares cross-sectional inventor-productivity dispersion in the data with the baseline steady state. These moments are not used as calibration targets. The model is close to the data in the standard deviation of log positive inventor productivity and in the upper-tail ratio, while it somewhat understates dispersion in the lower tail. Overall, the untargeted moments indicate that the baseline steady state generates a plausible amount of heterogeneity in inventor productivity.

5 Quantitative Exercises

This section quantifies how inventor mobility affects growth and welfare through the two margins emphasized by the model: the allocation of inventors across firms and knowledge diffusion carried by job-to-job moves. We first examine steady-state allocations under different values of bargaining power. We then analyze the transition after a non-compete ban. Next, we ask how close a fixed bargaining split can come to the planner allocation. Finally, we assess the impact of the observed decline in inventor mobility in Germany on knowledge diffusion, internal R&D activities, and economic growth. The numerical implementation of the HJB and Kolmogorov forward equations and the

consumption-equivalent welfare calculation are described in Appendix C.3, Appendix C.4, and Appendix C.5.

5.1 Bargaining Power, Mobility, and Growth in Steady State

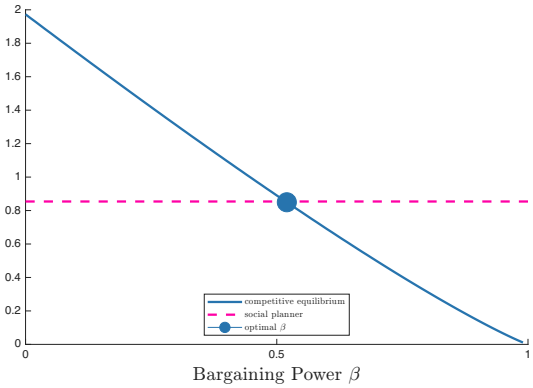
This subsection studies steady-state comparative statics with respect to bargaining power β, which summarizes the strength of non-compete enforcement in the model. Lower β raises the poacher's share of match surplus, increasing vacancy posting, labor market tightness, and inventor mobility. The same force strengthens knowledge diffusion but also raises recruiting resource costs. The exercise therefore clarifies the resource-cost, mobility, and growth channels before turning to transition welfare.

Figure 1 summarizes the aggregate steady-state implications. Moving from higher to lower β raises vacancy posting, which increases the vacancy-cost-to-output ratio (Figure 1a), labor market tightness (Figure 1b), and the job-to-job transition rate (Figure 1c). Across bargaining powers, aggregate growth responds mainly through the knowledge-diffusion margin. At the baseline calibration with $\beta = 0.5$, knowledge diffusion accounts for 16.2% of total growth, while the remaining 83.8% is driven by internal R&D. However, knowledge-diffusion growth (Figure 1e) is much more sensitive to inventor mobility than internal-R&D growth. Consequently, the decline in aggregate growth caused by a higher β (Figure 1f) is driven primarily by weaker knowledge diffusion.

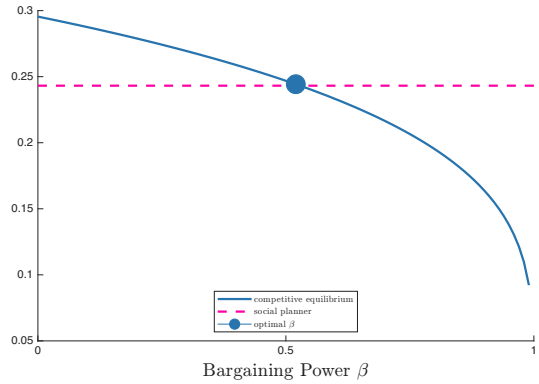
The internal-R&D contribution to growth, shown in Figure 1d, corresponds to $\delta_{RD}(t)$ in Section 3. With an inelastic supply of inventors, internal-R&D growth is determined by the allocative efficiency of inventors. It is maximized when inventor productivity $z = x^{-1} = k/n$ is equalized across firms. In equilibrium, however, inventor misallocation arises from search frictions, idiosyncratic shocks, and the concavity of the internal-R&D technology. By reducing inventor mobility, a higher β worsens this allocation and lowers internal-R&D growth.

Figure 2 explains the firm-level mechanism behind these aggregate effects. The steady-state vacancy policy is U-shaped in inventor productivity. Posting more vacancies allows firms to hire more inventors and, at the same time, to acquire more knowledge capital through knowledge diffusion. The vacancy policy therefore depends on both the marginal value of inventors and the marginal value of knowledge capital. When the number of inventors n is small relative to knowledge capital k, inventor productivity $z = x^{-1} = k/n$ is high and the marginal value of inventors is large (Figure 2c). Conversely, when knowledge capital k is small relative to the number of inventors n, inventor productivity is low and the marginal value of knowledge capital is large (Figure 2d). These two forces generate the

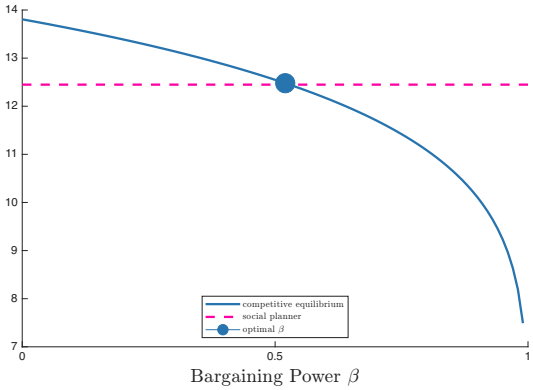
Figure 1: Steady-State Aggregates by Bargaining Power



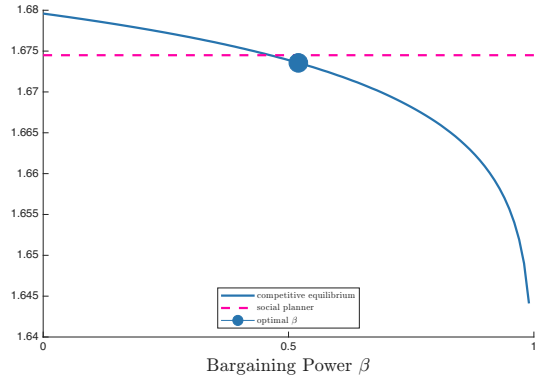
(a) Vacancy-Cost-to-Output Ratio



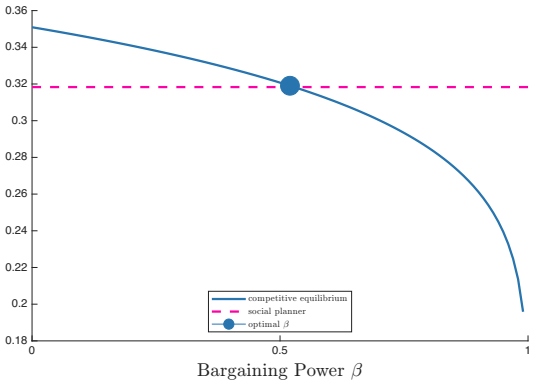
(b) Labor Market Tightness



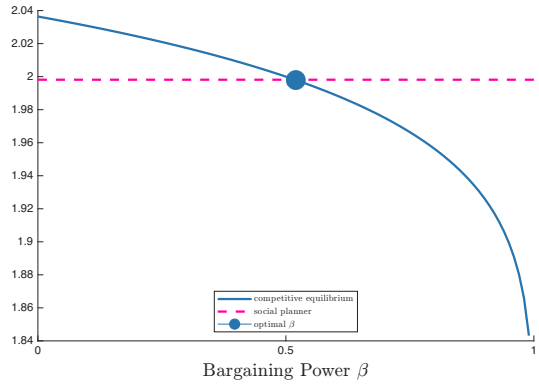
(c) Job-to-Job Transition Rate



(d) Internal-R&D Growth



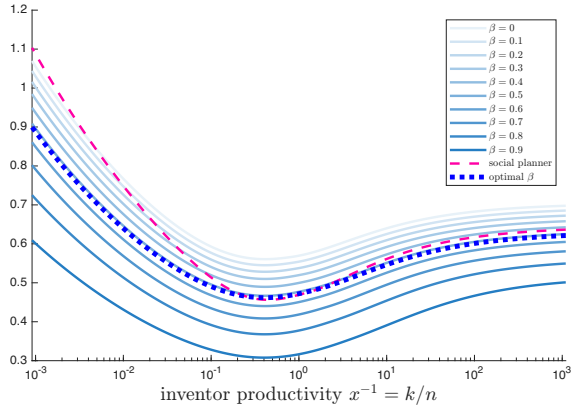
(e) Knowledge-Diffusion Growth



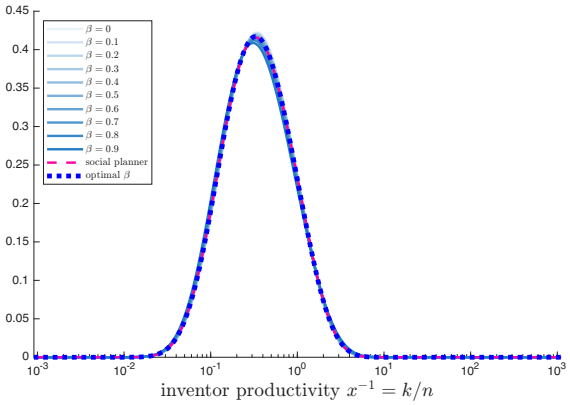
(f) Economic Growth Rate

Notes: Panels plot aggregate steady-state outcomes against the fixed bargaining power β. The solid blue line reports decentralized equilibria, the dashed magenta line reports the planner allocation, and the filled blue circle marks the computed second-best fixed bargaining power, $\beta = 0.52$, discussed in Section 5.3. A larger β represents stronger non-compete enforcement. Growth rates and the job-to-job transition rate are annualized percentages; labor market tightness is unitless.
Alt text: Six line charts compare steady-state outcomes as bargaining power changes. Stronger non-compete enforcement is associated with lower job-to-job mobility and weaker growth, mainly through lower knowledge-diffusion growth. Planner and second-best fixed-split benchmarks are shown for comparison.

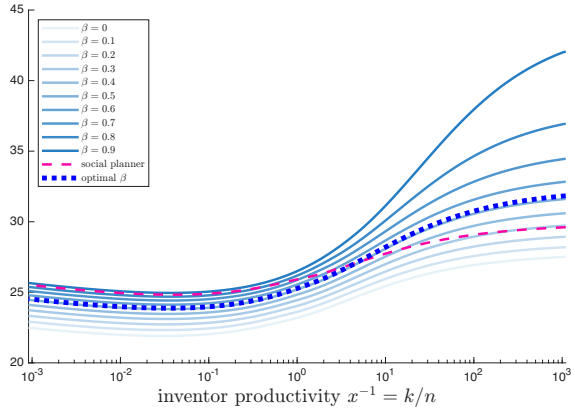
Figure 2: Steady-State Policy Functions by Bargaining Power



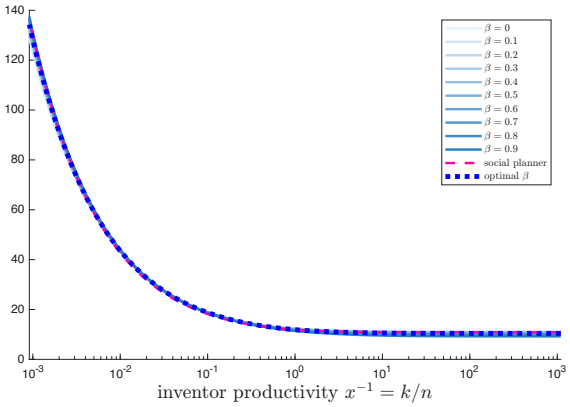
(a) Vacancy Posting Policy



(b) Inventor-Weighted Distribution



(c) Marginal Inventor Value



(d) Marginal Knowledge Value

Notes: Panels plot steady-state objects against inventor productivity $z = x^{-1} = k/n$. Solid blue curves report decentralized equilibria for fixed bargaining powers $\beta = 0, 0.1, \ldots, 0.9$, where a larger β represents stronger non-compete enforcement. The dashed magenta curve reports the planner allocation, and the dotted blue curve reports the computed second-best fixed bargaining power, $\beta = 0.52$.
Alt text: Four line charts show steady-state policy functions and distributions by inventor productivity. Vacancy posting is U-shaped in inventor productivity, and the marginal values of inventors and knowledge capital are highest in opposite parts of the state space.

U-shaped vacancy policy.

Figure 2a also shows that vacancy posting decreases as β increases. A higher β reduces the share of match surplus captured by the poaching firm, directly lowering the return to vacancy posting. In the calibrated model, the direct incentive effect dominates the equilibrium feedbacks from lower vacancy posting by other firms and from the higher surplus received by firms whose inventors are poached.

The same figures also report the planner allocation and the computed second-best fixed β as benchmarks. We return to the interpretation of this fixed split in Section 5.3.

5.2 A Non-Compete Ban: Growth and Welfare

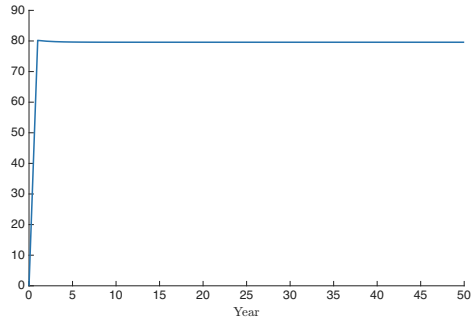
This subsection evaluates a ban on non-compete clauses. The ban raises long-run annual growth, mainly by strengthening knowledge diffusion, but it lowers welfare because the induced increase in vacancy posting raises recruiting resource costs on impact.

The exercise is motivated by recent policy debates over restrictions on non-compete clauses. We model the policy as a transition in which the economy starts on the BGP under the baseline bargaining power $\beta = 0.5$. At $t = 0$, bargaining power unexpectedly and permanently falls to $\beta = 0$, and agents have perfect foresight thereafter. When $\beta = 0$, a firm that matches with an inventor can successfully poach by compensating the inventor for their marginal value at the current employer. Thus, the experiment isolates the effect of increasing the poacher's share of match surplus.

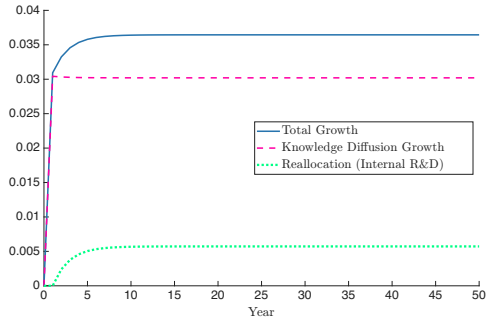
Figure 3 illustrates the transition dynamics relative to the initial BGP. First, the ban raises growth. When β falls to zero, the poacher's share of match surplus rises, so firms post more vacancies and inventor mobility increases. Long-run annual growth rises by 0.036 percentage points: 0.030 points come from stronger knowledge diffusion and 0.006 points from higher internal R&D growth. Thus, most of the growth effect comes from stronger knowledge diffusion (Figure 3b). The smaller reallocation component reflects that diffusion-driven poaching need not move inventors toward firms with the highest marginal inventor value.

Second, the welfare implication differs from the growth implication. The same increase in vacancy posting that raises mobility also raises recruiting resource costs: the vacancy-cost-to-output ratio rises by about 80% relative to the initial BGP (Figure 3a). Output rises over time (Figure 3c), but more final goods are initially devoted to vacancy posting, so consumption falls on impact (Figure 3d). The initial consumption loss outweighs the later gains from faster growth. As a result, consumption-equivalent welfare, defined in Appendix C.5, falls by 0.38% relative to the no-policy-change path.

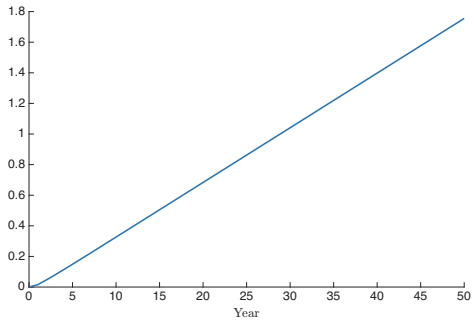
Figure 3: Transition Dynamics after a Non-Compete Ban



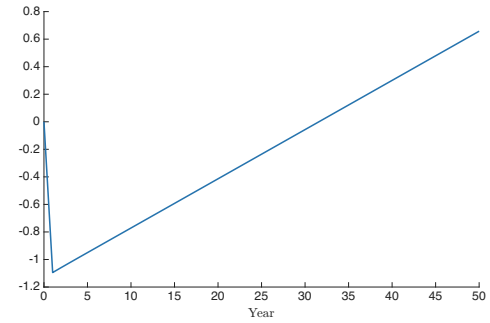
(a) Vacancy-Cost-to-Output Ratio



(b) Growth-Rate Decomposition



(c) Log Output



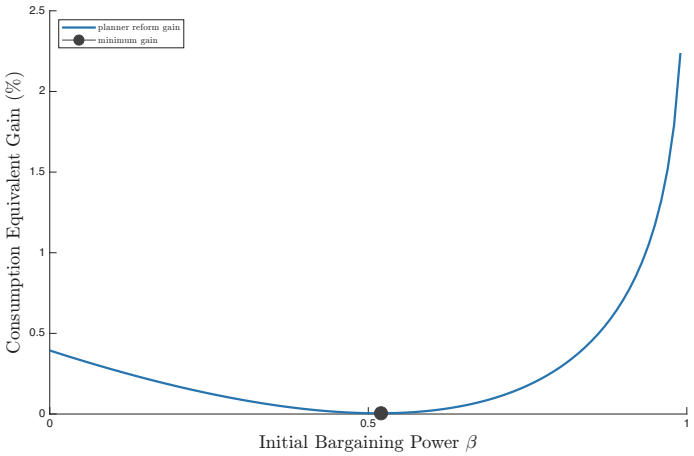
(d) Log Consumption

Notes: The economy is initially on the balanced growth path (BGP) under the baseline bargaining power $\beta = 0.5$. At year 0, β permanently falls to 0, which is interpreted as a ban on non-compete clauses; agents have perfect foresight thereafter. All panels report transition paths relative to the initial BGP. The growth panel decomposes total growth into internal-R&D and knowledge-diffusion components. Vacancy costs, log output, and log consumption are reported in percent; growth-rate deviations are percentage points.
Alt text: Four transition charts show the effects of lowering bargaining power from 0.5 to 0. Vacancy costs jump on impact, growth rises mainly through knowledge diffusion, output gradually increases, and consumption initially falls before recovering.

5.3 Fixed Bargaining Splits and the Planner Benchmark

The steady-state figures above include the planner allocation as a benchmark. This subsection quantifies how much welfare remains to be gained from replacing each fixed-β competitive equilibrium with the planner allocation. We measure this gap by the consumption-equivalent gain from a planner-reform transition. For a fixed value of β, the economy starts from the corresponding competitive steady state. At $t = 0$, it unexpectedly transitions to the planner allocation, and agents have perfect foresight thereafter. We then compare welfare along this transition with welfare from remaining forever in the initial competitive steady state.

Figure 4: Planner-Reform Welfare Gain by Bargaining Power



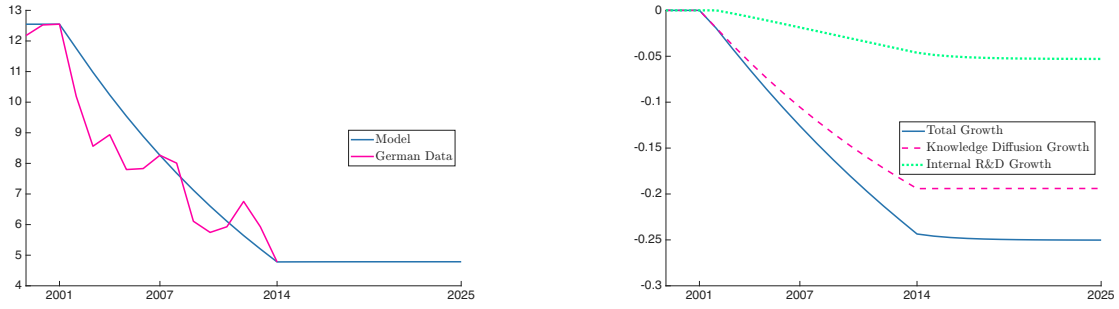
Notes: The figure reports the consumption-equivalent welfare gain, in percent, from the planner-reform transition for each fixed bargaining power $\beta = 0, 0.01, \ldots, 0.99$. The filled circle marks the minimum-gain point, $\beta = 0.52$.
Alt text: A line chart plots the consumption-equivalent welfare gain from moving from each fixed bargaining-power equilibrium to the planner allocation. The gain is positive throughout and reaches its minimum near bargaining power 0.52.

The exercise is diagnostic rather than an implementation result. In this heterogeneous economy, the planner-implied vacancy-posting wedge is state-dependent, so no single constant split exactly decentralizes the planner allocation. We therefore interpret the fixed β that minimizes the planner-reform gain as the second-best fixed bargaining power.

Figure 4 shows that the planner-reform gain is positive for all fixed bargaining powers, but it is minimized near the calibrated baseline $\beta = 0.5$. At $\beta = 0.52$, the gain from transitioning to the planner is only 0.0046% of consumption. Thus, within the class of constant bargaining powers, a near-symmetric split comes very close to the planner benchmark in welfare terms.

This welfare proximity should not be interpreted as exact decentralization. The Hosios-like benchmark for the vacancy-posting margin is state-dependent in this heterogeneous economy, and the innovation wedge remains a separate margin. Hence no single constant bargaining rule exactly implements the planner allocation. The value $\beta = 0.52$ is best interpreted as a second-best fixed split that comes quantitatively close to the planner allocation. The derivation of the state-dependent Hosios-like benchmark and related numerical diagnostics are reported in Appendix B.4.

Figure 5: Inventor Mobility Decline and Growth



(a) Job-to-Job Transition Rate

(b) Growth-Loss Decomposition

Notes: The left panel plots the German inventor job-to-job transition rate and the calibrated model path. The matching-efficiency path is calibrated so that the model matches the German inventor transition rates in 2001, 2007, and 2014; after 2014, matching efficiency remains fixed while the economy continues to converge to the new BGP. The right panel reports growth-rate changes relative to the initial BGP and decomposes total growth into internal-R&D growth, δ_{RD}, and knowledge-diffusion growth, δ_{KD}. Job-to-job transition rates are in percent, and growth-rate changes are percentage points.
Alt text: Two charts compare the observed and model-implied decline in German inventor job-to-job mobility and the associated growth effects. The mobility decline reduces growth mainly by weakening knowledge diffusion rather than internal R&D.

5.4 Declining Inventor Mobility, Growth, and Welfare

This subsection analyzes how the observed decline in inventor mobility affects growth and welfare. The central result is that lower mobility reduces annual growth by about 0.25 percentage points, and most of this loss comes from weaker knowledge diffusion rather than from the reallocation margin for internal R&D. As shown in Figure 5a, German inventor mobility fell sharply over the sample period: the job-to-job transition rate decreased from 12.6% in 2001 to 4.8% in 2014. The underlying annual rates are reported in Appendix A.3.

The same broad decline is also visible in U.S. inventor data. Using the Inventor Employment History, Akcigit and Goldschlag (2025) show that the hiring rate for inventors fell from 7% in 2000, the first observable year, to 3.5% in 2016, the last observable year.[20] Because our German data are at the establishment level whereas their U.S. data are at the firm level, the levels are not directly comparable. We therefore treat the U.S. evidence as corroborating evidence of a broader decline in inventor mobility, while calibrating the counterfactual exercise to the German inventor transition rates.

To quantify the implications of this decline, we feed the model a smooth deterministic path of matching efficiency that matches the German inventor job-to-job transition rates in 2001, 2007, and 2014. The exercise maps $t = 0$ to 2001 and $t = T$ to 2014, uses the baseline steady-state value for $\overline{q}_0$, and fixes matching efficiency after 2014 while the economy

[20] See Figure 11 in their paper.

continues to converge to the new BGP. Agents do not anticipate the change before $t = 0$, and have perfect foresight thereafter. The functional form and calibration details are reported in Appendix C.2. Along with the targeted German inventor job-to-job transition rates, Figure 5a plots the corresponding model path.

Figure 5b decomposes the resulting growth loss. Lower mobility reduces growth through both channels in the model. It worsens the allocation of inventors across firms, which lowers internal-R&D growth, and it weakens knowledge diffusion between firms. In the long run, growth from internal R&D declines by 0.05 percentage points, while growth from knowledge diffusion drops by 0.19 percentage points. The total decline in annual growth is therefore about 0.25 percentage points, or roughly 13% of the calibrated initial annual growth rate of 2%. Consumption-equivalent welfare decreases by 3.49% relative to welfare along the initial BGP, where matching efficiency remains constant.

6 Conclusion

In this paper, we examine how inventor mobility between firms shapes knowledge diffusion, long-run growth, and welfare. Using matched German patent and employer-employee records, we document three patterns consistent with knowledge diffusion through inventor mobility. Patent-listed inventors often move to establishments with lower citation productivity and lower mean wages; inventor-proxy workers are more likely than comparable workers to move to establishments with lower mean wages while receiving larger wage gains; and inventor inflows from higher-inventor-productivity sources predict faster incumbent inventor-productivity growth at destination establishments. A source-shock exercise based on severe source-employment contractions allocated through lagged inventor hiring networks supports interpreting this relationship as reflecting knowledge diffusion rather than only destination-side selection. Taken together, these patterns point beyond a simple wage-ladder account of job-to-job mobility: inventors may move toward establishments that rank lower on conventional establishment-level measures while carrying knowledge that is valuable at the destination.

Motivated by these patterns, we build an endogenous economic growth model in which inventors advance technology both through in-house R&D and by transporting knowledge when they switch employers. The model formalizes two channels through which mobility affects growth. It reallocates inventive labor across firms, changing internal-R&D growth, and it transfers knowledge from source to destination firms, contributing directly to creative destruction. Because firms choose vacancies and inventors move in response to poaching incentives, inventor job flows are equilibrium objects rather than exogenous

shocks. Calibrated to German data, the framework replicates the mobility dynamics in the data and allows us to assess labor-market policies that shape these two margins.

The counterfactual exercises deliver three key insights. First, policies that relax mobility restrictions can raise growth but need not raise welfare. Banning non-compete clauses raises annual growth by 0.036 percentage points, with 0.030 points from stronger knowledge diffusion and 0.006 points from a modest increase in internal R&D growth. Thus, most of the growth effect operates through knowledge diffusion. However, the ban also raises recruiting resource costs through increased vacancy posting, so consumption falls on impact and consumption-equivalent welfare falls by 0.38%. Second, no constant bargaining rule exactly decentralizes the planner allocation, because the Hosios-like vacancy-posting benchmark is state-dependent and the innovation wedge remains. Nonetheless, the planner-reform gain is minimized near $\beta = 0.52$, where the gain from transitioning to the planner is only 0.0046% of consumption. Third, the observed decline in inventor mobility since the early 2000s lowers growth from internal R&D by 0.05 percentage points and growth from knowledge diffusion by 0.19 percentage points, cutting overall growth by about 0.25 percentage points and reducing welfare by 3.49% relative to the initial balanced growth path.

Taken together, the results suggest that inventor mobility is not only a labor-market outcome but also a channel through which knowledge is reallocated across firms. Policies that restrict or encourage such mobility therefore affect growth through two margins: the allocation of inventive labor and the diffusion of knowledge embodied in mobile inventors. The policy implication is therefore not that more mobility is always desirable, but that policies governing high-skilled mobility should be evaluated jointly through their effects on knowledge diffusion, inventor allocation, and recruiting resource costs. Although our analysis focuses on patent-listed inventors in Germany and abstracts from firm entry, off-the-job search, and non-patented knowledge, the results suggest that similar mechanisms may matter for other high-skilled workers whose movement carries firm-specific knowledge across organizations.

References

Acemoglu, Daron. 2008. "Introduction to Modern Economic Growth."

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. 2018. "Innovation, Reallocation, and Growth." *The American Economic Review* 108 (11): 3450–3491.

Akcigit, Ufuk, and Sina T Ates. 2023. "What Happened to US Business Dynamism?" *The journal of political economy* 131 (8): 2059–2124.

Akcigit, Ufuk, Santiago Caicedo, Ernest Miguelez, Stefanie Stantcheva, and Valerio Sterzi. 2018. "Dancing with the Stars: Innovation Through Interactions."

Akcigit, Ufuk, and Nathan Goldschlag. 2025. "Measuring the Characteristics and Employment Dynamics of U.S. Inventors." *Journal of Economic Growth* 30 (2): 237–269. 10.1007/s10887-025-09251-9.

Albagli, Elias, Mario Canales, Chad Syverson, Matias Tapia, and Juan Wlasiuk. 2025. "Productivity Growth and Workers' Job Transitions: Evidence from Census Microdata." *The Economic Journal* 135 (666): 405–429. 10.1093/ej/ueae081.

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

Arrow, Kenneth J. 1962. *Economic Welfare and the Allocation of Resources for Invention*. 609–626, Princeton: Princeton University Press.

Atkeson, Andrew, and Ariel Burstein. 2019. "Aggregate Implications of Innovation Policy." *The Journal of Political Economy* 127 (6): 2625–2683.

Babalievsky, Fil. 2023. "Misallocation in the Market for Inventors." Manuscript.

Benhabib, Jess, Jesse Perla, and Christopher Tonetti. 2021. "Reconciling Models of Diffusion and Innovation: A Theory of the Productivity Distribution and Technology Frontier." *Econometrica* 89 (5): 2261–2301.

Bilal, Adrien, Niklas Engbom, Simon Mongey, and Giovanni L Violante. 2022. "Firm and Worker Dynamics in a Frictional Labor Market." *Econometrica* 90 (4): 1425–1462.

Bilal, Adrien, Niklas Engbom, Simon Mongey, and Giovanni L Violante. 2023. "Labor Market Dynamics When Ideas Are Harder to Find." In *The Economics of Creative Destruction: New Research on Themes from Aghion and Howitt*, edited by Ufuk Akcigit, and John Van Reenen, 221–252, Harvard University Press. 10.4159/9780674293052-010.

Braunerhjelm, Pontus, Ding Ding, and Per Thulin. 2020. "Labour Market Mobility, Knowledge Diffusion and Innovation." *European Economic Review* 123 103386.

Breschi, Stefano, and Francesco Lissoni. 2009. "Mobility of Skilled Workers and Co-Invention Networks: An Anatomy of Localized Knowledge Flows." *Journal of Economic Geography* 9 (4): 439–468.

Buera, Francisco J, and Robert E Lucas. 2018. "Idea Flows and Economic Growth." *Annual Review of Economics* 10 (1): 315–345.

Buera, Francisco J, and Ezra Oberfield. 2020. "The Global Diffusion of Ideas." *Econometrica* 88 (1): 83–114.

Cavenaile, Laurent, Murat Alp Celik, and Xu Tian. 2023. "Are markups too high?

Competition, strategic innovation, and industry dynamics." *SSRN Electronic Journal*.

Crews, Levi. 2023. "A Dynamic Spatial Knowledge Economy."

Dauth, Wolfgang, and Johann Eppelsheimer. 2020. "Preparing the sample of integrated labour market biographies (SIAB) for scientific analysis: a guide." *Journal for Labour Market Research* 54 (10).

Dorner, Matthias, Dietmar Harhoff, Fabian Gaessler, Karin Hoisl, and Felix Poege. 2018. "Linked Inventor Biography Data 1980-2014." FDZ-Datenreport 03/2018 en.

Elsby, Michael W L, and Axel Gottfries. 2021. "Firm Dynamics, On-the-Job Search, and Labor Market Fluctuations." *The Review of Economic Studies* 89 (3): 1–50.

Federal Trade Commission. 2024. "Noncompete Rule." `https://www.ftc.gov/legal-library/browse/rules/noncompete-rule`, September, Accessed: 2026-06-12.

Frodermann, Corinna, Alexandra Schmucker, Stefan Seth, and Philipp vom Berge. 2021. "Sample of Integrated Labour Market Biographies (SIAB) 1975–2019." FDZ-Datenreport01/2021 (en), Institute for Employment Research (IAB), Nuremberg. `10.5164/IAB.FDZD.2101.en.v1`.

Gilson, Ronald J. 1999. "The Legal Infrastructure of High Technology Industrial Districts: Silicon Valley, Route 128, and Covenants Not to Compete." *New York University Law Review* 74 (3): 575–629.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg. 2001. "The NBER patent citation data file: Lessons, insights and methodological tools.." NBER Working Paper No. 8498.

Haltiwanger, John C., Henry R. Hyatt, Lisa B. Kahn, and Erika McEntarfer. 2018. "Cyclical Job Ladders by Firm Size and Firm Wage." *American Economic Journal: Macroeconomics* 10 (2): 52–85.

Hawkins, William B. 2015. "Bargaining with Commitment Between Workers and Large Firms." *Review of Economic Dynamics* 18 (2): 350–364.

Herkenhoff, Kyle, Jeremy Lise, Guido Menzio, and Gordon M Phillips. 2024. "Production and Learning in Teams." *Econometrica: Journal of the Econometric Society* 92 (2): 467–504.

Hoisl, Karin. 2007. "Tracing Mobile Inventors: The Causality between Inventor Mobility and Inventor Productivity." *Research Policy* 36 (5): 619–636.

Hopenhayn, Hugo, Liyan Shi, and Koki Okumura. 2026. "Knowledge Creation and Diffusion with Limited Appropriation."

Hosios, Arthur J. 1990. "On the Efficiency of Matching and Related Models of Search and Unemployment." *The Review of Economic Studies* 57 (2): 279–298.

Jaffe, Adam B, Manuel Trajtenberg, and Rebecca Henderson. 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations." *The Quarterly Journal of Economics* 108 (3): 577–598.

Jarosch, Gregor, Ezra Oberfield, and Esteban Rossi-Hansberg. 2021. "Learning from Coworkers." *Econometrica* 89 (2): 647–676.

Jovanovic, Boyan, and Zhu Wang. 2025. "Idea Diffusion and Property Rights." 10.21144/wp20-11, Working paper, revised November 2025.

Kaas, Leo, and Philipp Kircher. 2015. "Efficient Firm Dynamics in a Frictional Labor Market." *The American Economic Review* 105 (10): 3030–3060.

Kaiser, Ulrich, Hans Christian Kongsted, and Thomas Rønde. 2015. "Does the Mobility of R&D Labor Increase Innovation?" *Journal of Economic Behavior & Organization* 110 91–105.

Klette, Tor Jakob, and Samuel Kortum. 2004. "Innovating Firms and Aggregate Innovation." *The Journal of Political Economy* 112 (5): 986–1018.

Kortum, Samuel S. 1997. "Research, Patenting, and Technological Change." *Econometrica: Journal of the Econometric Society* 65 (6): 1389.

Lentz, Rasmus, and Dale T Mortensen. 2008. "An empirical model of growth through product innovation." *Econometrica: journal of the Econometric Society* 76 (6): 1317–1373.

Lentz, Rasmus, and Dale T Mortensen. 2022. "Labor Market Friction, Firm Heterogeneity, and Aggregate Employment and Productivity."

Lucas, Robert E, and Benjamin Moll. 2014. "Knowledge Growth and the Allocation of Time." *The Journal of Political Economy* 122 (1): 1–51.

Mawdsley, John K, and Deepak Somaya. 2016. "Employee Mobility and Organizational Outcomes: An Integrative Conceptual Framework and Research Agenda." *Journal of Management* 42 (1): 85–113.

McCrary, Sean. 2022. "A Job Ladder Model of Firm, Worker, and Earnings Dynamics."

Pakes, A. S. 1986. "Patents as options: Some estimates of the value of holding European patent stocks." *Econometrica* 54 (4): 755–784.

Perla, Jesse, and Christopher Tonetti. 2014. "Equilibrium Imitation and Growth." *The Journal of Political Economy* 122 (1): 52–76.

Petrongolo, Barbara, and Christopher A Pissarides. 2001. "Looking into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature* 39 (2): 390–431.

Postel-Vinay, Fabien, and Jean-Marc Robin. 2002. "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity." *Econometrica* 70 (6): 2295–2350.

Prato, Marta. 2025. "The Global Race for Talent: Brain Drain, Knowledge Transfer, and Growth." *The Quarterly Journal of Economics* 140 (1): 165–238. 10.1093/qje/qjae040.

Rahko, Jaana. 2017. "Knowledge Spillovers through Inventor Mobility: The Effect on Firm-Level Patenting." *The Journal of Technology Transfer* 42 (3): 585–614.

Rosenkopf, Lori, and Paul Almeida. 2003. "Overcoming Local Search through Alliances

and Mobility." *Management Science* 49 (6): 751–766.

Schaal, Edouard. 2017. "Uncertainty and Unemployment." *Econometrica* 85 (6): 1675–1721.

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

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

Song, Jaeyong, Paul Almeida, and Geraldine Wu. 2003. "Learning-by-Hiring: When Is Mobility More Likely to Facilitate Interfirm Knowledge Transfer?" *Management Science* 49 (4): 351–365.

Stanley, Michael H R, Luís A N Amaral, Sergey V Buldyrev, Shlomo Havlin, Heiko Leschhorn, Philipp Maass, Michael A Salinger, and H Eugene Stanley. 1996. "Scaling Behaviour in the Growth of Companies." *Nature* 379 (6568): 804–806.

Appendix

Appendix A Empirical Appendix

A.1 Data and Sample Construction

Our analyses use two administrative data sets, "Linked Inventor Biography Data 1980–2014" (INV-BIO) and "Sample of Integrated Labor Market Biographies" (Stichprobe der Integrierten Arbeitsmarktbiografien—SIAB). Both data sets are constructed by the Institute for Employment Research (IAB); see Dorner et al. (2018) for INV-BIO and Frodermann et al. (2021) for SIAB 7519.

The patent information contained in the INV-BIO data set is sourced from register data recorded in PATSTAT, which includes bibliographical, procedural, and legal status information on patent applications handled by the European Patent Office. Additionally, data from DPMAregister, the online patent register of the German Patent and Trademark Office, is incorporated to enhance the PATSTAT data extract. The DPMAregister provides exclusive records of national patent applications that are not transferred to the European Patent Office or filed under the PCT (Patent Cooperation Treaty) route. As a result, the INV-BIO data set comprises inventors who are listed on patent filings at the European Patent Office (EPO) between 1999 and 2011 and have been successfully linked with IEB (for more details, see Dorner et al., 2018).

Within INV-BIO, the inventor-flow analysis uses inventor job-to-job moves with observed origin and destination establishments, no intervening unemployment spell, and an establishment change. The inventor-productivity-growth analysis uses establishment-year observations with the diffusion exposure, lagged inventor productivity, net inventor inflows, and fixed-effect variables observed, after the trimming described in the notes to Table A.1 and Table A.5.

The SIAB data set is a 2% random sample from the Integrated Employment Biographies (IEB). The IEB combines data from five different sources, each of which may contain information from various administrative procedures. It comprises all individuals in Germany who hold at least one of the following employment statuses: employment subject to social security, marginal part-time employment, receipt of benefits according to the German Social Code III or II, official registration as a job seeker at the German Federal Employment Agency, and (planned) participation in programs of active labor market policies (for more details, see Frodermann et al., 2021; Dauth and Eppelsheimer, 2020). Within SIAB, the inventor-proxy worker regressions use job-to-job moves with no

intervening unemployment spell, observed origin and destination establishments, and the required lagged worker controls. Observations that are singletons under the absorbed fixed effects are dropped.

Table A.1 reports summary statistics for the empirical samples used in the main analysis. Panel A describes the INV-BIO establishment-year samples used for the inventor-flow and inventor-productivity-growth exercises. Panel B describes the SIAB job-to-job sample used in the inventor-proxy worker regressions.

A.2 Robustness Checks for Inventor Flow Patterns

Table A.2 repeats the inventor-proxy worker flow regressions in the subsample of job-to-job moves accompanied by an increase in the worker's own wage. This restriction checks whether moves to lower-mean-wage establishments simply reflect individual wage losses. Table A.3 applies the same restriction to patent-listed inventor flows, showing that many flows from more productive establishments to less productive ones remain even when the worker's own wage rises after the move.

Together, these wage-increase restrictions support the interpretation that lower-rank inventor moves are not merely wage-loss moves.

A.3 Inventor Mobility Rates

Table A.4 reports annual job-to-job transition rates for inventor person-year observations. The rate falls from 12.6% in 2001 to 4.8% in 2014. The 2001 rate provides the baseline mobility target for the balanced-growth-path calibration in Section 4. The 2001, 2007, and 2014 rates are also used to discipline the matching-efficiency path in the inventor-mobility-decline counterfactual in Section 5.

Appendix B Theoretical Appendix

This appendix collects the basic derivations used in the model, the planner value representation, and the state-dependent Hosios-like benchmark for the search margin.

Table A.1: Summary Statistics for Empirical Samples

Panel A. INV-BIO

Establishment-level variables	Mean	S.D.	N, thousand
Establishment-year sample			
N of inventors (n_{et})	5.18	29.53	678.9
N of employees	300.76	1088.36	678.9
Mean wage, Euro	100.83	54.63	670.3
Three-year backward average citation productivity ($\bar{z}_{et}^{(3)}$)	0.13	0.66	465.6
Total inventor inflows	0.40	2.99	678.9
Inventor-productivity-growth sample			
Inventor productivity growth	-0.02	1.09	13.9
Knowledge-diffusion inflow term	0.16	0.34	13.9
Net inventor inflow rate	0.06	0.14	13.9

Panel B. SIAB

	Full job-to-job		
Worker-level variables	Mean	S.D.	N, th.
D_{it}^{wage}	0.45	–	652
$\Delta \log w_{it}$	0.173	0.657	652
Wage increase indicator	0.74	–	652
$I_{i,t-1}$	0.10	–	652
Age	33.3	11.3	652
Female, %	42.7	–	652
University-entrance qualification, %	27.0	–	652
Log wage	3.877	0.871	652

Notes: Panel A reports INV-BIO establishment-year observations; its N column is measured in thousands of establishment-years. Rows under "Establishment-year sample" use available establishment-year observations. Three-year backward average citation productivity is the establishment's current and two lagged annual citation-productivity values averaged within establishment. Rows under "Inventor-productivity-growth sample" use the trimmed establishment-year sample used for the inventor-productivity growth regressions in Table A.5, conditional on a positive knowledge-diffusion inflow term. The knowledge-diffusion inflow term is total lagged source inventor productivity brought by inventor inflows, divided by lagged destination inventor productivity and the lagged number of destination inventors. Mean wage is the establishment-level mean wage of full-time employees. Panel B reports the SIAB full job-to-job estimation sample used in Table 3, after fixed-effect singleton observations are dropped; its N column is measured in thousands of person-year observations. D_{it}^{wage} equals one if worker i's destination establishment has lower previous-year mean wage than the origin establishment. $I_{i,t-1}$ and the controls are measured in the previous year. Dashes indicate standard deviations are not reported for indicator variables.

Table A.2: Estimation Results for Inventor-Proxy Worker Flows: Wage-Increase Moves

	Wage-increase job-to-job sample	
	(1)	(2)
Dependent variable	D_{it}^{wage}	$\Delta \log w_{it}$
$I_{i,t-1}$	0.0094***	0.0535***
	(0.0025)	(0.0027)
Controls	Yes	Yes
Origin estab. FE	Yes	Yes
Dest. estab. FE	Yes	Yes
Year FE	Yes	Yes
Observations	404,722	404,722
Adjusted R^2	0.5678	0.7002

Notes: Column (1) reports a fixed-effects linear probability model. Column (2) reports fixed-effects OLS estimates of wage changes. The sample is restricted to job-to-job moves with no intervening unemployment spell and an increase in the worker's own wage. D_{it}^{wage}, $I_{i,t-1}$, controls, fixed effects, singleton drops, clustering, and significance stars are defined as in Table 3.

Table A.3: Transition Matrices of Inventor Job-to-Job Flows with Wage Increases (Percent)

(A) Rank by Citation Productivity

	Destination rank bin									
Origin rank bin	1	2	3	4	5	6	7	8	9	10
1	12	11	10	10	8	9	9	10	10	10
2	11	11	11	9	9	9	9	9	10	10
3	9	11	11	11	9	11	9	10	10	10
4	11	10	11	11	10	10	11	9	9	8
5	8	8	9	8	9	11	16	10	9	10
6	7	8	10	7	11	10	10	12	11	14
7	5	7	6	14	12	9	12	11	11	12
8	5	6	7	6	8	13	12	15	15	14
9	4	5	11	6	10	8	9	17	12	18
10	4	7	5	6	5	10	10	11	16	26

(B) Rank by Mean Wage

	Destination rank bin									
Origin rank bin	1	2	3	4	5	6	7	8	9	10
1	17	16	13	12	10	9	7	7	5	4
2	10	16	13	12	10	9	8	8	7	6
3	7	12	13	12	10	9	9	9	9	9
4	5	9	11	11	11	11	11	11	11	9
5	5	8	9	11	10	11	10	12	12	13
6	3	7	8	9	10	13	13	12	13	12
7	3	6	7	8	9	11	13	13	15	16
8	2	5	5	7	10	11	14	14	15	18
9	2	4	5	5	9	9	14	14	16	22
10	2	3	4	4	5	7	11	12	18	34

Notes: Entries are row-normalized percentages and rows sum to 100 up to rounding. The sample is restricted to inventor job-to-job moves accompanied by an increase in the worker's own wage. Job-to-job moves require an establishment change with no intervening unemployment spell. Rank-bin definitions follow Table 1; origin and destination ranks use previous-year establishment measures. The sample covers 1980–2014.

Table A.4: German Inventor Job-to-Job Transition Rates by Year

Year	Inventor person-years	Job-to-job moves	Rate (%)
1999	117,904	14,358	12.2
2000	126,076	15,790	12.5
2001	129,050	16,194	12.6
2002	130,510	13,301	10.2
2003	130,597	11,180	8.6
2004	129,933	11,611	8.9
2005	129,425	10,093	7.8
2006	129,400	10,133	7.8
2007	129,446	10,699	8.3
2008	129,127	10,340	8.0
2009	127,244	7,773	6.1
2010	124,402	7,146	5.7
2011	122,456	7,260	5.9
2012	120,339	8,130	6.8
2013	118,021	6,998	5.9
2014	115,710	5,531	4.8
Total	2,009,640	166,537	8.3

Notes: The table reports annual inventor person-years, job-to-job moves, and the corresponding transition rate after restricting to years 1999 and later. Job-to-job moves require current and lagged establishments to be observed, no intervening unemployment spell, and a change in establishment. The rate equals 100 times job-to-job moves divided by inventor person-years and is rounded to one decimal place.

B.1 Household's Problem

The representative household has the preference

$$\int_0^\infty e^{-\rho t} \log C(t) dt,$$

and faces the budget constraint

$$\frac{d}{dt} A(t) = r(t)A(t) + w(t)L(t) + I(t) - P(t)C(t),$$

where $I(t)$ is the aggregate income of inventors. The current-value Hamiltonian is

$$\mathcal{H}^H\left(t, A(t), C(t), \lambda(t)\right) = \log C(t) + \lambda(t)\left[r(t)A(t) + w(t)L(t) + I(t) - P(t)C(t)\right],$$

where $\lambda(t)$ is the costate variable. We obtain the first-order conditions

$$\frac{1}{C(t)} = \lambda(t)P(t) \tag{20}$$

$$\rho\lambda(t) - \dot{\lambda}(t) = \lambda(t)r(t), \tag{21}$$

and the transversality condition

$$\lim_{t\to\infty}\left[\exp\left(-\rho t\right)\lambda(t)A(t)\right]=0.$$

Combining (20) and (21) gives the nominal Euler equation

$$r(t)=\rho+\frac{d}{dt}\log\left(P(t)C(t)\right).$$

Now focus on a balanced growth path (BGP). The normalization $P(t)Y(t)=1$ keeps nominal output constant, and on a BGP the consumption-output ratio is time-invariant, so $P(t)C(t)$ is constant as well. Hence the Euler equation reduces to $r=\rho$ on the BGP. Along a BGP all nominal quantities relevant here are stationary, so $P(t)C(t)$, $w(t)L(t)$, $I(t)$, and $r(t)$ are constant. The household budget constraint therefore implies $\dot{A}(t)=0$ and $A(t)=\bar{A}$ where $\bar{A}$ is a constant. The transversality condition reduces to $\lim_{t\to\infty}\left[\exp\left(-\rho t\right)\bar{A}\right]=0$, which holds whenever $\rho>0$.

B.2 Final Good Producer's Problem

Given prices, competitive final-good producers maximize profits

$$\Pi(t)=P(t)Y(t)-\int_0^1 p(\omega,t)l(\omega,t)d\omega,$$

subject to the technology

$$\begin{aligned}\log Y(t)&=\int_0^1\log\left(a(\omega,t)l(\omega,t)\right)d\omega\\&=\int_0^1\log a(\omega,t)d\omega+\int_0^1\log l(\omega,t)d\omega.\end{aligned}$$

The first-order condition with respect to $l(\omega,t)$ is $P(t)\partial Y(t)/\partial l(\omega,t)=p(\omega,t)$. Since $\partial\log Y(t)/\partial l(\omega,t)=1/l(\omega,t)$, this condition implies

$$p(\omega,t)l(\omega,t)=P(t)Y(t).$$

The choice of the numeraire $P(t)Y(t)=1$ therefore gives

$$p(\omega,t)l(\omega,t)=1,$$

for all $\omega \in [0,1]$ and $t > 0$.

B.3 Social Planner's Problem

This subsection records the formal characterization used in Definition 2; it is not an existence proof for the infinite-dimensional control problem. Rather than work with density costates directly, it is more transparent to characterize the planner's allocation through the marginal social value of an infinitesimal coalition of product lines and inventors.

Planner Value Representation

Let $G(k,n,t)$ denote the firm measure in the (k,n)-space; when this measure admits a density, that density corresponds to $g(k,n,t)$ in the main text. As in the main text, G induces the product density $f(x,t)$, the inventor-weighted firm measure G_n, the vacancy-weighted firm measure $G_{\tilde{v}}$, and the product-level weighted densities $f_x(x,t)$ and $f_v(x,t)$. Let

$$\Psi(t) \equiv 1 - \int_0^\infty c(v(x,t)) f(x,t) dx$$

denote the share of output available for consumption. The planner's continuation value can then be written as

$$\mathcal{V}\left(t, \log Y, G\right) \equiv \sup_{\{v, \mathbb{1}_P\}_{s \geq t}} \int_t^\infty e^{-\rho(s-t)} \left[\log Y(s) + \log \Psi(s)\right] ds,$$

subject to (3), (12), (6), (11), and the current firm distribution $G_t = G$. Because $\log Y$ enters additively,

$$\mathcal{V}\left(t, \log Y, G\right) = \frac{\log Y}{\rho} + \mathcal{U}(t, G),$$

where $\mathcal{U}(t,G)$ is the continuation value net of the current output level.

To characterize optimal reallocation, consider admissible signed perturbations $h(dk, dn)$ that preserve the aggregate mass of product lines and inventors to first order:

$$\int k \, h(dk, dn) = 0, \qquad \int n \, h(dk, dn) = 0.$$

If the directional derivative of $\mathcal{U}$ admits an integral representation on this tangent space, write

$$D_G \mathcal{U}(t, G)[h] = \frac{1}{\Psi(t)} \int \Omega(k, n, t) \, h(dk, dn).$$

The object $\Omega(k, n, t)$ is therefore the $\Psi(t)$-scaled marginal social value of an infinitesimal coalition organized as a firm with knowledge capital k and inventor employment n.

At the functional level, the planner's Bellman equation is

$$\rho\mathcal{V}\left(t, \log Y, G\right) = \max_{\{v, \mathbb{1}_P\}} \left\{\log Y + \log \Psi + \mathcal{V}_t + \mathcal{V}_{\log Y} \dot{\log} Y + D_G \mathcal{V}[\dot{G}_t]\right\},$$

where $\dot{G}_t$ is the law of motion for the firm distribution induced by the current controls. Taking a first variation of this Bellman equation in a feasible coalition direction yields the Bellman equation for Ω.

Firm-Level Planner HJB

Let $\varpi(t)$ denote the normalized instantaneous social gain from one product line, expressed in the same current-value normalization as $\Omega(k, n, t)$. Over a short interval dt, a coalition with state (k, n) contributes the current flow

$$\varpi(t)k - c\left(\hat{v}(k, n, t)/k\right) k.$$

Its knowledge capital evolves as

$$dk = k\left[\mu\left(n/k\right) - \delta(t)\right] dt + \sigma k \, dW_t,$$

so Ito's formula contributes

$$\left[\mu\left(n/k\right) - \delta(t)\right] k\Omega_k(k, n, t) + \frac{\sigma^2}{2} k^2 \Omega_{kk}(k, n, t)$$

to the drift of the marginal value.

When one of the firm's vacancies meets an inventor from a firm with state (k', n'), the planner's net gain is the match surplus

$$\hat{M}((k, n), (k', n'), t) = \max\left[\Omega_n(k, n, t) + \hat{\alpha}(k/n, k'/n')\Omega_k(k, n, t) - \Omega_n(k', n', t), 0\right].$$

Using the inventor-weighted firm distribution, successful poaching by the current firm contributes

$$q(\theta(t))\hat{v}(k, n, t) \int \hat{M}((k, n), (k', n'), t) \, dG_n(k', n', t).$$

Likewise, each inventor currently employed by the firm meets vacancies at rate $\theta(t)q(\theta(t))$. Since the destination firm is drawn from $G_{\tilde{v}}$, reallocations of the firm's own inventors

contribute

$$\theta(t)q(\theta(t))n \int \hat{M}((k',n'),(k,n),t)\, dG_{\tilde{v}}(k',n',t).$$

Finally, a marginal increase in $\hat{v}$ raises aggregate tightness by $1/N$, so the planner internalizes the congestion effect

$$\hat{v}(k,n,t)\theta(t)q'(\theta(t)) \int\int \hat{M}((k'',n''),(k',n'),t)\, dG_n(k',n',t)\, dG_{\tilde{v}}(k'',n'',t).$$

This term is absent from the decentralized firm HJB because an individual firm takes $\theta(t)$ as given.

Because $\Omega(k,n,t)$ is defined in the $\Psi(t)$-scaled current-value normalization, the Bellman equation for Ω inherits the extra term $\dot{\Psi}(t)/\Psi(t)$ on the left-hand side. Collecting the short-horizon terms yields the firm-level planner HJB

$$\begin{aligned}
&\left(\rho + \frac{\dot{\Psi}(t)}{\Psi(t)}\right)\Omega(k,n,t) - \Omega_t(k,n,t) \\
&= \max_{\hat{v}\geq 0}\Bigg\{\omega(t)k - c(\hat{v}/k)k + \left[\mu(n/k) - \delta(t)\right]k\Omega_k(k,n,t) \\
&\quad + \frac{\sigma^2}{2}k^2\Omega_{kk}(k,n,t) + q(\theta(t))\hat{v}\int \hat{M}((k,n),(k',n'),t)\, dG_n(k',n',t) \\
&\quad + \theta(t)q(\theta(t))n\int \hat{M}((k',n'),(k,n),t)\, dG_{\tilde{v}}(k',n',t) \\
&\quad + \hat{v}\theta(t)q'(\theta(t))\int\int \hat{M}((k'',n''),(k',n'),t)\, dG_n(k',n',t)\, dG_{\tilde{v}}(k'',n'',t)\Bigg\}.
\end{aligned}$$

Reduction to Product-Line Units

Because flow payoffs and state transitions are homogeneous of degree one in (k,n), the marginal social value is homogeneous of degree one. Write

$$\Omega(k,n,t) = kS(x,t), \qquad x \equiv n/k, \qquad \hat{v}(k,n,t) = kv(x,t).$$

Then

$$\Omega_n(k,n,t) = S_x(x,t), \qquad \Omega_k(k,n,t) = S(x,t) - xS_x(x,t), \qquad \Omega_{kk}(k,n,t) = \frac{x^2}{k}S_{xx}(x,t).$$

Using $\alpha(x, x') \equiv \hat{\alpha}(1/x, 1/x')$, the match surplus becomes the normalized expression in (8). Moreover,

$$\int \hat{M}((k,n),(k',n'),t)\, dG_n(k',n',t) = \int M(x,x',t)\, dF_x(x',t),$$

$$\int \hat{M}((k',n'),(k,n),t)\, dG_{\tilde{v}}(k',n',t) = \int M(x',x,t)\, dF_v(x',t),$$

and

$$\begin{aligned} &\int\int \hat{M}((k'',n''),(k',n'),t)\, dG_n(k',n',t)\, dG_{\tilde{v}}(k'',n'',t) \\ &= \int\int M(x'',x',t)\, dF_x(x',t)\, dF_v(x'',t). \end{aligned}$$

Dividing the firm-level HJB by k then yields (17). Differentiating (17) with respect to $v(x,t)$ gives the vacancy first-order condition in Definition 2.

Social Value of Innovation

Finally, because $\mathcal{V}_{\log Y} = 1/\rho$, a one-step innovation of size $\log\lambda$ raises continuation utility by $(\log\lambda)/\rho$. Expressed in the current-value normalization, this gain is $(\log\lambda/\rho)\Psi(t)$. Using

$$\Omega_k(k,n,t) = S(x,t) - xS_x(x,t),$$

the expected marginal social value of innovation is therefore

$$\int_0^\infty [S(x,t) - xS_x(x,t)]\, f(x,t)dx,$$

which gives (19).

B.4 A Hosios-like Benchmark for the Search Margin

This subsection derives a state-dependent Hosios-like benchmark by asking which bargaining weight would make the decentralized vacancy first-order condition coincide with the planner vacancy first-order condition at a given state x. The exercise isolates the search margin only; it does not eliminate the innovation wedge captured by $\varpi(t)$. Strictly speaking, the competitive and planner value functions solve different equations. When needed, we write S^{CE} and M^{CE} for the competitive objects and S^{SP} and M^{SP} for the planner objects. In the expressions below, we suppress these superscripts and evaluate the objects at the

allocation whose vacancy-posting margin is being decentralized.

Equating (9) and (18) at state x and solving for β yields

$$\beta_H(x,t) \equiv -\frac{\theta(t)q'(\theta(t))}{q(\theta(t))} \frac{\int_0^\infty \int_0^\infty M(x'',x',t)dF_x(x',t)dF_v(x'',t)}{\int_0^\infty M(x,x',t)dF_x(x',t)}.$$

For the Cobb–Douglas matching function $q(\theta) = \overline{q}\theta^{\eta-1}$, the benchmark reduces to

$$\beta_H(x,t) = (1-\eta)\frac{\int_0^\infty \int_0^\infty M(x'',x',t)dF_x(x',t)dF_v(x'',t)}{\int_0^\infty M(x,x',t)dF_x(x',t)}.$$

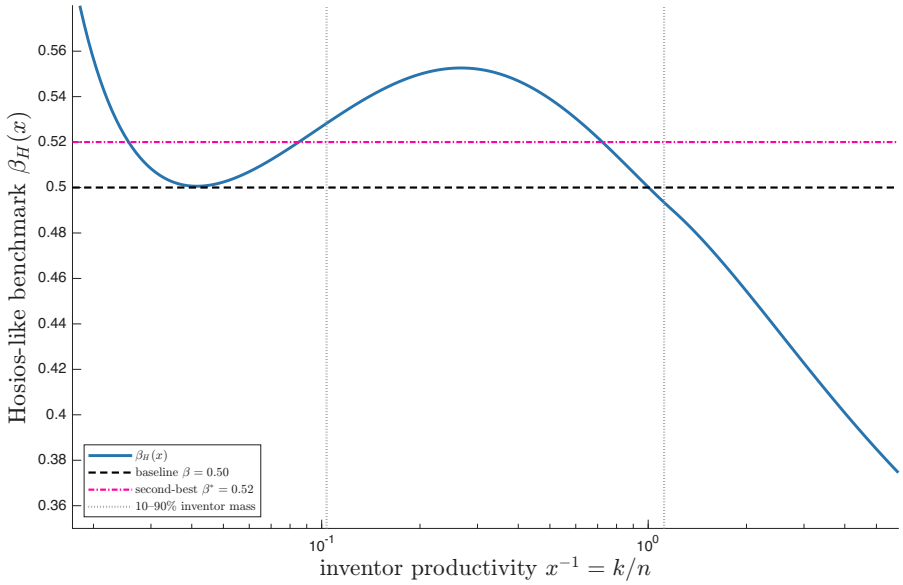
In the homogeneous DMP benchmark, $\int_0^\infty M(x,x',t)dF_x(x',t)$ is constant across states. The double integral in the numerator then equals this same constant, so $\beta_H = 1-\eta$, the usual Hosios condition. In our model, however, the expected match surplus varies with x because match surplus depends on the firm's state and on the distribution of potential partners. Hence $\beta_H(x,t)$ is generally state-dependent.

Figure A.1 evaluates this object at the planner steady state. The benchmark varies across inventor-productivity states, especially in the tails where $A(x)$ is small, so no constant bargaining weight can exactly implement the planner allocation. At the same time, over the 10th to 90th percentiles of the inventor-weighted planner distribution, $\beta_H(x)$ lies roughly between 0.49 and 0.55. This range overlaps the computed second-best bargaining weight and helps explain why a fixed bargaining weight close to 0.52 leaves only a small remaining planner-reform gain in the quantitative exercise, even though it does not decentralize the planner allocation exactly.

Interpretation

This benchmark is therefore not a full decentralization result. First, a constant β can coincide with $\beta_H(x,t)$ only in special cases where the expected match surplus $\int_0^\infty M(x,x',t)dF_x(x',t)$ is constant across states. Second, even then, the innovation wedge in (19) would also have to be absent or separately corrected, because a state-dependent $\beta_H(x,t)$ aligns only the vacancy-posting margin. The computed fixed-β optimum should thus be interpreted as a second-best approximation to this state-dependent search benchmark, not as an exact implementation of the planner allocation.

Figure A.1: State-Dependent Hosios-Like Benchmark



Notes: The solid line plots the Hosios-like benchmark $\beta_H(x)$ at the planner steady state as a function of inventor productivity $z = x^{-1} = k/n$. The horizontal dashed line shows the baseline bargaining power $\beta = 0.5$, the dash-dotted line shows the computed second-best fixed bargaining power $\beta^* = 0.52$, and the dotted vertical lines mark the 10th and 90th percentiles of the inventor-weighted planner distribution. For display, $\beta_H(x)$ is reconstructed from the planner value function using a common centered stencil.
Alt text: A line chart shows the state-dependent Hosios-like bargaining benchmark across inventor productivity. In the central range of the planner distribution, the benchmark is close to 0.5 and overlaps the computed second-best fixed bargaining power of 0.52.

Numerical Diagnostics

In the baseline steady state, the planner-implied normalized social gain from a product line is 0.198, whereas the private instantaneous profit per product line in the competitive equilibrium is 0.231, illustrating that the innovation wedge remains. In the planner-reform metric plotted in Figure 4, moving from the baseline $\beta = 0.5$ to the computed second-best $\beta = 0.52$ lowers the remaining planner-reform gain by about 0.0006 percentage points.

Appendix C Quantitative Appendix

This appendix reports the quantitative inputs and numerical procedures used in the calibration and counterfactual exercises. It first reports the establishment-level inventor-productivity-growth regression results used to discipline the diffusion, internal-R&D, and shock parameters, together with robustness, source-decomposition, source-contraction, and source-side exercises for this moment. It then describes the matching-efficiency path used in the mobility-decline counterfactual. The remaining subsections record the

Table A.5: Inventor Productivity Growth Regressions

	Incumbent inventor productivity growth		
	(1)	(2)	(3)
Knowledge-diffusion inflow term	0.5259***	0.3056***	0.2499***
	(0.0706)	(0.0729)	(0.0887)
Inverse root lagged inventor productivity	0.2546***	0.3465***	0.4400***
	(0.0234)	(0.0243)	(0.0290)
Establishment FE	Yes	Yes	Yes
Year FE	Yes	No	No
Industry × region × year FE	No	Yes	Yes
Establishment-specific linear trend	No	No	Yes
Root MSE	0.702	0.641	0.653
Observations	12,690	10,812	10,812

Notes: The unit of observation is an establishment-year. The dependent variable is incumbent inventor-productivity growth from $t-1$ to t. Inventor inflows are job-to-job moves from $t-1$ to t. The knowledge-diffusion inflow term is total lagged source inventor productivity brought by inventor inflows, divided by lagged destination inventor productivity and the lagged number of destination inventors. Inverse root lagged inventor productivity is $(z_{e,t-1})^{-1/2}$. The net inventor inflow rate is included in all columns but not reported. All columns are weighted by the number of incumbent/stayer inventors used to construct the dependent variable. The regression sample is conditional on a positive knowledge-diffusion inflow term. The dependent variable, the inverse-root lagged inventor-productivity term, and the net inventor inflow rate are trimmed at the 1st and 99th percentiles; the knowledge-diffusion inflow term is trimmed only above the 99th percentile. Industry × region × year fixed effects use industry classifications from the BHP establishment module and federal-state regions. Standard errors clustered by establishment are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

numerical HJB and Kolmogorov forward equations and the consumption-equivalent welfare calculation.

C.1 Inventor-Productivity-Growth Regression Results

This subsection reports the establishment-level regressions used to discipline the inventor-productivity-growth parameters in the quantitative model. The baseline specification provides the empirical moment used for $\overline{\alpha}$, $\overline{\mu}$, and σ; the additional specifications assess robustness to richer fixed effects, zero-exposure establishment-years, and lagged outcome controls. The auxiliary exercises then examine three aspects of the moment: the source-decomposition exercise separates source-establishment productivity from mover quality, the source-contraction exercise uses lagged hiring networks to reduce reliance on contemporaneous destination choices, and the source-side exercise checks whether departures are followed by productivity losses at origin establishments.

The estimation results are presented in Table A.5. We use column (1) as the baseline calibration specification because it corresponds most directly to the model-implied moment

Table A.6: Robustness of Inventor Productivity Growth Regressions

	Zero-exposure sample			Lagged-outcome control	
	(1)	(2)	(3)	(4)	(5)
Knowledge-diffusion inflow term	0.3151***	0.1803***	0.1213**	0.3127***	0.2491**
	(0.0519)	(0.0514)	(0.0613)	(0.0811)	(0.1001)
Lagged incumbent growth				-0.1902***	-0.1896***
				(0.0108)	(0.0218)
Establishment FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	No	No	No	No
Industry × region × year FE	No	Yes	Yes	Yes	Yes
Establishment-specific linear trend	No	No	Yes	No	Yes
Observations	32,588	29,907	29,907	8,860	8,860

Notes: The unit of observation is an establishment-year. The dependent variable is incumbent inventor-productivity growth from $t-1$ to t. Columns (1)–(3) add establishment-years with zero diffusion exposure; columns (4) and (5) use the positive-exposure sample and add lagged incumbent growth. All specifications include inverse root lagged inventor productivity and the net inventor inflow rate, are weighted by the number of incumbent/stayer inventors used to construct the dependent variable, and have standard errors clustered by establishment. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

after controlling for establishment and aggregate year effects. The more saturated specifications in columns (2) and (3) assess sensitivity to local industry-region-year shocks and establishment-specific trends. Using column (1), the baseline calibration sets the monthly model parameters as follows:

$$\overline{\alpha} = 0.526, \qquad \overline{\mu} = \frac{0.2546}{12} = 0.0212, \qquad \sigma = \frac{0.7020}{\sqrt{12}} = 0.203,$$

The diffusion coefficient is taken directly from the coefficient on the annual diffusion exposure, while the divisions by 12 and $\sqrt{12}$ convert the internal-R&D drift and shock standard deviation into monthly model units.

Table A.6 reports two checks on the inventor-productivity-growth moment. The baseline estimates in Table A.5 use establishment-years with positive diffusion exposure. Columns (1)–(3) expand the sample to include establishment-years with no inventor inflow between $t-1$ and t, setting the diffusion exposure to zero when lagged destination inventor productivity and inventor employment are observed. Columns (4) and (5) return to the positive-exposure sample and add lagged incumbent-inventor productivity growth, checking whether the diffusion coefficient is driven by persistence in destination productivity growth. The coefficient on the diffusion exposure remains positive in both exercises.

Table A.7 separates the source-productivity inflow term from the moving inventors' own lagged productivity. The leave-one-mover-out source term recomputes source-

Table A.7: Source Productivity and Mover Quality in Inventor Productivity Growth

	(1)	(2)	(3)
Source-productivity inflow term	0.4779***	0.2733***	0.2099**
	(0.0703)	(0.0796)	(0.0994)
Mover-own-productivity inflow term	0.0479	0.0399	0.0362
	(0.0302)	(0.0389)	(0.0492)
Establishment FE	Yes	Yes	Yes
Year FE	Yes	No	No
Industry × region × year FE	No	Yes	Yes
Establishment-specific linear trend	No	No	Yes
Observations	12,628	10,757	10,757

Notes: The unit of observation is an establishment-year. The dependent variable is incumbent inventor-productivity growth from $t-1$ to t. The source-productivity inflow term uses a leave-one-mover-out measure that excludes the moving inventor from the source-establishment productivity average. The mover-own-productivity inflow term is constructed from the moving inventors' own lagged productivity. All columns include inverse root lagged inventor productivity and the net inventor inflow rate. Regressions are weighted by the number of incumbent/stayer inventors used to construct the dependent variable. Standard errors clustered by establishment are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

establishment productivity after removing the moving inventor's own lagged citation productivity from the source average. This removes the mechanical overlap between the source-productivity measure and the mover-own-productivity term. The leave-one-mover-out source-productivity coefficient remains positive and sizable after directly controlling for mover own productivity, while the mover-own-productivity coefficient is small and imprecise. We therefore interpret the baseline coefficient as a source-productivity moment, not simply as a mover-quality premium.

As a complementary source-shock exercise, Table A.8 replaces contemporaneous realized destination choices with lagged inventor hiring links. We define a severe source-employment contraction as an 80% decline in BHP establishment employment between adjacent June 30 reference dates. We then allocate these source-year shocks to potential destinations using inventor job-to-job hiring links from two to six years before the shock and construct a predicted source-productivity-weighted exposure with the same destination normalization as the baseline knowledge-diffusion inflow term.

The predicted exposure is positively associated with incumbent inventor-productivity growth at destinations across the three fixed-effect specifications. The association is also positive when we additionally control for the predicted displaced-inventor-count exposure, indicating that the pattern is not driven only by the predicted number of incoming inventors. The predicted exposure is also positively related to the baseline knowledge-diffusion inflow term in a first-stage-style regression. This evidence supports a causal interpretation of the baseline knowledge-diffusion moment by reducing reliance on realized destination

Table A.8: Lagged-Network Source-Contraction Exposure Regressions

	(1)	(2)	(3)
Panel A. Reduced-form: incumbent inventor-productivity growth			
Predicted severe-contraction exposure	0.0515**	0.0482**	0.0449
	(0.0200)	(0.0187)	(0.0295)
Predicted exposure, count control	0.0708***	0.0611***	0.0612*
	(0.0234)	(0.0219)	(0.0338)
Panel B. First-stage-style: baseline knowledge-diffusion inflow term			
Predicted severe-contraction exposure	0.0236**	0.0245*	0.0363
	(0.0117)	(0.0125)	(0.0231)
Establishment FE	Yes	Yes	Yes
Year FE	Yes	No	No
Industry × region × year FE	No	Yes	Yes
Establishment-specific linear trend	No	No	Yes
Observations	26,521	24,633	24,633

Notes: The unit of observation is an establishment-year. The predicted exposure is constructed from severe source-employment contractions allocated through lagged inventor hiring links and is not used as the calibrated estimate of $\overline{\alpha}$. Panel A uses incumbent inventor-productivity growth from $t-1$ to t as the dependent variable; the count-control row additionally controls for predicted displaced-inventor-count exposure. Panel B uses the baseline knowledge-diffusion inflow term as the dependent variable. All specifications include inverse root lagged inventor productivity and the net inventor inflow rate, are weighted by the number of incumbent/stayer inventors used to construct the productivity-growth outcome, and have standard errors clustered by establishment. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

choices.

Finally, Table A.9 turns to the source side and examines whether inventor departures are associated with subsequent productivity losses at the source establishment. These source-side regressions are descriptive rather than causal, since inventor departures are endogenous. The coefficients on inventor outflow rates are small and imprecise. Quality-weighted outflow coefficients are positive and statistically significant in the specifications with industry-by-region-by-year controls. Thus, these regressions do not show clear evidence that inventor departures reduce subsequent incumbent-inventor productivity growth at origin establishments.

C.2 Counterfactual Matching-Efficiency Path

The mobility-decline counterfactual uses a deterministic path of matching efficiency so that the model matches the observed decline in German inventor job-to-job transition rates. Following Akcigit and Ates (2023), we assume that the path of matching efficiency takes the form

$$\overline{q}_t = \overline{q}_0 + \frac{\exp\left(-(t/T)\nu\right) - 1}{\exp\left(-\nu\right) - 1}\left(\overline{q}_T - \overline{q}_0\right), \quad \text{for all } t \in [0, T].$$

Table A.9: Source-Side Productivity-Loss Diagnostics

	Next-period source incumbent productivity growth		
	(1)	(2)	(3)
Inventor outflow rate	0.0718	-0.1392	-0.0326
	(0.2099)	(0.1928)	(0.2274)
Quality-weighted inventor outflow term	0.0939	1.0918***	1.0371***
	(0.0888)	(0.0852)	(0.0925)
Establishment FE	Yes	Yes	Yes
Year FE	Yes	No	No
Industry × region × year FE	No	Yes	Yes
Establishment-specific linear trend	No	No	Yes
Observations: outflow-rate sample	33,179	30,488	30,488
Observations: quality-weighted sample	33,161	30,467	30,467

Notes: The unit of observation is an establishment-year. The dependent variable is next-period incumbent inventor-productivity growth at the source establishment. The two reported outflow variables are estimated in separate regressions with the same fixed-effect structure. The inventor outflow rate is the number of job-to-job inventor departures divided by the number of source-establishment inventors. The quality-weighted inventor outflow term is the summed productivity of departing inventors, divided by source-establishment inventor productivity and the number of source-establishment inventors. Each row reports separate regressions. All regressions include inverse root source inventor productivity and the source net inventor inflow rate, are weighted by the number of next-period incumbent/stayer inventors used to construct the dependent variable, and have standard errors clustered by establishment. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

During the period from $t = 0$ to $t = T$, matching efficiency moves smoothly from $\overline{q}_0$ to $\overline{q}_T$, with ν governing the speed of adjustment. Until $t = 0$, agents do not anticipate the changes in $\overline{q}_t$; after $t = 0$, they have perfect foresight. After $t = T$, matching efficiency remains fixed at $\overline{q}_T$, but the economy need not be on the new BGP at $t = T$ and continues to converge thereafter.

In the exercise, $t = 0$ corresponds to 2001 and $t = T$ corresponds to 2014. We set $\overline{q}_0$ equal to the value obtained from the baseline steady-state calibration in Section 4, which targets the 2001 job-to-job transition rate. The remaining parameters, ν and $\overline{q}_T$, are chosen so that the model matches the German inventor job-to-job transition rates in 2007 and 2014. The targeted annual rates are 8.3% in 2007 and 4.8% in 2014, and the calibrated values are $\nu = 0.583$ and $\overline{q}_T = 0.00861$, compared with $\overline{q}_0 = 0.0198$.

C.3 Numerical HJB Equation

This subsection records the implicit update used to solve the HJB equation. The derivation has three steps. First, we rewrite the nonlocal surplus terms so that the local terms in S, S_x, and S_{xx} can be treated implicitly while the remaining nonlocal terms are evaluated using the previous iterate. Second, we express the update in the log state. Third, we discretize

the resulting equation on an evenly spaced log-state grid.

The two nonlocal surplus terms can be written as

$$\begin{aligned}
&\int \left[S_x(x) + \alpha\left(x, x'\right)\left\{S(x) - xS_x(x)\right\} - S_x(x')\right]^+ dF_x(x') \\
&= \int \mathbb{1}_P(x, x')\left[\alpha\left(x, x'\right) S(x) + \left\{x^{-1} - \alpha\left(x, x'\right)\right\} xS_x(x) - S_x(x')\right] dF_x(x') \\
&= \int \mathbb{1}_P(x, x')\alpha\left(x, x'\right) dF_x(x')S(x) \\
&\quad + \int \mathbb{1}_P(x, x')\left\{x^{-1} - \alpha\left(x, x'\right)\right\} dF_x(x')xS_x(x) \\
&\quad - \int \mathbb{1}_P(x, x')S_x(x')dF_x(x'),
\end{aligned}$$

and

$$\begin{aligned}
&\int \left[S_x(x') + \alpha\left(x', x\right)\left\{S(x') - x'S_x(x')\right\} - S_x(x)\right]^+ dF_v(x') \\
&= \int \mathbb{1}_P(x', x)\left[\alpha\left(x', x\right) S(x') + \left\{1 - \alpha\left(x', x\right) x'\right\} S_x(x') - S_x(x)\right] dF_v(x') \\
&= -\int \mathbb{1}_P(x', x)dF_v(x')S_x(x) \\
&\quad + \int \mathbb{1}_P(x', x)\left\{\alpha\left(x', x\right) S(x') + \left\{1 - \alpha\left(x', x\right) x'\right\} S_x(x')\right\} dF_v(x').
\end{aligned}$$

Then, the HJB equation is rewritten as

$$\begin{aligned}
\rho S(x) =& \pi - c\left(v(x)\right) \\
&+ \left\{\mu(x) - \delta\right\}\left\{S(x) - xS_x(x)\right\} + \frac{\sigma^2}{2}x^2S_{xx}(x) \\
&+ \left\{(1-\beta)q(\theta)v(x)\int \mathbb{1}_P(x, x')\alpha\left(x, x'\right) dF_x(x')\right\} S(x) \\
&+ \left[(1-\beta)q(\theta)v(x)\int \mathbb{1}_P(x, x')\left\{x^{-1} - \alpha\left(x, x'\right)\right\} dF_x(x')\right. \\
&\qquad \left. - \beta\theta q(\theta)\int \mathbb{1}_P(x', x)dF_v(x')\right] xS_x(x) \\
&- (1-\beta)q(\theta)v(x)\int \mathbb{1}_P(x, x')S_x(x')dF_x(x') \\
&+ \beta\theta q(\theta)x\int \mathbb{1}_P(x', x)\left\{\alpha\left(x', x\right) S(x') + \left\{1 - \alpha\left(x', x\right) x'\right\} S_x(x')\right\} dF_v(x').
\end{aligned}$$

We solve the HJB equation using an implicit method. Let Δ_{num} denote the pseudo-time step used in the numerical iteration, and let τ index the iteration. Then, given $S^{\tau-1}(x)$, the implicit method gives an update

$$
\begin{aligned}
&\rho S^{\tau}(x) - \frac{1}{\Delta_{\text{num}}}\left[S^{\tau-1}(x) - S^{\tau}(x)\right] \\
&= \pi - c\left(v(x)\right) + \left\{\mu(x) - \delta\right\}\left\{S^{\tau}(x) - xS_x^{\tau}(x)\right\} + \frac{\sigma^2}{2}x^2 S_{xx}^{\tau}(x) \\
&\quad + (1-\beta)q(\theta)v(x)\left[\int \mathbb{1}_P(x,x')\alpha(x,x')\,dF_x(x')\right] S^{\tau}(x) \\
&\quad + \Bigg[(1-\beta)q(\theta)v(x)\int \mathbb{1}_P(x,x')\left\{x^{-1} - \alpha(x,x')\right\}\,dF_x(x') \\
&\qquad\quad - \beta\theta q(\theta)\int \mathbb{1}_P(x',x)\,dF_v(x')\Bigg] x S_x^{\tau}(x) \\
&\quad - (1-\beta)q(\theta)v(x)\int \mathbb{1}_P(x,x')S_x^{\tau-1}(x')\,dF_x(x') \\
&\quad + \beta\theta q(\theta)x\int \mathbb{1}_P(x',x)\Big[\alpha(x',x)S^{\tau-1}(x') + \left\{1 - \alpha(x',x)x'\right\}S_x^{\tau-1}(x')\Big]\,dF_v(x').
\end{aligned}
$$

$$
\begin{aligned}
&\left[\frac{1}{\Delta_{\text{num}}} + \rho - \left\{\mu(x) - \delta\right\} - (1-\beta)q(\theta)v(x)\int \mathbb{1}_P(x,x')\alpha(x,x')\,dF_x(x')\right] S^{\tau}(x) \\
&\quad - \Bigg[-\mu(x) + \delta + (1-\beta)q(\theta)v(x)\int \mathbb{1}_P(x,x')\left\{x^{-1} - \alpha(x,x')\right\}\,dF_x(x') \\
&\qquad\quad - \beta\theta q(\theta)\int \mathbb{1}_P(x',x)\,dF_v(x')\Bigg] x S_x^{\tau}(x) - \frac{\sigma^2}{2}x^2 S_{xx}^{\tau}(x) \\
&= \frac{1}{\Delta_{\text{num}}}S^{\tau-1}(x) + \pi - c\left(v(x)\right) - (1-\beta)q(\theta)v(x)\int \mathbb{1}_P(x,x')S_x^{\tau-1}(x')\,dF_x(x') \\
&\quad + \beta\theta q(\theta)x\int \mathbb{1}_P(x',x)\Big[\alpha(x',x)S^{\tau-1}(x') + \left\{1 - \alpha(x',x)x'\right\}S_x^{\tau-1}(x')\Big]\,dF_v(x').
\end{aligned}
$$

Let $\tilde{x} = \log x$. Then,

$$\begin{aligned}
&\left[\frac{1}{\Delta_{\text{num}}} + \rho - \left\{\tilde{\mu}\left(\tilde{x}\right) - \delta\right\} - (1-\beta)q(\theta)\tilde{v}\left(\tilde{x}\right)\int \mathbb{1}_P(\tilde{x},\tilde{x}')\tilde{\alpha}\left(\tilde{x},\tilde{x}'\right)\, d\tilde{F}_{\tilde{x}}(\tilde{x}')\right]\tilde{S}^{\tau}(\tilde{x}) \\
&- \Bigg[-\left\{\tilde{\mu}\left(\tilde{x}\right) - \delta + \tfrac{\sigma^2}{2}\right\} + (1-\beta)q(\theta)\tilde{v}\left(\tilde{x}\right)\int \mathbb{1}_P(\tilde{x},\tilde{x}') \\
&\qquad \left\{\exp\left(-\tilde{x}\right) - \tilde{\alpha}\left(\tilde{x},\tilde{x}'\right)\right\}\, d\tilde{F}_{\tilde{x}}(\tilde{x}') - \beta\theta q(\theta)\int \mathbb{1}_P(\tilde{x}',\tilde{x})\, d\tilde{F}_{\tilde{v}}(\tilde{x}')\Bigg]\tilde{S}^{\tau}_{\tilde{x}}(\tilde{x}) \\
&- \frac{\sigma^2}{2}\tilde{S}^{\tau}_{\tilde{x}\tilde{x}}(\tilde{x}) \\
&= \frac{1}{\Delta_{\text{num}}}\tilde{S}^{\tau-1}(\tilde{x}) + \pi - c\left(\tilde{v}\left(\tilde{x}\right)\right) \\
&\quad - (1-\beta)q(\theta)\tilde{v}\left(\tilde{x}\right)\int \mathbb{1}_P(\tilde{x},\tilde{x}')\exp\left(-\tilde{x}'\right)\tilde{S}^{\tau-1}_{\tilde{x}}(\tilde{x}')\, d\tilde{F}_{\tilde{x}}(\tilde{x}') \\
&\quad + \beta\theta q(\theta)\exp\left(\tilde{x}\right)\int \mathbb{1}_P(\tilde{x}',\tilde{x})\Big[\tilde{\alpha}\left(\tilde{x}',\tilde{x}\right)\tilde{S}^{\tau-1}(\tilde{x}') \\
&\qquad + \left\{\exp\left(-\tilde{x}'\right) - \tilde{\alpha}\left(\tilde{x}',\tilde{x}\right)\right\}\tilde{S}^{\tau-1}_{\tilde{x}}(\tilde{x}')\Big]\, d\tilde{F}_{\tilde{v}}(\tilde{x}'). \qquad (22)
\end{aligned}$$

We now discretize $\tilde{x}$ on an evenly spaced N_x grid, ordered as

$$\begin{pmatrix} \tilde{x}_1 \\ \tilde{x}_2 \\ \vdots \\ \tilde{x}_{N_x} \end{pmatrix}.$$

Then, (22) can be rewritten in vector form as follows:

$$H_1 S^{\tau} - H_2 S^{\tau}_{\tilde{x}} - \frac{\sigma^2}{2}S^{\tau}_{\tilde{x}\tilde{x}} = \frac{1}{\Delta_{\text{num}}}S^{\tau-1} + H_0,$$

where $S^{\tau}_i = \tilde{S}^{\tau}(\tilde{x}_i)$, $(S^{\tau}_{\tilde{x}})_i = \tilde{S}^{\tau}_{\tilde{x}}(\tilde{x}_i)$, and $(S^{\tau}_{\tilde{x}\tilde{x}})_i = \tilde{S}^{\tau}_{\tilde{x}\tilde{x}}(\tilde{x}_i)$. The ith element of the N_x-vector H_0 is

$$\begin{aligned}
H_{0,i} = {} & \pi - c\left(\tilde{v}\left(\tilde{x}\right)\right) - (1-\beta)q(\theta)\tilde{v}\left(\tilde{x}\right)\int \mathbb{1}_P(\tilde{x},\tilde{x}')\exp\left(-\tilde{x}'\right)\tilde{S}^{\tau-1}_{\tilde{x}}(\tilde{x}')\, d\tilde{F}_{\tilde{x}}(\tilde{x}') \\
& + \beta\theta q(\theta)\exp\left(\tilde{x}\right)\int \mathbb{1}_P(\tilde{x}',\tilde{x})\Big[\tilde{\alpha}\left(\tilde{x}',\tilde{x}\right)\tilde{S}^{\tau-1}(\tilde{x}') + \left\{\exp\left(-\tilde{x}'\right) - \tilde{\alpha}\left(\tilde{x}',\tilde{x}\right)\right\}\tilde{S}^{\tau-1}_{\tilde{x}}(\tilde{x}')\Big]\, d\tilde{F}_{\tilde{v}}(\tilde{x}'),
\end{aligned}$$

Here, H_0 collects the flow payoff and the nonlocal terms that are evaluated explicitly using

the previous iterate $S^{\tau-1}$. By contrast, the vectors h_1 and h_2 collect the coefficients on, respectively, the zeroth-order and first-derivative terms in the unknown current iterate S^{τ}. With this split, the nonlocal terms remain on the right-hand side, while the local differential operator is treated implicitly. Let h_1 be the N_x-vector with elements

$$\frac{1}{\Delta_{\text{num}}} + \rho - \left\{\tilde{\mu}\left(\tilde{x}\right) - \delta\right\} - (1-\beta)q(\theta)\tilde{v}\left(\tilde{x}\right) \int \mathbb{1}_P(\tilde{x}, \tilde{x}')\tilde{\alpha}\left(\tilde{x}, \tilde{x}'\right)\, d\tilde{F}_{\tilde{x}}(\tilde{x}'),$$

and let h_2 be the N_x-vector with elements

$$-\left\{\tilde{\mu}\left(\tilde{x}\right) - \delta + \frac{\sigma^2}{2}\right\} + (1-\beta)q(\theta)\tilde{v}\left(\tilde{x}\right) \int \mathbb{1}_P(\tilde{x}, \tilde{x}') \left\{\exp\left(-\tilde{x}\right) - \tilde{\alpha}\left(\tilde{x}, \tilde{x}'\right)\right\}\, d\tilde{F}_{\tilde{x}}(\tilde{x}')$$
$$- \beta\theta q(\theta) \int \mathbb{1}_P(\tilde{x}', \tilde{x})\, d\tilde{F}_{\tilde{v}}(\tilde{x}').$$

In the definitions of $H_{0,i}$, $h_{1,i}$, and $h_{2,i}$, the unprimed state is evaluated at $\tilde{x} = \tilde{x}_i$. Define the diagonal matrices $H_1 = \text{diag}(h_1)$ and $H_2 = \text{diag}(h_2)$.

Let $D_{\tilde{x}}$ be the $N_x \times N_x$ matrix that, when pre-multiplying S^{τ}, gives an approximation of $S^{\tau}_{\tilde{x}}$, and let $D_{\tilde{x}\tilde{x}}$ be the $N_x \times N_x$ matrix that, when pre-multiplying S^{τ}, gives an approximation of $S^{\tau}_{\tilde{x}\tilde{x}}$:

$$S^{\tau}_{\tilde{x}} = D_{\tilde{x}} S^{\tau}$$
$$S^{\tau}_{\tilde{x}\tilde{x}} = D_{\tilde{x}\tilde{x}} S^{\tau}.$$

To compute the derivative matrix $D_{\tilde{x}}$, we follow an upwind scheme. That is, we use a forward approximation when the drift of the state variable is positive, and a backward approximation when the drift of the state variable is negative.

At the boundary nodes, we close the second-derivative operator with a curvature-matching rule. Let $S^{\tau-1}$ denote the previous iterate. We replace the first and last rows of $D_{\tilde{x}\tilde{x}}$ by

$$D_{\tilde{x}\tilde{x}}(1, \cdot) = D_{\tilde{x}\tilde{x}}(2, \cdot)\frac{S^{\tau-1}_1}{S^{\tau-1}_2}, \qquad D_{\tilde{x}\tilde{x}}(N_x, \cdot) = D_{\tilde{x}\tilde{x}}(N_x - 1, \cdot)\frac{S^{\tau-1}_{N_x}}{S^{\tau-1}_{N_x-1}}.$$

Equivalently, the normalized curvature $D_{\tilde{x}\tilde{x}}S/S$ at the boundary is set equal to that at the adjacent interior node. This provides a smooth extrapolation of the transformed surplus in the tail and, on the maintained grid, has negligible effects on steady-state aggregates.

The implicit method works by updating S^τ through the following equation:

$$S^\tau = \left\{ H_1 - H_2 D_{\tilde{x}} - \frac{\sigma^2}{2} D_{\tilde{x}\tilde{x}} \right\}^{-1} \left\{ \frac{1}{\Delta_{\text{num}}} S^{\tau-1} + H_0 \right\}.$$

Thus, each HJB iteration reduces to solving a linear system with coefficient matrix $H_1 - H_2 D_{\tilde{x}} - \frac{\sigma^2}{2} D_{\tilde{x}\tilde{x}}$.

C.4 Numerical Kolmogorov Forward Equation

This subsection records the forward equation used to update the product-line density on the same log-state grid. Given the policy functions and match indicators from the HJB step, the density is updated with an implicit scheme using the drift of the log state.

Applying the log-state transformation used in the KF equation, define the drift term entering the forward operator for the product-line density $\tilde{f}(\tilde{x}, t)$ as

$$\begin{aligned} b_{\tilde{x}}(\tilde{x}, t) = q(\theta(t))\tilde{v}(\tilde{x}, t) \int \mathbb{1}_P(\tilde{x}, \tilde{x}', t) \left\{ \exp(-\tilde{x}) - \tilde{\alpha}(\tilde{x}, \tilde{x}') \right\} d\tilde{F}_{\tilde{x}}(\tilde{x}', t) \\ - \theta(t) q(\theta(t)) \int \mathbb{1}_P(\tilde{x}', \tilde{x}, t)\, d\tilde{F}_{\tilde{v}}(\tilde{x}', t) - \left\{ \tilde{\mu}(\tilde{x}) - \delta(t) + \frac{\sigma^2}{2} \right\}. \end{aligned}$$

That is, $b_{\tilde{x}}(\tilde{x}, t)$ is the net drift of the log state. The first term captures the effect of successful poaching inflows, the second term captures inventor outflows to other firms, and the third term collects internal R&D, creative destruction, and the Ito correction from the log-state transformation. Its sign determines the direction of mass transport in the forward equation and therefore the upwind discretization used below. As with the HJB equation, given $\tilde{f}^{\tau-1}(\tilde{x})$, let $b_{\tilde{x}}^{\tau-1}(\tilde{x})$ denote the drift term above evaluated with $\tilde{F}_{\tilde{x}}^{\tau-1}$ and $\tilde{F}_{\tilde{v}}^{\tau-1}$. The implicit method provides the update:

$$\begin{aligned} &\frac{1}{\Delta_{\text{num}}} \left[\tilde{f}^{\tau-1}(\tilde{x}) - \tilde{f}^{\tau}(\tilde{x}) \right] \\ &\quad = \left\{ q(\theta)\tilde{v}(\tilde{x}) \int \mathbb{1}_P(\tilde{x}, \tilde{x}') \tilde{\alpha}(\tilde{x}, \tilde{x}')\, d\tilde{F}_{\tilde{x}}^{\tau-1}(\tilde{x}') + \tilde{\mu}(\tilde{x}) - \delta \right\} \tilde{f}^{\tau}(\tilde{x}) \\ &\qquad - \frac{\partial}{\partial \tilde{x}} \left[b_{\tilde{x}}^{\tau-1}(\tilde{x}) \tilde{f}^{\tau}(\tilde{x}) \right] + \frac{\sigma^2}{2} \frac{\partial^2}{\partial \tilde{x}^2} \tilde{f}^{\tau}(\tilde{x}). \end{aligned}$$

This can be rewritten in vector form as follows:

$$
\begin{aligned}
-\frac{1}{\Delta_{\text{num}}}\left(f^{\tau}-f^{\tau-1}\right) &= B_1 f^{\tau} - D_{\tilde{x}} B_2 f^{\tau} + \frac{\sigma^2}{2} D_{\tilde{x}\tilde{x}} f^{\tau} \\
\frac{1}{\Delta_{\text{num}}} f^{\tau-1} &= \left\{\frac{1}{\Delta_{\text{num}}} I_{N_x} + B_1 - D_{\tilde{x}} B_2 + \frac{\sigma^2}{2} D_{\tilde{x}\tilde{x}}\right\} f^{\tau} \\
f^{\tau} &= \left\{\frac{1}{\Delta_{\text{num}}} I_{N_x} + B_1 - D_{\tilde{x}} B_2 + \frac{\sigma^2}{2} D_{\tilde{x}\tilde{x}}\right\}^{-1} \frac{1}{\Delta_{\text{num}}} f^{\tau-1},
\end{aligned}
$$

where $f_i^{\tau} = \tilde{f}^{\tau}(\tilde{x}_i)$. The ith diagonal element of B_1 is

$$
q(\theta)\tilde{v}\left(\tilde{x}\right) \int \mathbb{1}_P(\tilde{x}, \tilde{x}')\tilde{\alpha}\left(\tilde{x}, \tilde{x}'\right)\, d\tilde{F}_{\tilde{x}}^{\tau-1}(\tilde{x}') + \tilde{\mu}\left(\tilde{x}\right) - \delta,
$$

and the ith diagonal element of B_2 is

$$
\begin{aligned}
& q(\theta)\tilde{v}\left(\tilde{x}\right) \int \mathbb{1}_P(\tilde{x}, \tilde{x}') \left\{\exp\left(-\tilde{x}\right) - \tilde{\alpha}\left(\tilde{x}, \tilde{x}'\right)\right\}\, d\tilde{F}_{\tilde{x}}^{\tau-1}(\tilde{x}') \\
& - \theta q(\theta) \int \mathbb{1}_P(\tilde{x}', \tilde{x})\, d\tilde{F}_{\tilde{v}}^{\tau-1}(\tilde{x}') - \left\{\tilde{\mu}\left(\tilde{x}\right) - \delta + \frac{\sigma^2}{2}\right\}.
\end{aligned}
$$

In the definitions of $B_{1,ii}$ and $B_{2,ii}$, the unprimed state is evaluated at $\tilde{x} = \tilde{x}_i$. To construct the derivative matrices for the forward equation, we use a backward approximation when the drift of the state variable is positive, and a forward approximation when the drift of the state variable is negative. This is the adjoint counterpart of the upwind choice in the HJB equation: the HJB discretizes the derivative of the value function along characteristics, whereas the KF equation discretizes the transport of probability mass.

C.5 Consumption-Equivalent Welfare Gains

This subsection describes how we compute welfare from the simulated transition paths. We first recover consumption from output and vacancy costs, then attach a balanced-growth tail after the finite transition horizon, and finally convert welfare differences into consumption-equivalent units.

Consumption Path $\{C(t)\}_{t\geq 0}$

From the resource constraint (13),

$$
\log C(t) = \log Y(t) + \log\left\{1 - \int_0^{\infty} c(v(x,t))dF(x,t)\right\}.
$$

Inserting (2), we obtain

$$\log C(t) = \log \lambda \int_0^t \delta(\tau) d\tau + \log L + \log \left\{ 1 - \int_0^\infty c(v(x,t)) dF(x,t) \right\}. \tag{23}$$

Therefore, we calculate the consumption path $\{C(t)\}_{t\geq 0}$ from the path of creative destruction rate $\{\delta(t)\}_{t\geq 0}$ and the current vacancy posting $v(x,t)$ and distribution $F(x,t)$ using (23). Because $\log L$ is time-invariant, it cancels out of all consumption-equivalent welfare comparisons.

Welfare $V(\{C(t)\}_{t\geq 0})$

Let T_{sim} denote the finite transition horizon used in the simulation. Assume that after T_{sim}, $C(t)$ grows at the economic growth rate in the final steady state $\overline{g}$. Then, we calculate $V(\{C(t)\}_{t\geq 0})$ numerically from $\{C(t)\}_{t\in[0,T_{\text{sim}}]}$ and $\overline{g}$ as follows:

$$\begin{aligned}
V(\{C(t)\}_{t\geq 0}) &= \int_0^\infty e^{-\rho t} \log C(t) dt \\
&= \int_0^{T_{\text{sim}}} e^{-\rho t} \log C(t) dt + \int_{T_{\text{sim}}}^\infty e^{-\rho t} \log C(t) dt \\
&= \int_0^{T_{\text{sim}}} e^{-\rho t} \log C(t) dt + \int_{T_{\text{sim}}}^\infty e^{-\rho t} \log \left\{ C(T_{\text{sim}}) e^{(t-T_{\text{sim}})\overline{g}} \right\} dt \\
&= \int_0^{T_{\text{sim}}} e^{-\rho t} \log C(t) dt + \frac{1}{\rho} e^{-\rho T_{\text{sim}}} \left\{ \log C(T_{\text{sim}}) + \frac{\overline{g}}{\rho} \right\}.
\end{aligned}$$

For a balanced growth path with constant growth rate g and constant vacancy cost share cv, this expression simplifies to

$$V^{ss} = \frac{\log L + \log(1-cv)}{\rho} + \frac{g}{\rho^2}.$$

Since $\log L$ is common across the steady-state benchmark comparisons in the paper, those comparisons depend only on

$$\frac{\log(1-cv)}{\rho} + \frac{g}{\rho^2}.$$

Consumption-Equivalent Welfare Gain $\mathcal{L}$

Let $\{C(t)\}_{t\geq 0}$ denote the consumption path in the economy without policy change and $\{\widetilde{C}(t)\}_{t\geq 0}$ denote the consumption path in the economy with policy change. The consumption-equivalent welfare gain from the policy change is the scalar $\mathcal{L}$ such that the household is indifferent between the consumption path $\{\mathcal{L}C(t)\}_{t\geq 0}$ and the policy-change consumption path $\{\widetilde{C}(t)\}_{t\geq 0}$.

Let $V\left(\{C(t)\}_{t\geq 0}\right)$ denote welfare:

$$V\left(\{C(t)\}_{t\geq 0}\right) \equiv \int_0^\infty e^{-\rho t}\log C(t)dt.$$

Then,

$$\begin{aligned} V\left(\mathcal{L}\{C(t)\}_{t\geq 0}\right) &= \int_0^\infty e^{-\rho t}\log(\mathcal{L}C(t))dt \\ &= \frac{\log\mathcal{L}}{\rho} + \int_0^\infty e^{-\rho t}\log C(t)dt \\ &= \frac{\log\mathcal{L}}{\rho} + V\left(\{C(t)\}_{t\geq 0}\right). \end{aligned}$$

The indifference condition $V\left(\mathcal{L}\{C(t)\}_{t\geq 0}\right) = V\left(\{\widetilde{C}(t)\}_{t\geq 0}\right)$ therefore implies

$$\frac{\log\mathcal{L}}{\rho} = V\left(\{\widetilde{C}(t)\}_{t\geq 0}\right) - V\left(\{C(t)\}_{t\geq 0}\right),$$

so the consumption-equivalent welfare gain from policy change is calculated as

$$\mathcal{L} = \exp\left[\rho\left\{V\left(\left\{\widetilde{C}(t)\right\}_{t\geq 0}\right) - V\left(\{C(t)\}_{t\geq 0}\right)\right\}\right].$$

Welfare changes reported in percent are $100(\mathcal{L}-1)$.