NBER WORKING PAPER SERIES

GOOD RENTS VERSUS BAD RENTS:
R&D MISALLOCATION AND GROWTH

Philippe Aghion
Antonin Bergeaud
Timo Boppart
Peter J. Klenow
Huiyu Li

Working Paper 34190
http://www.nber.org/papers/w34190

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2025

We thank Nicolas Crouzet, Maarten De Ridder, Yifan Zhu, and participants at the Conference in Memory of the work of Emmanuel Farhi, the NBER Summer Institute Macroeconomics and Productivity workshop, and the NBER EFG meeting for their helpful comments. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco, the Board of Governors of the Federal Reserve System, or the National Bureau of Economic Research. This work is supported by a public grant overseen by the French National Research Agency (ANR) as part of the "Investissements d'Avenir" program (reference: ANR-10-EQPX-17 – Centre d'accès securisé aux données – CASD), and has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 786587). Philippe Aghion also acknowledges the support of EUR grant ANR-17-EURE-0001. Timo Boppart acknowledges the support of SNSF consolidator grant TMCG-1 223114. Erin Crust provided excellent research assistance.

Good Rents versus Bad Rents: R&D Misallocation and Growth
Philippe Aghion, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li
NBER Working Paper No. 34190
August 2025
JEL No. O31, O38, O41, O52

ABSTRACT

Firm price-cost markups may reflect (a) bigger step sizes from quality innovations that confer significant knowledge spillovers onto other firms, and/or (b) higher process efficiency than competing firms or other factors which bear no obvious knowledge externality. We write down an endogenous growth model with innovation step size and process efficiency as alternative sources of markup heterogeneity. Compared with the laissez-faire equilibrium, the social planner wants to reallocate research towards high step size firms but not high process efficiency firms. We then use price and productivity data across firms in French manufacturing to infer firm step sizes and process efficiency. We find that the planner could achieve faster growth by reallocating research toward high step size firms, and more so if high step size firms could freely license their innovations to high process efficiency firms.

Philippe Aghion
Collège de France
and INSEAD, and London School
of Economics
philippe.aghion@insead.edu

Antonin Bergeaud
HEC Paris
antonin.bergeaud@gmail.com

Timo Boppart
IIES, Stockholm University
Department of Economics
timo.boppart@iies.su.se

Peter J. Klenow
Stanford University
Department of Economics
and NBER
Klenow@Stanford.edu

Huiyu Li
Federal Reserve Bank of San Francisco
Huiyu.Li@sf.frb.org

A data appendix is available at http://www.nber.org/data-appendix/w34190

1 Introduction

Economists weigh the static distortions created by price-cost markups against the dynamic incentives that such markups provide for private sector innovation. Classic examples include Dasgupta and Stiglitz (1980) and Tirole (1988). More recently, the debate has evolved to include firm heterogeneity in markups and innovation effort.[1] However, less attention has been paid to the possibility that the *source* of markups may differ across firms—with implications for the optimal allocation of research. Firms may differ in the step size of their quality-improving innovations, as in Klette and Kortum (2004) or Akcigit and Kerr (2018). Such quality innovation can confer significant knowledge spillovers onto other firms. Alternatively, markups may vary because firms differ in their process efficiency (which arguably spills over to other firms less easily) or in their ability to circumvent regulations, obtain permits, or impose entry barriers on potential competitors (which bear no obvious knowledge externality).[2]

In this paper, we analyze the optimal allocation of R&D in an economy with multiple sources of markup heterogeneity across firms, and we contrast it with the market equilibrium under *laissez-faire*.[3] In the model we develop, firms differ both in their quality advantage over other firms and in their degree of proprietary process efficiency. Firms can improve quality through innovating, whereas the

[1]Evidence for heterogeneous markups includes Edmond, Midrigan, and Xu (2015, 2023), Haltiwanger, Kulick, and Syverson (2018), De Loecker, Eeckhout, and Unger (2020), , Baqaee and Farhi (2020), and Autor, Dorn, Katz, Patterson, and Van Reenen (2020). Other papers provide evidence for heterogeneous markups in the form of incomplete exchange rate pass-through. See Gopinath and Rigobon (2008), Gopinath, Itskhoki, and Rigobon (2010), Fitzgerald and Haller (2014), and Amiti, Itskhoki, and Konings (2019). Among reasons for heterogeneous research intensity are firm differences in process efficiency (Cavenaile, Celik, and Tian, 2023; Voronina, 2022), research productivity (De Ridder, 2024), or ability to implement innovations (Akcigit, Celik, and Greenwood, 2016; Ma, 2024).

[2]One paper which documents heterogeneity in the source of markups is Akcigit, Baslandze, and Lotti (2023). Using firm-level evidence from Italy, these authors contrast small firms (whose rents stem from innovation) with market leaders (whose rents rely importantly on political connections). Only the former source of rents generates knowledge spillovers onto future innovators in their environment.

[3]We focus on the allocation of research across firms and fix total research. That is, we set aside the question of whether the market devotes too little labor to research versus production.

level of process efficiency of each firm is given once and for all. Both big quality steps and high process efficiency enable a firm to charge above-average markups, yet only quality steps bestow knowledge spillovers on subsequent innovators.

Our analysis sheds light on whether the allocation of research under *laissez-faire* is excessively or insufficiently tilted towards high-rent firms, depending upon whether the main source of markup heterogeneity across firms is quality steps versus process efficiency. The planner wishes to undo the static misallocation of production labor created by markup dispersion, but also endeavors to allocate research labor optimally. In particular, compared to *laissez faire*, the planner wants to shift innovation effort toward big quality step firms.

We use data on French manufacturing from 2012 to 2019 to infer the extent to which firms differ in their quality step sizes and process efficiency. According to our theory, big quality steps allow firms to raise their prices to charge high markups. High process efficiency, in contrast, allows firms to charge high markups by *not* passing through their lower marginal cost into their prices. Thus, among firms with higher markups, high-price firms take big quality steps and low-price firms enjoy higher process efficiency.

We find that French manufacturing firms do indeed differ not only in the size but also in the source of their price-cost markups (as measured by the ratio of their revenues to their total costs). High markup firms sometimes have high prices (consistent with big step sizes) and sometimes have high process efficiency (high physical output relative to their inputs). Consistent with our model, firm differences in process efficiency are quite persistent and the average growth rate of process efficiency is small.

Compared to the laissez-faire equilibrium, we first consider the case of a planner who is constrained to carry out production at the same firm that develops a given innovation. In that case the planner's solution achieves about 8 basis points faster growth per year (from 1.20 to 1.28% per year) by increasing the share of innovation and products at firms with big step sizes. This translates to about a 2% gain in permanent consumption-equivalent welfare.[4] We find further that, because firms are heterogeneous in the source of their markups and size, a size-dependent R&D subsidy of the kind in place in France can achieve only 1 basis point faster growth than under *laissez faire*.

If the planner can instead decouple innovation and production, then the planner will concentrate research at the big step size firms and production at the high process efficiency firms. Such an unconstrained planner attains a growth rate of 1.51% per year (versus 1.20% in the decentralized equilibrium), which translates to 8.2% higher welfare. This might be achievable in a decentralized equilibrium in which firms can freely license their innovations to other firms for production, in the spirit of the BioNTech and Pfizer collaboration to mass produce COVID vaccines. We leave a full analysis of licensing and M&A (mergers and acquisitions) to future work.

This paper relates to several strands of literature: First, to the literature on competition, R&D, and growth such as Dasgupta and Stiglitz (1980), Aghion, Harris, and Vickers (1997), Aghion, Harris, Howitt, and Vickers (2001), and Aghion, Bloom, Blundell, Griffith, and Howitt (2005). Acemoglu and Akcigit (2012) use a step-by-step innovation model with leaders and followers to analyze the growth and welfare implications of various patent protection policies. They conclude that patents should provide stronger protection to firms with bigger technological leads over their rivals. There is only one source of markups in their model, however, namely quality differences. We contribute to this literature by considering multiple sources of markup heterogeneity at once, and by proposing an empirical method to infer the sources of markup heterogeneity in an economy.

Second, our paper overlaps with the recent literature on whether market power and markup dispersion are inhibiting growth. Examples include Peters

[4] In addition to the dynamic gains, the planner reaps static gains from eliminating the distorted allocation of production labor due to markup dispersion across products. However, the planner's allocation entails lower average process efficiency and higher overhead costs because it tilts activity to the subset of firms with big step sizes.

(2020), Liu, Mian, and Sufi (2022), Akcigit and Ates (2023), De Ridder (2024), and Aghion, Bergeaud, Boppart, Klenow, and Li (2023). Our analysis is distinct in encompassing two sources of markup heterogeneity—process efficiency and quality steps—which allows us to characterize the socially optimal R&D allocation across firms with different markup sources and to contrast it with the *laissez-faire* equilibrium.

Third, two recent papers focus on R&D misallocation across fields and sectors. Acemoglu (2023) analyzes, for example, green versus dirty innovations and health care inventions emphasizing prevention versus *ex post* cures. Liu and Ma (2024) compare the equilibrium and optimal R&D allocation across fields in a multi-sector environment involving dynamic cross-field knowledge spillovers. In contrast, our focus is on R&D misallocation across firms within industries.

Fourth and finally, recent papers analyzing R&D misallocation across firms include Babalievsky (2022), Ayerst (2023), and Lehr (2024, 2025). Babalievsky (2022) and Lehr (2025) model imperfections in the market for inventors. Lehr (2024) interprets dispersion in R&D intensity across firms as reflecting misallocation, and quantifies the extent to which growing R&D misallocation is undercutting U.S. growth. Ayerst (2023) contrasts private profitability (which leads to a high rate of patenting) with spillovers (proxied by patent citations). We focus instead on markup dispersion and its decomposition into quality (with spillovers) versus process efficiency (with no spillovers). We do not emphasize differences in patenting, patent citations, or R&D intensity across firms.

The rest of the paper proceeds as follows: Section 2 lays out our endogenous growth model of multiproduct firms with differing quality innovation step sizes and process efficiency levels. Section 3 describes the facts that motivate and quantify our theoretical channels. In Section 4 we calibrate the key model parameters based on the moments we document in the French data and compare the R&D allocation in the decentralized economy with the planner's solution. Section 5 concludes.

2 Theory

In this section, we provide a theory that features two distinct sources of markup heterogeneity across firms, with spillovers of quality improvements onto future innovators. We first lay out the model environment and then we describe and define the decentralized equilibrium. We contrast the decentralized outcome to the solutions of two versions of a planner's problem. In the first version, which we call the constrained planner's problem, the planner can freely allocate production labor and research labor, respectively, but cannot separate production and innovation. We remove this constraint in the second version.

2.1 Setup

We suppress time indices wherever it should not lead to confusion.

Household and preferences Time is discrete and there is a representative household with the following preferences over a final output good:

$$\sum_{t=0}^{\infty} \beta^t \log(C_t). \tag{1}$$

Furthermore, the household can supply L units of production labor and Z units of R&D labor each period without generating disutility.

Final goods production Final output, Y is a Cobb-Douglas bundle of a unit interval of intermediate goods which carry qualities $q(i)$:

$$Y = \exp\left(\int_0^1 \log\left(q(i)y(i)\right) di\right).$$

The final output good can be used for two purposes: consumption C or, as will be explained below, production overhead O.

Intermediate input production A "large" number J of firms can produce the intermediate goods $i \in [0,1]$. Each firm j can produce at a line-specific quality $q(i,j)$ and a firm-specific process efficiency $a(j)$:

$$y(i,j) = a(j) \cdot l(i,j), \tag{2}$$

where $l(i,j)$ denotes production labor used in line i by firm j. The initial distribution of quality levels across firms in each line is exogenously given.

The quality levels at which a firm could produce change endogenously over time as a result of R&D activity. Each firm has access to a linear R&D technology with heterogeneous step sizes. That is, if $x \cdot \psi_z$ units of research labor are used by a particular firm j, this firm innovates in x lines that are randomly drawn from the lines where firm j is not the highest quality producer yet. In each selected line, firm j raises the highest existing quality across firms—the best blue print so far—by a factor $\gamma(j) > 1$. The innovating firm j can then produce at this higher quality from the next period onward. As quality innovations build on each other, they confer positive spillovers on future innovators.

The J firms differ exogenously in their level of process efficiency $a(j)$ and their innovation step size $\gamma(j)$. In the following, we assume both dimensions take on two potential values (high a_H and low a_L in the case of process efficiency, and big γ_B and small γ_S in the case of step size). We further assume $\gamma_S > a_H/a_L \equiv \Delta$.[5] Given the binary heterogeneity in both γ and a, there are four types of firms indexed by $k \in \{HB, HS, LB, LS\}$. Here HS denotes firms with a high process efficiency and a small innovation step size and so on. In the following, we denote the ratio of big to small step sizes by $\Gamma \equiv \gamma_B/\gamma_S$.

In addition to the linear production cost, firms use resources for fixed production costs called "overhead". That is, in order to be active in a period, a firm that has the highest quality level in $n(j)$ lines needs to spend $\psi_{ok}\, n(j)^2 Y$

[5] This will ensure that the constrained planner finds it optimal to produce using the highest quality firm in each line. This condition also ensures production by the highest quality firm (irrespective of its type) in the decentralized economy we study below.

units of final output on overhead. Importantly, the overhead costs are convex in the number of lines and can depend on the type k of a firm. Allowing ψ_{ok} to vary by firm type allows the model to generate firm size differences that are not perfectly correlated with firm markups.

In the following we assume that all firms of the same type k start out with the same number of lines $n_{k,0}$ in which they have the highest quality. With the convex overhead cost schedule and the continuum of lines this ensures that it is optimal for the planner to maintain homogeneity within type and the planner's problem can be characterized in terms of four "representative" firms.[6] Hereafter we denote by ϕ_k the fraction of the J firms that are of type k.

Aggregates and resource constraints Aggregate resources used on overhead are given by $O = \sum_{j=1}^{J} \psi_{ok}\, n(j)^2 Y$ and the economy's resource constraint is

$$Y = C + O. \tag{3}$$

Next, we have the production labor and R&D labor resource constraints:

$$L = \sum_{j=1}^{J} \int_0^1 l(i,j)\, di, \tag{4}$$

$$Z = \psi_z \sum_{j=1}^{J} x(j) = \psi_z \sum_k J\phi_k x_k, \tag{5}$$

where the second equality exploits homogeneity within types. As the total number of lines is fixed at one, (5) implies that the aggregate rate of creative destruction—the share of lines innovated upon—is fixed and equal to $\sum_{j=1}^{J} x(j) = Z/\psi_z$. Finally, we have an accounting equation that says that the

[6]The same will also hold true in the decentralized equilibrium that we study. Along a balanced growth path, homogeneity within type is automatically fulfilled and this assumption does not put any additional restrictions.

total number of products produced must sum to 1 across firms:

$$\sum_k J\phi_k n_k = \sum_k S_k = 1, \tag{6}$$

where $S_k = \sum_k J\phi_k n_k$ denotes the share of lines produced by firms of type k.

2.2 Decentralized economy

For the decentralized economy, we assume competitive markets with the exception of intermediate goods production. Therefore, final output can be characterized as the behavior of a representative firm solving

$$\max_{y(i)} \left\{ P \cdot \exp\left(\int_0^1 \log\left(q(i)y(i)\right) di \right) - \int_0^1 p(i)y(i)\, di \right\}. \tag{7}$$

In the following, we normalize the price of the final output $P = \exp\left(\int_0^1 \log\left(p(i)/q(i)\right) di \right)$ to 1 in all periods.

The representative household inelastically supplies L units of production labor and Z units of research labor to the labor market and solves

$$\max_{\{C_t, A_{t+1}\}_{t=1}^{\infty}} \sum_{t=0}^{\infty} \beta^t \log(C_t), \tag{8}$$

subject to $A_{t+1} = A_t(1 + r_t) + w_t L + w_{z,t} Z - C_t$ for all periods and a standard no-Ponzi game condition. Here A denotes wealth, r the real interest rate, and w and w_z the wage rates of production and research labor, respectively.

Intermediate input production The J intermediate input producers own patents to produce at particular qualities in given lines. In each line the different firms that own patents to produce at different quality levels compete à la Bertrand. We solve for equilibrium pricing decisions under Bertrand competition, yielding an expression for the per-period profits of a firm. We then focus on the firm's dynamic problem.

As we assumed $\gamma_S > a_H/a_L$, the firm that has the highest quality patent in a line is the "leading" firm, i.e., the firm with the lowest quality-adjusted marginal cost. Due to the Cobb-Douglas structure, under Bertrand competition it is optimal for this leading firm $j(i)$ in a given line i to set its quality-adjusted price equal to the quality-adjusted marginal cost of the second-best producer $j'(i)$, i.e., $p(i, j(i), j'(i))/q(i, j(i)) = w/q(i, j'(i)) \cdot a(j')$. This price setting implies the markup over marginal cost charged by producer $j(i)$ in line i is

$$\mu(i, j(i), j'(i)) = \gamma_j \cdot \frac{a(j(i))}{a(j'(i))}. \tag{9}$$

The markup is equal to the step size (which depends on the identity of the producing firm) times the ratio of process efficiency for the producing firm relative to the second-best firm. Hence, the process efficiency of the second-best firm, which is either high or low, does influence markups and therefore profits from a given line. Due to the Cobb-Douglas structure of final good production, sales in each product line are given by Y and are independent of the quality level and the price. The operating profits of a producing firm in a given line (before overhead costs) are therefore given by $Y(1 - 1/\mu(i, j(i), j'(i)))$.

Total period profits of firm j then depends on the number of lines in which the firm has the highest quality patent, $n(j)$, and the share of these lines in which they face a high productivity second-best firm $h(j)$. These two variables $n(j)$ and $h(j)$ are the two individual state variables in the dynamic firm problem. Total firm profits after overhead expressed relative to output Y are given by

$$\pi\,(j, n, h) = nh\left(1 - \frac{a_H}{\gamma_j a_j}\right) + n(1-h)\left(1 - \frac{a_L}{\gamma_j a_j}\right) - \psi_{oj} n^2. \tag{10}$$

As we again assume homogeneity within types, (10) can be expressed as four type-specific profit functions $\pi_k(n, h)$ for $k \in \{HB, HS, LB, LS\}$ that depend on the individual state variables n and h. The profit functions are type-specific as they depend on a, γ, and ψ_o of the producing firm. The dynamic problem of a firm of type k is then given by

$$V_k = \max_{\{x_t, n_{t+1}, h_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} Y_t \left(\pi_k(n_t, h_t) - x_t \psi_z \frac{w_{z,t}}{Y_t} \right) \prod_{s=0}^{t} \left(\frac{1}{1+r_s} \right)$$

subject to

$$n_{t+1} = n_t(1 - Z/\psi_z) + x_t, \ \forall t, \tag{11}$$

$$h_{t+1} n_{t+1} = h_t n_t (1 - Z/\psi_z) + S_t x_t, \ \forall t, \tag{12}$$

for a given n_0 and h_0 and non-negativity constraints $x_t \geq 0$, $n_{t+1} \geq 0$. That is, each firm is choosing its R&D activity to maximize the net present value of its profits. Constraint (11) describes how the number of lines the firm produces evolves. The firms take the aggregate rate of creative destruction Z/ψ_z as exogenously given. Constraint (12) characterizes dynamics in the share of lines in which the firm faces a high process efficiency second-best firm. The firm takes into account that the newly innovated lines are drawn from the total pool of lines, in which a fraction S_t is currently served by high process efficiency firms.

Market clearing and aggregates We have the labor markets clearing conditions (4) and (5) and the asset market clearing condition

$$\sum_k J \phi_k V_k = A.$$

Finally, we have the accounting equation (6) and an accounting equation giving the aggregate share of lines served by high process efficiency firms

$$S_t \equiv S_{HB,t} + S_{HS,t} = J \left(\phi_{HB} n_{HB,t} + \phi_{HS} n_{HS,t} \right). \tag{13}$$

Equilibrium definition A decentralized equilibrium is defined as a sequence of quantities and prices that jointly solve the final producer problem, the intermediate producer problems, and the household problem, while satisfying market clearing and all the aggregate constraints.

2.3 Planner's problem

Constrained planner First, we consider a planner who can allocate production and R&D labor across firms under the constraint that R&D must be done in-house, i.e., a firm needs to have innovated in a specific line by itself in order to produce at a particular quality level. The planner maximizes welfare by choosing the allocation of research labor Z to the different J firms to innovate and then produce in the innovating firms with the available production labor L. This maximization is subject to the resource constraint which also includes the overhead cost. Online Appendix A states the problem formally.

Unconstrained planner To understand potential gains from licensing or M&A in our model, we also entertain a case where the planner can separate production and innovation. We call this the unconstrained planner's problem where the planner can assign R&D resources to any firm, have that firm randomly draw a line to innovate upon, and then take this blueprint to produce using a potentially different firm. The overhead cost is still incurred for the number of lines a firm produces. Formally, this planner's problem is the same as for the constrained planner's problem but without the constraint that innovations have to be developed in-house by the producing firm.

2.4 Definition of a BGP

We define a balanced growth path (BGP) in the standard way, i.e., as a path along which all quantities grow at constant rates. In the following we focus on an interior BGP in which growth is strictly positive and all firm types are active. Under some parameter restrictions such an interior BGP exists and is unique (both in the planner problems and in the decentralized equilibrium). We will denote variables along the BGP by upper bar, e.g., $\bar{r}$ will denote the interest rate along the decentralized BGP. Along a BGP, the distribution of the number of lines across firms will be stationary while aggregate output, consumption, the geometric mean of quality $Q_t = \exp\left(\int_0^1 \log\left(q_t(i)\right) di\right)$, and total resources used for overhead grow at an endogenous rate $\bar{g}$.

2.5 Welfare along the BGP

As consumption grows along the BGP at a constant endogenous rate $\bar{g}$, welfare in (1) can be rewritten as

$$\frac{1}{1-\beta}\left(\log(\bar{C}_0) + \frac{\beta}{1-\beta}\log(1+\bar{g})\right), \tag{14}$$

where $\bar{C}_0 \equiv C_t(1+\bar{g})^{-t}$ is the detrended consumption level along the BGP. That is, welfare can be written as a weighted sum of the logarithm of the detrended consumption *level* and the steady state rate of consumption *growth*, where the relative weight on the logarithm of the gross growth rate is $\frac{\beta}{1-\beta}$. As output and innovations use distinct types of labor which are in fixed supply, discrepancies in the consumption level $\bar{C}_0$ and the growth rate $\bar{g}$ between a planner's solution and a decentralized equilibrium solely arise from different allocations of production and research labor across firms. In this sense, our model shuts down potential distortions in the total amount of resources devoted to R&D and focuses entirely on misallocation, both statically and in terms of R&D resources.

More precisely, in the decentralized equilibrium as well as the planner's solution, the detrended consumption level along the BGP can be written as the following product

$$\bar{C}_0 = (1-\bar{o}) \cdot Q_0 \cdot \Phi \cdot \mathcal{M} \cdot L, \tag{15}$$

where $\bar{o} \equiv O/Y = \sum_k \psi_{ok} J \phi_k(\bar{n}_k)^2$ is the fraction of output used for overhead, $Q_0 = \exp\left(\int_0^1 \log\left(q_0(i)\right) di\right)$ is the geometric mean of initial quality levels, $\Phi = \exp\left(\sum_k \bar{S}_k \log\left(a_k\right)\right) = a_S \Delta^{\bar{S}}$ is the geometric average of process efficiencies, and $\mathcal{M}$ captures potential misallocation of production labor across lines. In the decentralized equilibrium this allocative efficiency term is the ratio of the geometric relative to the arithmetic average of the inverse markups across lines, or formally

$$\mathcal{M} = \frac{\exp\left(\int_0^1 \log \frac{1}{\mu(i,j(i),j'(i))}\, di\right)}{\int_0^1 \frac{1}{\mu(i,j(i),j'(i))}\, di} \leq 1. \tag{16}$$

$\mathcal{M}$ is below one due to markup dispersion across lines in the decentralized equilibrium, whereas it is one for the constrained and unconstrained planner.

The terms $\bar{o}$ and Φ are functions of the distribution of lines across firm types, $\bar{S}_k$. Overhead resources $\bar{o}$ are minimized if firm size is inversely related to ψ_{ok}, i.e., $\bar{n}_k = 1/(\psi_{ok} \sum_{k'} J\phi_{k'}/\psi_{ok'})$, $\forall k$. In contrast, aggregate process efficiency Φ is maximized if the high process efficient firms serve the whole market ($\bar{S} = 1$). As the decentralized equilibrium and the planners' solutions will give rise to different distributions of market shares along the BGP (and there is markup dispersion in the decentralized equilibrium) the detrended consumption level $\bar{C}_0$ in the decentralized equilibrium will deviate from the planner's solutions.

In both the decentralized equiliibrum and the constrained planner's solution, the BGP growth rate is a function of the market share distribution:

$$1 + \bar{g} = \exp\left(\frac{Z}{\psi_z} \sum_k J\phi_k \bar{n}_k \log\left(\gamma_k\right)\right) = \left(\bar{\gamma}\right)^{\frac{Z}{\psi_z}}, \tag{17}$$

where $\bar{\gamma} \equiv \Pi_k \gamma_k^{\bar{S}_k}$, and $\bar{S}_k \equiv J\phi_k \bar{n}_k$. Here, $\bar{\gamma}$ is the geometric mean of the step size γ_k weighted by the share of lines $\bar{S}_k$ served by type k firms. The growth rate is higher when more products are produced by firms with a big step size. The intuition for this result is as follows: As total research labor is fixed and all firms have the same linear R&D technology, the rate of creative destruction is fixed and equal to $\frac{Z}{\psi_z}$. In the constrained planner solution and the decentralized equilibrium, along a BGP the firm size distribution is stationary and all firms innovate on a number of lines that is proportional to their market share $\bar{S}_k$. The growth rate is then the market share-weighted geometric average of the step size raised to the rate of creative destruction $\frac{Z}{\psi_z}$.

For the unconstrained planner, R&D does not have to be done in-house and it is possible to achieve the maximum gross growth rate $\gamma_B^{\frac{Z}{\psi_z}}$ by having only high step size firms innovate. For the constrained planner and in the decentralized equilibrium, quality improvements have to be developed in-house, which creates an interesting trade-off. The constrained planner weighs level and growth effects highlighted in (14). The market share distribution along the decentralized BGP is determined by the relative profitability (capability to charge markups) across firms, which will in general lead to inefficiencies. We will next characterize the optimal market share distributions chosen by the two planners along the BGP and contrast these with the decentralized equilibrium.

2.6 Characterizing BGPs

2.6.1 Constrained planner's solution

How does the constrained planner determine the market shares $\bar{S}_k = J\phi_k\bar{n}_k$? As shown in Online Appendix A, the optimal differences between $\bar{n}_k$ and $\bar{n}_{k'}$ for any two types of firms that produce must satisfy

$$\tilde{\psi}_{ok}\bar{n}_k - \tilde{\psi}_{ok'}\bar{n}_{k'} = (1-\bar{o})\left(\log\left(\frac{a_k}{a_{k'}}\right) + \left(1 + \frac{Z/\psi_z}{1/\beta - 1}\right)\log\left(\frac{\gamma_k}{\gamma_{k'}}\right)\right), \quad (18)$$

where $\tilde{\psi}_{ok} \equiv 2\psi_{ok}$ and $\bar{o} = O/Y$ is the share of total overhead costs in output along the BGP.[7] The planner chooses a higher market share for firms with higher process efficiency a and bigger step size γ. Increasing the number of products in high a firms increases the aggregate level of process efficiency Φ in production, whereas increasing the number of products in big γ firms has a positive effect on the long-run growth rate $\bar{g}$.[8] As a consequence, the relative weight the planner puts on differences in step sizes versus process efficiency is increasing

[7] The planner may find it optimal to shut down some firm types. We fully characterize the solution in Online Appendix A.

[8] Because the constrained planner cannot separate production and innovation, choosing a high market share for one firm also implies that this firm uses a larger share of the R&D labor.

in the consumer's patience (β) and in the amount of effective research labor available (Z/ψ_z). The differences in the optimal number of products across firms are moderated by the scalar in front of the overhead cost curve. If the overhead cost curve shifts down (ψ_{ok} decreases for all k) the optimal long-run differences in firm size are magnified. The convex overhead cost is the reason why the planner does not allocate all products to firms with the highest combination of step size and process efficiency.[9]

As we noted in the setup of the planner's problem, the planner finds it optimal to allocate the same amount of production labor to each product line. Hence, $\bar{S}_k$ is also the share of production labor allocated to type-k firms ($L_k/L = \bar{S}_k$). Because the rate of innovation in each firm $\bar{x}_k/\bar{n}_k$ is equal to the aggregate rate of creative destruction and all firms have the same R&D efficiency, $\bar{S}_k$ is also the share of R&D labor allocated to type-k firms along the BGP.

2.6.2 Unconstrained planner's solution

In the case of the unconstrained planner, innovations and production can be decoupled and—as we show in the Online Appendix A—the difference in $\bar{n}_k$ of any two types of firms that produce satisfies

$$\tilde{\psi}_{ok}\bar{n}_k - \tilde{\psi}_{ok'}\bar{n}_{k'} = (1-\bar{o})\log\left(\frac{a_k}{a_{k'}}\right). \tag{19}$$

The unconstrained planner may let firms with lower production efficiency produce to save overhead costs. Unlike for the constrained planner, γ_k does not matter for the level of production.

For innovations, the unconstrained planner has only the high step size firms innovate and the growth rate is

$$1+\bar{g} = \gamma_B^{\frac{Z}{\psi_z}}.$$

[9]Since the overhead cost function is quadratic, $\bar{o}$ is a quadratic function of the $\bar{n}$, and we can derive a closed form solution for the optimal share of products produced by type-k firms. This is done in Online Appendix A.

2.6.3 Characterizing the BGP of the decentralized equilibrium

The first-order conditions for the J firms in the decentralized equilibrium imply[10]

$$\tilde{\psi}_{ok}\bar{n}_k - \tilde{\psi}_{ok'}\bar{n}_{k'} = \left(\frac{a_L}{\gamma_{k'}a_{k'}} - \frac{a_L}{\gamma_k a_k}\right)(\bar{S}\Delta + 1 - \bar{S}), \quad \bar{S} \equiv \bar{S}_{HS} + \bar{S}_{HB}. \tag{20}$$

Unlike the constrained planner's condition, firms do not distinguish between step size and process efficiency. We show in Online Appendix A that the number of products produced by a type-k firm along an interior BGP satisfies

$$\bar{n}_k = \frac{\bar{\tilde{\psi}}_o}{\tilde{\psi}_{ok}}\frac{1}{J\phi} + \frac{\bar{S}\Delta + 1 - \bar{S}}{\gamma_S}\frac{\tilde{\omega}_k}{\tilde{\psi}_{ok}} \tag{21}$$

where $\bar{\tilde{\psi}}_o \equiv 1/(\sum_{k'} J\phi_{k'}/\tilde{\psi}_{ok'})$ is the average overhead cost and

$$\tilde{\omega}_k \equiv \left(\sum_{k'}\phi_{k'}\frac{\bar{\tilde{\psi}}_o}{\tilde{\psi}_{ok'}}\frac{\gamma_S a_L}{\gamma_{k'}a_{k'}}\right) - \frac{\gamma_S a_L}{\gamma_k a_k}$$

is the markup advantage of a type-k firm relative to the average firm. Firms with higher markups, either due to bigger step size or higher process efficiency, have higher than average share of products. Let $\phi^H \equiv \phi_{HB} + \phi_{HS}$ be the share of firms with high process efficiency and $\bar{\tilde{\psi}}_o^H$ the average overhead cost for these firms. The product shares in the decentralized equilibrium can be shown to satisfy

$$\bar{S}_k = \frac{\bar{\tilde{\psi}}_o}{\tilde{\psi}_{ok}}\phi_k\left(1 + \frac{\bar{S}(\Delta-1)+1}{\gamma_S}\frac{\tilde{\omega}_k}{\bar{\tilde{\psi}}_o/J}\right), \tag{22}$$

with

$$\frac{\bar{S}(\Delta-1)+1}{\gamma_S} = \frac{(\Delta-1)\frac{\bar{\tilde{\psi}}_o}{\bar{\tilde{\psi}}_o^H}\phi^H + 1}{\gamma_S - \frac{J}{\bar{\tilde{\psi}}_o}(\Delta-1)\frac{\bar{\tilde{\psi}}_o}{\bar{\tilde{\psi}}_o^H}\phi^H(\sum_{k,a_k=a_H}\phi_k^H\frac{\bar{\tilde{\psi}}_o^H}{\tilde{\psi}_{ok}}\tilde{\omega}_k)}.$$

[10]As in the planner's solutions, some firms may find it optimal to shut down in the decentralized equilibrium. Here we provide intuition using the first order conditions of producing firms. We fully characterize the equilibrium in Online Appendix A.

The decentralized product shares in (22) differ from the expression for the constrained planner's product shares in (18). In contrast to the planner's solution, the relative market shares are pinned down by the relative markups and are independent of the weight β consumers put on growth vs. level effects. In addition, as an increase in the minimum step size γ_S shifts—for a given Δ and Γ—the markup distribution to the right, γ_S has an influence on the decentralized market shares along the BGP. Such an effect is not in the planner's solution, where the optimal market shares only depends on relative step size and relative process efficiency Γ and Δ.

Furthermore, the relative labor share of two firm types k and k' along the BGP is given by

$$\frac{\lambda_k(\bar{S})}{\lambda_{k'}(\bar{S})} = \frac{\gamma_{k'} a_{k'}}{\gamma_k a_k}, \tag{23}$$

as all firms face the same share of high process efficiency competitors $\bar{S}$. Since production wages are the same across firms and $\bar{S}_k$ is equal to the sales share of type-k firms, the relative employment of type k firms is given by

$$\frac{L_k}{L_{k'}} = \frac{\lambda_k(\bar{S})\bar{S}_k}{\lambda_{k'}(\bar{S})\bar{S}_{k'}} = \frac{\gamma_{k'} a_{k'}}{\gamma_k a_k}\frac{\bar{S}_k}{\bar{S}_{k'}}. \tag{24}$$

This expression says that firms with high markups (high $\gamma_k a_k$) have lower employment shares than sales and product shares. This differs from the planner's solution, where employment and market shares coincide for all firms.

There is also heterogeneity in labor allocated to the different production lines *within* a firm in the decentralized equilibrium BGP. This is because markups vary across production lines within a firm by a factor Δ depending on whether the second-best firm is of high or low process efficiency. The amount of labor a firm devotes to line i where it faces a high efficiency competitor relative to the amount of labor it devotes to line i' where it faces a low efficiency competitor is therefore

$$l(i,j) = \Delta\, l(i',j). \tag{25}$$

This gives rise to an additional static efficiency loss as the planner would equalize the amount of production labor across all lines within a firm. The heterogeneity in labor allocation across lines lowers the level of detrended consumption relative to the planner's allocation and shows up as $\mathcal{M} < 1$ in our welfare decomposition (15). The allocative efficiency term along the BGP becomes

$$\mathcal{M} = \frac{\Delta^{\bar{S}}}{\Delta \bar{S} + 1 - \bar{S}} \frac{\exp\left(-\sum_k \bar{S}_k \log(a_k \gamma_k)\right)}{\sum_k \bar{S}_k \frac{1}{a_k \gamma_k}} < 1.$$

Finally, as in the constrained planner's problem, the share of research labor allocated to type-k firms is the same as the share of products produced by type-k firms because the arrival rate of new products matches the aggregate rate of creative destruction Z/ψ_z for all firms.

2.7 Planner's solutions vs. the decentralized equilibrium

As detailed above, the unconstrained and constrained planner's allocations differ from the decentralized equilibrium. First, the product share of each type of firm in the decentralized equilibrium will deviate from the planner's product shares, leading to static and/or dynamic misallocation. The dynamic misallocation arises because the market shares of firms determine their innovation activity along a BGP and therefore impacts the long-run growth rate.

Second, conditional on the product shares of each type of firm, the planner and decentralized equilibrium differ in the allocation of employment across product lines. Equations (18) and (24) imply that the constrained planner always wants more employment in the firms with higher $\gamma_k a_k$ than the decentralized equilibrium. This is because the planner internalizes a dynamic spillover of a larger innovation step size (as innovations build on a quality ladder) whereas in the decentralized equilibrium the profitability (markup) of firms is allocative. In addition, because there is a distribution of types of second-best firms, in the decentralized equilibrium, there will be heterogeneity in markups and production labor across lines within firms. In contrast, the constrained and unconstrained planner want to allocate the same amount of labor to each line.

To illustrate why the constrained planner chooses a larger market share for high step size firms, consider the following special case: suppose the product $a_k \cdot \gamma_k$ and overhead costs are the same across all firms. This is the case when $\Delta = \Gamma$ and process efficiency and the step size are perfectly negatively correlated, i.e., $\phi_{HB} = \phi_{LS} = 0$. In this case, as markups are equalized across firms the number of lines per firm is equalized along the decentralized BGP or $\bar{n}^D_{LB} = \bar{n}^D_{HS} = 1/J$. Interestingly, this is the allocation that minimizes overhead cost $\bar{o}$. However, the decentralized equilibrium does not take into account the dynamic positive externality generated by big step size firms. The constrained planner would choose a larger market share of the big step size firms LB and a lower market share for the high process efficiency firms HS. Formally, we have

$$\bar{n}^P_{LB} - \bar{n}^P_{HS} = \frac{1-\bar{o}}{2\psi_o}\frac{Z/\psi_z}{1/\beta - 1}\log(\Gamma) > 0. \tag{26}$$

Hence the constrained planner would sacrifice some static process efficiency and increase the overhead cost $\bar{o}$ to exploit the larger growth potential of the big step size firms. The extent to which this is done depends on the shape of the overhead cost function (ψ_o), the growth potential of the economy determined by the available R&D labor Z/ψ_z and the discount rate β.

3 Data and some facts

In this section we describe the data used to quantify channels of misallocation in the model of the previous section.

3.1 Datasets

Our main source of data is "Fichier Approché des Résultats d'ESANE" (FARE) from 2012 to 2019, which contains the financial statements of all firms that are subject to the standard corporate tax scheme in France.[11] From this data, we construct firm-level measures of nominal sales ($P^Q \cdot Q$), total intermediate input spending ($P^X \cdot X$),the wage bill (W), and net tangible asset value ($P^K \cdot K$).

We focus on the manufacturing sector where product prices are available in the "Enquête Annuelle de Production" (EAP) dataset, which surveys manufacturing firms and covers all firms with more than 20 employees and a stratified sample below this threshold. For each firm, it splits annual sales into product-level sales and units sold, which allows us to construct unit price as product sales divided by units sold.[12]

We merge FARE with the matched employer-employee data "Déclarations Annuelles des Données Sociales" (DADS) to calculate additional moments used for calibration. Our cleaned dataset, with information on unit price by firm product, firm revenue, and inputs, includes 13,942 firms with an average of 2.4 product-level observations per year. More details about the dataset, cleaning procedures, and variable construction are given in Online Appendix B.

3.2 Estimating unit price and productivity

We construct a firm-level price estimate $\hat{p}(j,t)$ in the same way as De Ridder, Grassi, and Morzenti (2025). We first standardize unit prices within each product category by calculating the sales share of each firm that sold the product in year t. Then we calculate the sales-weighted geometric average of prices in the category. We divide unit prices by this average price for standardization. The price level of a firm $\hat{p}(j,t)$ is the geometric average of the standardized unit prices of the firm across its products, where we weight each product by its share of firm sales.

[11] We use the "consolidated" version of FARE which aggregates legal units (or "siren") into enterprise groups ("entreprises profilées"). The scope of this dataset was expanded in 2012 to provide better coverage of large firm balance sheets, given the largest groups are often made up of numerous legal entities, each of them corresponding to an entry in the unconsolidated version of FARE. For this reason, we restrict our sample to 2012–2019.

[12] Our product groups correspond to the classification used by the PRODCOM survey conducted by Eurostat. There are about 4,000 different product codes.

We define a firm's revenue productivity (TFPR) as its revenue/inputs

$$TFPR_{jt} = \frac{P^Q_{jt} \cdot Q_{jt}}{(P^K_{jt} K_{jt})^{\alpha_{s,t}\theta_{s,t}} (W_{jt})^{(1-\alpha_{s,t})\theta_{s,t}} (P^X_{jt} X_{jt})^{(1-\theta_{s,t})}}, \tag{27}$$

where $\alpha_{s,t}$ and $\theta_{s,t}$ are sector-year factor shares.[13] We define physical productivity (TFPQ) as TFPR divided by the standardized firm price

$$TFPQ_{jt} = \frac{TFPR_{jt}}{\hat{p}(j,t)}.$$

3.3 Empirical patterns

First, we find substantial heterogeneity in firm price and TFPQ. Some of the dispersion could be due to measurement error or temporary fluctuations outside of our model. To address this, we regress log price and log TFPQ in each year on their lagged values from 5 years ago and industry-year fixed effects. We then calculate the interquartile range of the projected value (75th/25th percentile ratios in levels). Row 1 of Table 1 shows that the interquartile range of fitted TFPQ is 1.501 while row 2 shows the IQ range of fitted price is 1.427.[14]

Variation in price and TFPQ are only weakly related to firm size. Rows 3-5 of Table 1 report the market share of firms with above average firm prices, TFPQ, and TFPR, respectively. Rows 6-9 report the market share of firms by four price and TFPQ partitions that mimic the four types in our model.[15] Firms with above-average prices are 53% of sales, for example. According to rows 6-9, only

[13] For each sector s and year t and using R=15%,

$$\alpha_{s,t} = \frac{\sum_{j\in s} R \cdot (P^K_{jt} \cdot K_{jt})}{\sum_{j\in s} \left(R \cdot (P^K_{jt} \cdot K_{jt}) + W_{jt} \right)} \quad \text{and} \quad \theta_{s,t} = \frac{\sum_{j\in s} \left(R \cdot (P^K_{jt} \cdot K_{jt}) + W_{jt} \right)}{\sum_{j\in s} \left(R \cdot (P^K_{jt} \cdot K_{jt}) + W_{jt} + P^X_{jt} \cdot X_{jt} \right)}.$$

This formula assumes a common α for firms in the same industry. In the data, we find that the capital to wage bill ratio does not vary significantly with firm hours.

[14] Online Appendix B explains in detail how we computed all the moments in Table 1.

[15] To calculate these shares, we compute the sales or firm share within each sector-year and then average across sector-years giving equal weights to each sector-year.

Table 1: Empirical moments

	Moment	Value
1.	Interquartile ratio of fitted TFPQ	1.501
2.	Interquartile ratio of fitted price	1.427
3.	Market share of firms above the average unit price	0.532
4.	Market share of firms above the average TFPQ	0.440
5.	Market share of firms above the average TFPR	0.398
6.	Share of firms above mean price / above mean TFPQ	0.062
7.	Share of firms below mean price / above mean TFPQ	0.405
8.	Share of firms above mean price / below mean TFPQ	0.466
9.	Share of firms below mean price / below mean TFPQ	0.066

Notes: Details on the construction and of these moments are given in Online Appendix B

6.2%, of firms have both above-average prices and above-average TFPQ. Similarly, only 6.6% of firms have below-average prices and below-average TFPQ. For the majority of firms, price and TFPQ are negatively correlated.

Our model assumes growth comes from quality innovation while process efficiency is fixed for each firm. This simplification is motivated by another important empirical pattern: within-firm TFPQ exhibits little trend. Annual TFP growth averages 2.3% in French manufacturing over 1995–2019 (EU-KLEMS). This might stem in part from process efficiency growth. However, we find that within-firm TFPQ growth averages less than 0.1% per year in French manufacturing from 2012–2019 (see Online Appendix B). Thus it appears that little growth comes from process innovation. As growth from new varieties is seldom captured in measured TFP growth (see for example Moulton, 2018 and Aghion, Bergeaud, Boppart, Klenow, and Li, 2019), this leaves quality growth as the primary driver of measured TFP growth. This suggests that the priority is to model firm-level quality improvements.

As mentioned above, we find that TFPQ is heterogeneous across firms. So firms *do* seem to differ in their process efficiency. To see whether such differences are transitory or persistent, we ranked firms by their TFPQ and looked at the change in the rankings five years later. Figure 1 displays the transition matrix where the vertical axis is the firm's decile of TFPQ in year t and the horizontal axis is the firm's decile in year $t-5$. The color on the grid represents the share of firms in the final state given in the y-axis among all firms that started in the same initial state shown in the x-axis. Firms concentrate along the diagonal, tending to stay in the same decile of TFPQ over a 5-year window from 2014 to 2019. Combined with the aforementioned evidence of low average growth in firm TFPQ, our data suggest that firm TFPQ is highly persistent. This motivates us to model firm process efficiency as given once and for all.[16]

Figure 1: Firm TFPQ transition

(a) Heatmap (b) Probabilities

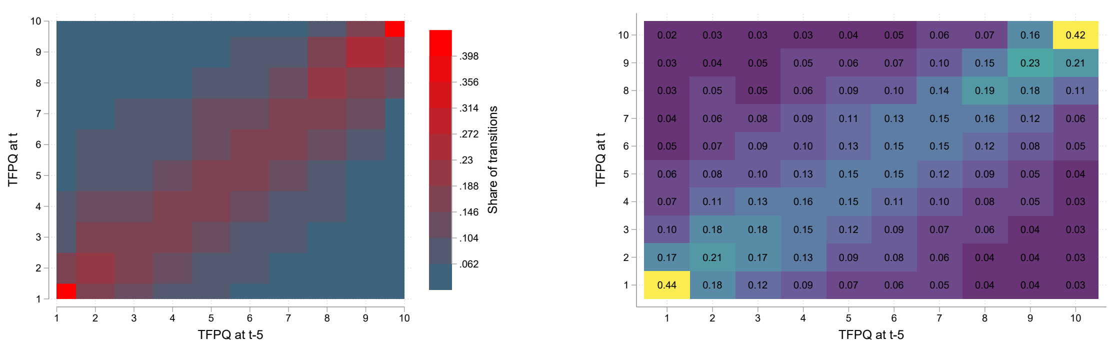

Notes: Underlying data is French manufacturing firm-year observations from 2012–2019. TFPQ is constructed from the FARE dataset. Observations are assigned a decile based on the level of TFPQ within industry-year. Probabilities are calculated based on the decile at $t-5$.

Finally, we document two other empirical patterns to use as targets in the calibration (See Online Appendix B for details). First, overhead cost as a share of

[16]Figure B.1 in Online Appendix B shows the transition matrix for 2012 to 2019.

sales is only weakly correlated with firm size.[17] A 1% increase in sales raises the overhead-to-sales ratio by only 0.001 percentage points, starting from a baseline average of 2.9%. This near-zero semi-elasticity is informative about how the overhead cost shifter correlates with step size and process efficiency across firms. Second, the relationship between R&D expenditures and firm size is ambiguous. As shown in Figure B.2, the relationship between R&D intensity and size depends on whether we include firms that report zero R&D. With these firms included, the relationship appears negative; excluding them, it turns positive. Because this pattern is sensitive to the reporting threshold, we do not use the semi-elasticity of R&D with respect to size as a target moment. Instead, we target the aggregate R&D share of output in the manufacturing sector.

4 Calibration

4.1 Decentralized equilibrium with R&D tax policies

Before calibrating, we extend the model to allow for R&D subsidies so that we can incorporate the effect of these programs on targeted data moments. We will later consider alternative R&D subsidy schedules in our counterfactuals.

In France, businesses face a flat income tax rate $\tau = 33\%$. They receive a subsidy of $\underline{\tau}_{RD} = 30\%$ for R&D expenditure up to 100 million Euros and a subsidy of $\overline{\tau}_{RD} = 5\%$ for R&D expenditure in excess of the cutoff. Incorporating these elements, firm post-tax profits relative to Y become

$$\begin{aligned}
& (1-\tau)\left(\pi_k(n,h) - \frac{\psi_z w_z}{Y}x\right) && \text{post income tax} \\
+ \quad & \underline{\tau}_{RD} \min\left\{\frac{\psi_z w_z}{Y}x,\, \underline{RD}\right\} && \text{subsidy for R\&D below threshold} \\
+ \quad & \overline{\tau}_{RD} \max\left\{\frac{\psi_z w_z}{Y}x - \underline{RD},\, 0\right\} && \text{subsidy for R\&D in excess of threshold}
\end{aligned}$$

[17] We measure overhead costs as payment to employees in occupations not directly associated with production, management, or R&D.

where $\underline{RD}$ is the ratio of the 100 million Euros relative to average firm sales. We assume any net tax revenue or deficit is rebated lump sum to the household.[18]

The R&D subsidy adds an additional determinant of firm size to our model. For a firm that draws x lines, the marginal subsidy rate can be expressed as

$$\tau_{RD}(x) = \begin{cases} \underline{\tau}_{RD} & \text{if } \frac{\psi_z w_z}{Y} x < \underline{RD} \\ \overline{\tau}_{RD} & \text{if } \frac{\psi_z w_z}{Y} x > \underline{RD} \end{cases} . \tag{28}$$

The first-order conditions for the firms become

$$\begin{aligned} \tilde{\psi}_{ok}\bar{n}_k - \tilde{\psi}_{ok'}\bar{n}_{k'} &= \overbrace{\frac{\bar{S}\Delta + 1 - \bar{S}}{\gamma_S}\left(\frac{\gamma_S a_L}{\gamma_{k'} a_{k'}} - \frac{\gamma_S a_L}{\gamma_k a_k}\right)}^{\text{markup difference}} \\ &+ \underbrace{\frac{\psi_z w_z}{Y}\left(\frac{1}{\beta} - 1 + \frac{Z}{\psi_z}\right)\frac{\tau_{RD}(Z/\psi_z \bar{n}_k) - \tau_{RD}(Z/\psi_z \bar{n}_{k'})}{1-\tau}}_{\text{R\&D subsidy}}. \end{aligned}$$

If the marginal R&D subsidy is the same for all firms, the second line above disappears and the subsidy program does not affect the firm size distribution. However, when $\overline{\tau}_{RD}(x) < \underline{\tau}_{RD}(x)$ at the equilibrium x, the marginal subsidy is lower for firms with high enough $\bar{n}_k$ and the subsidy schedule makes larger firms smaller compared to an economy with a flat subsidy.

4.2 Intuition for calibration

We calibrate 12 parameters—$Z/\psi_z, \beta, \gamma_S, \gamma_B, \Delta, \{\psi_{ok}\}, \{\phi_k\}$—along the BGP of the decentralized equilibrium with the R&D policy.[19] We choose parameters jointly to fit (i) the interquartile range of price and productivity (TFPQ) across firms; (ii) market shares of firms with above-average price, TFPQ, and TFPR,

[18] Aghion, Antonin, and Bunel (2021, chapter 12) describe French R&D subsidies. See Online Appendix A for the full model with R&D subsidies.

[19] Mimicking the French R&D policy, we set $\tau = 33\%$, $\underline{\tau}_{RD} = 30\%$, $\bar{\tau}_{RD} = 5\%$. We set $\underline{RD} = 11.96$ since the threshold of 100 million Euros is 11.96 times bigger than the average of firm sales.

respectively; (iii) the share of firms in price and TPFQ bins; (iv) the semi-elasticity of the overhead cost share with respect to firm sales; and (v) the aggregate productivity growth rate, interest rate, and R&D share of output.

We will first provide some intuition for how certain data moments are informative of particular model parameters.[20] Then, we report the calibration results and quantitatively compare allocations under the decentralized equilibrium with the planner's solution.

Recall that, in the decentralized equilibrium, the price of good i is

$$p(i, j(i), j'(i)) = \frac{w \cdot \gamma(j(i))}{a(j'(i))},$$

where $j(i)$ and $j'(i)$ index the producing and the second-best firm. We calculate a firm j's price index as the sales-weighted average of the prices of all products produced by the firm. Firms innovate upon a randomly drawn line. Hence, along the BGP, for any firm the share of products where the second-best producer is of type k is equal to the share of products produced by type k firms ($\bar{S}_k$). Thus the "average" price level for a firm j is proportional to its step size:

$$p_{j,t} = \bar{w}_0 (1 + \bar{g})^t \cdot \gamma_j \sum_{k'} \frac{\bar{S}_{k'}}{a_{k'}} \propto \gamma_j. \tag{29}$$

Equation (29) says that cross-firm price variation is entirely driven by cross-firm differences in the step size of innovations γ_j. The heterogeneity in process efficiency does not explain any of price level differences across firms because, under Bertrand competition, productivity differences affect markups but leave prices unaffected. This means the ratio of the q_3 to q_1 percentile (75th/25th) of the price distribution $p_{j,t}$ is a function of $\Gamma \equiv \gamma_B / \gamma_S$. and firm type shares:

$$1 + (\Gamma - 1) \mathbb{1}(q_1 < \phi_{HB} + \phi_{LB} < q_3), \tag{30}$$

[20]The parameters are calibrated jointly and the model is nonlinear in parameters. Hence the effect of a parameter on outcomes, in general, depends on the value of other parameters.

For the calibration, we will use the ratio of the 75^{th} to 25^{th} percentiles, which will shed light on the value of Γ given the share of firms with big step sizes $\phi_{HB} + \phi_{LB}$.

Next, we define TFPR as firm-level revenue over firm-level costs and TFPQ as TFPR divided by the firm-level price index we described earlier. TFPR is proportional to firm markup, which is the inverse of firm-level labor share (wage bill divided by revenue). Recall that the labor share of a firm j along the BGP is given by

$$\lambda_j = \frac{\bar{S}a_H + (1 - \bar{S})a_L}{\gamma_j a_j} \tag{31}$$

where $\bar{S} = \bar{S}_{HB} + \bar{S}_{HS}$ is the sales-share of high process efficiency firms. Hence, cross-firm differences in TFPR are driven both by differences in the step sizes of innovation and by process efficiency heterogeneity:

$$\text{TFPR}_j \propto \gamma_j a_j. \tag{32}$$

As a consequence, the TFPQ of a firm—defined as its TFPR divided by its price level—is proportional to its process efficiency:

$$\text{TFPQ}_j \equiv \frac{\text{TFPR}_j}{p_j} \propto a_j. \tag{33}$$

In turn, the ratio of the q_3 percentile to the q_1 percentile (75th/25th) of the TFPQ distribution across firms in the model is

$$1 + (\Delta - 1)\mathbb{1}(q_1 < \phi_{HB} + \phi_{HS} < q_3), \tag{34}$$

and is informative of the gap in process efficiency $\Delta \equiv a_H/a_L$.

In addition, the market share of firms with above-average prices is equal to $S_{HB} + S_{LB}$ in the model, and the market share of firms with above-average TFPQ will be equal to $\bar{S}$. Finally, the share of firms below/above mean price and below/above mean TFPQ directly relates to the value of $\phi_{HB}, \phi_{HS}, \phi_{LB}$ and ϕ_{LS}.

4.3 Baseline calibration

Table 2 displays the 12 data moments we use to calibrate the parameters of the model (the 13th moment is redundant because firm shares sum to 1). The firm-level moments correspond to those in Table 1 of the previous section. We obtain the remaining target values from external sources. We take the average annual growth rate of manufacturing productivity and the R&D expenditure share of output from EU-KLEMS over 1995–2019. We use the real interest rate for the U.S. from 1996 to 2016 as estimated by Farhi and Gourio (2018).

Table 3 displays the 12 calibrated parameter values. The calibrated level of γ_S is 1.527 while the step size gap is 1.427 (so $\gamma_B = 2.179$). The model chooses $\Delta = 1.501$ to fit the interquartile range of TFPQ. The calibrated overhead cost is the highest for the firms with big step sizes and high process efficiency to explain why size is not strongly positively correlated with TFPR in the data. We need the discount factor β to be 0.97 to fit the interest rate target and need R&D costs ψ_z/Z to be 16.8 to fit the productivity growth rate (given step sizes).

4.4 Welfare decomposition

We next evaluate the welfare criterion (14) along the BGP at the calibrated parameter values, to compare welfare in the decentralized equilibrium to welfare in the constrained and unconstrained social planner's solutions. Tables 4 and 5 display the components of (14) compared to the constrained and unconstrained planner solutions, respectively. Table 6 compares the allocation of products by firm types. The first rows of Tables 4 and 5 display the distance of the growth rate in the decentralized economy to that of the planner's allocation in consumption-equivalent terms. Namely, we calculate the percent increase ξ in the consumption level $\bar{C}_0$ in the decentralized equilibrium such that welfare is the same as the planner's allocation:

$$\log(1+\xi) = \frac{\beta}{1-\beta} \log\left(\frac{1+\bar{g}^P}{1+\bar{g}^D}\right).$$

Table 2: Baseline calibration

	Targets	Data	Model
1.	Interquartile ratio of fitted TFPQ	1.501	1.501
2.	Interquartile ratio of fitted price	1.427	1.427
3.	Market share of firms above mean price	0.532	0.549
4.	Market share of firms above mean TFPQ	0.440	0.443
5.	Market share of firms above mean TFPR	0.398	0.443
6.	Share of firms above mean price and above mean TFPQ	0.062	0.062
7.	Share of firms below mean price and above mean TFPQ	0.405	0.405
8.	Share of firms above mean price and below mean TFPQ	0.466	0.466
9.	Share of firms below mean price and below mean TFPQ	0.066	0.067
10.	Semi-elasticity of overhead cost share wrt sales	0.001	0.001
11.	Aggregate R&D share of output (percent)	5.5	5.5
12.	Productivity growth rate (ppt year)	1.2	1.2
13.	Interest rate (ppt/year)	5.2	5.2

Notes: Sources for 1 to 11: authors' calculations from DADS, EAP and FARE, French manufacturing, 2012–2019 (see Table 1 and Online Appendix B). Source for 12: EU-KLEMS, French manufacturing, TFP growth in labor-augmenting form, 1995–2019. Source for 13: Farhi and Gourio (2018), U.S. all economy, 1996–2016.

The constrained planner is able to achieve an 8 basis points higher growth rate (1.28%) than in the decentralized equilibrium (1.20%) by allocating a higher share of products (66.3%) to big step size firms than in the decentralized equilibrium (54.9%), as shown in Table 6. This raises welfare by 2.01% in consumption-equivalent terms. The *unconstrained* planner is able to raise the growth rate to 1.51%, or 31 basis points higher than the decentralized equilibrium, because the unconstrained planner can allocate all R&D to the highest step size firms without incurring higher overhead costs. This brings about 8.2% higher welfare in consumption-equivalent terms.

The second rows of Tables 4 and 5 show the welfare gain from changes in the

Table 3: Baseline calibrated parameters

Parameters	Parameter definitions	Values
$\psi_{o\,HB}$	Overhead cost of firms with high process efficiency and big step size	0.38
$\psi_{o\,HS}$	Overhead cost of firms with high process efficiency and small step size	0.19
$\psi_{o\,LB}$	Overhead cost of firms with low process efficiency and big step size	0.16
$\psi_{o\,LS}$	Overhead cost of firms with low process efficiency and small step size	0.07
ϕ_{HB}	Share of firms with high process efficiency and big step size	0.062
ϕ_{HS}	Share of firms with high process efficiency and small step size	0.405
ϕ_{LB}	Share of firms with low process efficiency and big step size	0.466
ϕ_{LS}	Share of firms with low process efficiency and small step size	0.067
γ_S	Small step size	1.527
$\Gamma \equiv \gamma_B/\gamma_S$	Step size gap	1.427
$\Delta \equiv \varphi_H/\varphi_L$	Process efficiency gap	1.501
β	Discount factor	0.962
ψ_Z/Z	R&D cost relative to R&D labor	51.8

level of consumption ξ given by

$$\log(1+\xi) = \log\left(\frac{1+\bar{C}^P}{1+\bar{C}^D}\right).$$

We need to take a stand on $\bar{Q}_0$ to calculate the overall welfare gain. To be conservative, we assume $\bar{Q}_0$ is the same. The constrained planner raises consumption by removing markup dispersion across product lines to improve allocative efficiency.

To achieve higher growth, the planner shifts the share of products away from high process efficiency firms (39.6% vs. 44.3% as shown in Table 6, and incurs bigger overhead costs by giving more products to high step size firms (66.3% vs. 54.9%). As a result, the consumption level under the constrained planner is lower than in the decentralized equilibrium, which incurs a welfare cost of 1.96% in consumption-equivalent units. In contrast, the unconstrained planner can separate production from innovation and is able to shift more production to high process efficiency firms (60.6% vs. 44.3% in Table 6) and increase welfare by 6.76% through a higher detrended consumption level.

Table 4: Welfare comparison, decentralized vs. constrained planner

Growth gain in consumption-equivalent terms	2.01%
Level gain in consumption-equivalent terms	-1.96%
Difference in the growth rate $\bar{g}^P - \bar{g}^D$	0.080 ppt
Relative consumption level $\bar{C}_0^P / \bar{C}_0^D$	0.980
Relative consumption share $(1 - \bar{o}^P)/(1 - \bar{o}^D)$	0.972
Relative process efficiency $\bar{\Phi}^P / \bar{\Phi}^D$	0.981
Relative allocative efficiency $\bar{\mathcal{M}}^P / \bar{\mathcal{M}}^D$	1.028

Table 5: Welfare comparison, decentralized vs. unconstrained planner

Growth gain in consumption- equivalent terms	8.15%
Level gain in consumption-equivalent terms	6.76%
Difference in the growth rate $\bar{g}^P - \bar{g}^D$	0.314 ppt
Relative consumption level $\bar{C}_0^P / \bar{C}_0^D$	1.068
Relative consumption share $(1 - \bar{o}^P)/(1 - \bar{o}^D)$	0.972
Relative process efficiency $\bar{\Phi}^P / \bar{\Phi}^D$	1.069
Relative allocative efficiency $\bar{\mathcal{M}}^P / \bar{\mathcal{M}}^D$	1.028

Table 6: Product share (in percent) and product per firm (relative to the average), planner vs. decentralized

	S_{HB}	S_{HS}	S_{LB}	S_{LS}	high proc. eff. share	big step size share
Decentralized	4.4	39.9	50.5	5.2	44.3	54.9
Constrained Planner	6.0	33.6	60.3	0.0	39.6	66.3
Unconstrained Planner	4.3	56.3	29.6	9.7	60.6	33.9

	n_{HB}	n_{HS}	n_{LB}	n_{LS}	high proc. eff. avg n	big step size avg n
Decentralized	0.70	0.99	1.08	0.77	0.84	1.18
Constrained Planner	0.97	0.83	1.29	0.00	0.75	1.42
Unconstrained Planner	0.70	1.39	0.64	1.45	1.15	0.73

Notes: The first panel displays the share of products produced by firms of each type. "high proc. eff. share" = $\bar{S}_{HB} + \bar{S}_{HS}$ and "big step size share" = $\bar{S}_{HB} + \bar{S}_{LB}$. The $\bar{n}_k$'s are the number of products produced by each firm type relative to the mean $1/J$.

4.5 Alternative R&D policy

The current French R&D policy features one R&D expenditure cutoff and two marginal R&D subsidy rates with a lower rate above the cutoff. We can use our framework to ask whether alternative policies can deliver a better allocation of R&D. We consider policies with the same structure as the current policy, namely one R&D expenditure cutoff and two subsidy rates.

We set the rate below the cutoff to 0.3, the same rate as in the current policy. We then look for a rate above the cutoff $\bar{\tau}$ that maximizes the growth rate. We impose an upper bound of one minus the business income tax on $\bar{\tau}$ so that firms are not being paid to do R&D. We obtain a growth-maximizing subsidy of $\bar{\tau} = 0.37$ for R&D spending relative to average sales above 0.059. This alternative $\bar{\tau}$ is higher than the $\bar{\tau}$ in the current policy, while the cutoff is much lower. Under our alternative policy, only the largest firms—which have low process efficiency, high step size, and low overhead costs—are above the cutoff and enjoy a higher

R&D subsidy. As a result, the largest firms expand while the rest of the firms shrink under our alternative R&D policy.

Table 7 compares the allocation under this alternative policy to both the baseline calibrated economy under the current R&D policy and to the two planner allocations. Compared to the current R&D policy, this alternative policy directs more R&D to firms with bigger step sizes and low process efficiency and raises the growth rate by 1 basis point. The gain is small because firm size is not highly correlated with step size. The alternative policy expands the market share of big step size and low process efficiency firms but shrinks the market share of big step size and high process efficiency firms. The policy is thus less effective at raising the growth rate than the two planner solutions, which can directly target high step size firms. This exercise illustrates that size-based R&D policy may not be effective at raising the growth rate when firm size is not a good proxy for its generation of knowledge spillovers.

Table 7: Product share and growth under an alternative R&D policy

	S_{HB}	S_{HS}	S_{LB}	S_{LS}	Growth rate
Baseline	4.4	39.9	50.5	5.2	1.20
Alternative	4.3	39.0	52.0	4.8	1.21
Constrained planner	6.0	33.6	60.3	0.0	1.28
Unconstrained planner	4.3	56.3	29.6	9.7	1.51

Notes: product shares are in percent, growth in percentage points per year.

5 Conclusion

We formulated an endogenous growth model with two sources of markup dispersion across products and firms: heterogeneous step sizes and heterogeneous process efficiency, respectively. We contrasted the social planner's solution with the decentralized equilibrium. We did this for both a constrained planner (who must pair innovation with production) and an unconstrained planner (who can separate innovation from production).

We mapped out how to infer the sources of markup dispersion using data on prices and physical productivity (TFPQ) across firms. We then implemented this mapping with data on French manufacturing firms from 2012 to 2019. Viewed through the prism of our model, the data patterns suggested ample dispersion in both step sizes and process efficiency.

We illustrated potential magnitudes for R&D misallocation, static misallocation of production labor due to markup dispersion, and the resulting static and dynamic gains from addressing these. We emphasized that size-based R&D subsidies, such as those seen in France, are a poor approximation to the social planner's solution (constrained or unconstrained). A corollary is that research subsidies should not favor all large or small firms in the context of the French economy—only those with big step sizes.

Future work could usefully examine the role of licensing and M&A activity in boosting the level of process efficiency and the growth rate of product quality. Another promising route would be to endogenize total R&D labor. We conjecture this will increase the gains to tilting R&D subsidies toward high step-size firms.

An open question is whether the empirical patterns we find in France extend to other economies. Unfortunately, the coverage of industries and firms with quantities and prices in the U.S. is notably narrower (and shrinking precipitously) relative to France. Quantity and price data are not available for the computer and motor vehicle industries in the U.S., for example. Such quantity and price data are essential for disciplining the importance of process efficiency vs. quality improvements in our model. Other economies, both advanced and developing, may prove to have more suitable data.

Another direction for future research would be to investigate optimal R&D subsidies across industries. Industries differ persistently in their productivity growth rates. Depending on whether such differences stem from arrival rates or step sizes, fast growth industries might warrant bigger subsidies than slow growth industries. This is the spirit of Liu and Ma (2024), who focus on heterogeneity of cross-industry spillovers rather than within-industry spillovers.

There are additional sources of firm heterogeneity that we did not model that can also affect R&D misallocation. This includes firm differences in research efficiency as in Luttmer (2011), about which the planner may have less information as in Akcigit, Hanley, and Stantcheva (2022). This could generate size differences that are unrelated to markup differences.

Finally, our paper features a representative consumer, but could be extended to feature heterogeneity in firm ownership. In this way our theory and empirics could be extended to connect to a growing literature on income and wealth inequality such as Boar and Midrigan (2024), Aghion, Akcigit, Bergeaud, Blundell, and Hémous (2019), Song, Price, Guvenen, Bloom, and Von Wachter (2019), Piketty (2018) and Piketty and Saez (2003). We leave this and other extensions of our analysis for future research.

References

Acemoglu, D. (2023). Distorted innovation: Does the market get the direction of technology right? In *AEA Papers and Proceedings*, Volume 113, pp. 1–28.

Acemoglu, D. and U. Akcigit (2012). Intellectual property rights policy, competition and innovation. *Journal of the European Economic Association 10*(1), 1–42.

Aghion, P., U. Akcigit, A. Bergeaud, R. Blundell, and D. Hémous (2019). Innovation and top income inequality. *Review of Economic Studies 86*(1), 1–45.

Aghion, P., C. Antonin, and S. Bunel (2021). *The power of creative destruction*. Harvard University Press.

Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li (2019). Missing growth from creative destruction. *American Economic Review 109*(8), 2795–2822.

Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li (2023). A theory of falling growth and rising rents. *Review of Economic Studies 90(6)*, 2675–2702.

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005). Competition and innovation: An inverted-u relationship. *Quarterly Journal of Economics 120*(2), 701–728.

Aghion, P., C. Harris, P. Howitt, and J. Vickers (2001). Competition, imitation and growth with step-by-step innovation. *Review of Economic Studies 68*(3), 467–492.

Aghion, P., C. Harris, and J. Vickers (1997). Competition and growth with step-by-step innovation: An example. *European Economic Review 41*(3-5), 771–782.

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

Akcigit, U., S. Baslandze, and F. Lotti (2023). Connecting to power: political connections, innovation, and firm dynamics. *Econometrica 91*(2), 529–564.

Akcigit, U., M. A. Celik, and J. Greenwood (2016). Buy, keep, or sell: Economic growth and the market for ideas. *Econometrica 84*(3), 943–984.

Akcigit, U., D. Hanley, and S. Stantcheva (2022). Optimal taxation and r&d policies. *Econometrica 90*(2), 645–684.

Akcigit, U. and W. R. Kerr (2018). Growth through heterogeneous innovations. *Journal of Political Economy 126*(4), 1374–1443.

Amiti, M., O. Itskhoki, and J. Konings (2019). International shocks, variable markups, and domestic prices. *Review of Economic Studies 86*(6), 2356–2402.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020). The fall of the labor share and the rise of superstar firms. *Quarterly Journal of Economics 135*(2), 645–709.

Ayerst, S. (2023). Innovator heterogeneity, r&d misallocation and the productivity growth slowdown. working paper.

Babalievsky, F. (2022). Misallocation in the market for inventors. working paper.

Baqaee, D. R. and E. Farhi (2020). Productivity and misallocation in general equilibrium. *Quarterly Journal of Economics 135*(1), 105–163.

Boar, C. and V. Midrigan (2024). Markups and inequality. *Review of Economic Studies*, rdae103.

Cavenaile, L., M. A. Celik, and X. Tian (2023). Are markups too high? competition, strategic innovation, and industry dynamics. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3459775.

Dasgupta, P. and J. Stiglitz (1980). Industrial structure and the nature of innovative activity. *Economic Journal 90*(35x8), 266–293.

De Loecker, J., J. Eeckhout, and G. Unger (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics 135*(2), 561–644.

De Ridder, M. (2024). Market power and innovation in the intangible economy. *American Economic Review 114*(1), 199–251.

De Ridder, M., B. Grassi, and G. Morzenti (2025). The hitchhiker's guide to markup estimation. *Conditionally accepted at the Econometrica*.

Edmond, C., V. Midrigan, and D. Y. Xu (2015). Competition, markups, and the gains from international trade. *American Economic Review 105*(10), 3183–3221.

Edmond, C., V. Midrigan, and D. Y. Xu (2023). How costly are markups? *Journal of Political Economy 131*(7), 1619–1675.

Farhi, E. and F. Gourio (2018). Accounting for macro-finance trends: Market power, intangibles, and risk premia. *Brookings Papers on Economic Activity ""*(147), 147.

Fitzgerald, D. and S. Haller (2014). Pricing-to-market: evidence from plant-level prices. *Review of Economic Studies 81*(2), 761–786.

Gopinath, G., O. Itskhoki, and R. Rigobon (2010). Currency choice and exchange rate pass-through. *American Economic Review 100*(1), 304–36.

Gopinath, G. and R. Rigobon (2008). Sticky borders. *Quarterly Journal of Economics 123*(2), 531–575.

Haltiwanger, J., R. Kulick, and C. Syverson (2018). Misallocation measures: The distortion that ate the residual. Working Paper w24199, National Bureau of Economic Research.

Klette, T. J. and S. S. Kortum (2004). Innovating Firms and Aggregate Innovation. *Journal of Political Economy 112*(5), 986–1018.

Lehr, N. H. (2024). Did r&d misallocation contribute to slower growth.

Lehr, N. H. (2025). Does monopsony matter for innovation? Technical report, Working Paper.

Liu, E. and S. Ma (2024). Innovation networks and r&d allocation. Working Paper w29607, National Bureau of Economic Research.

Liu, E., A. Mian, and A. Sufi (2022). Low interest rates, market power, and productivity growth. *Econometrica 90*(1), 193–221.

Luttmer, E. G. (2011). On the mechanics of firm growth. *Review of Economic Studies 78*(3), 1042–1068.

Ma, Y. (2024). Specialization in a knowledge economy.

Moulton, B. (2018). The measurement of output, prices, and productivity: What's changed since the boskin commission? The Hutchins Center on Fiscal and Monetary Policy at the Brookings Institution.

Peters, M. (2020). Heterogeneous markups, growth, and endogenous misallocation. *Econometrica 88*(5), 2037–2073.

Piketty, T. (2018). *Capital in the twenty-first century*. Harvard University Press.

Piketty, T. and E. Saez (2003). Income inequality in the united states, 1913–1998. *Quarterly journal of economics 118*(1), 1–41.

Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. Von Wachter (2019). Firming up inequality. *Quarterly journal of economics 134*(1), 1–50.

Tirole, J. (1988). *The theory of industrial organization*. MIT press.

Voronina, M. (2022). Endogenous growth and optimal market power. Mimeo Harvard.