On the nonparametric identification of productivity growth in the presence of selection

Jonas Gathen*†

19 Sep 2023

Abstract

How much of aggregate productivity growth is driven by common productivity improvements across firms and how much is driven by the better selection of firms? In this short paper, I study the non-parametric identification of these two sources of productivity growth. I propose a framework that nests various endogenous and exogenous growth models and requires only (mild) restrictions on exit behavior and the shocks that drive heterogeneity in productivity. In this framework, separate identification of selection and a common, time-varying productivity growth term involves solving two selection biases. The first is a static or compositional selection bias whereby average productivity can increase due to entry and exit in the absence of any within-firm changes in productivity. The second, dynamic selection bias, arises from the persistence of productivity shocks and is driven by mean reversion. I show how a weighted average of within-plant productivity changes allows separate identification. Weights are chosen such that the dynamic selection bias exactly cancels and can be found by constructing a stationary distribution of the underlying productivity shocks from a synthetic panel of firms over time. I show how the identification approach can be extended to studying cohort effects and more general forms of heterogeneous productivity growth.

*Toulouse School of Economics (TSE), Email: jonas.gathen@tse-fr.eu.

†I am grateful for useful comments by Jihyun Kim, Yongseok Shin, Zhang Chen and Matteo Bobba. All remaining errors are my own.

Introduction

Productivity growth is widely seen as the major driver of economic growth.[1] Productivity growth, in turn, is driven by changes in the productivity distribution across all firms in the economy, which can be decomposed into (1) changes in productivity among surviving firms and (2) changes in productivity due to the exit of incumbent firms and the entry of new firms (selection).

How much of aggregate productivity growth is driven by common productivity improvements across firms and how much is driven by the better selection of firms? This question is at the heart of modern theories of economic growth (e.g. Acemoglu et al. 2018; Luttmer 2007). As I formalize in this paper, separate identification of selection and productivity growth is difficult because of at least two different forms of selection bias: a "static" and a "dynamic" selection bias. The "static" selection bias arises from changes in the productivity distribution due to changes in the composition of firms. For example, if less productive firms are more likely to exit or entering firms are more productive, average productivity can increase in the absence of any firm-level changes in productivity. A natural response would be to look only at within-firm productivity changes, which indeed can fully control for static selection bias under the right conditions on the timing of decisions and revelation of information. However, within-firm productivity changes are still biased estimates of common productivity growth if selection is on a persistent component. Intuitively, with persistence in productivity, surviving firms that have seen a history of very good productivity shocks are more likely to mean revert in the future. Similarly, young entrants may enter with lower productivity and mean revert upwards over time. Given a strong prior that productivity changes show persistence, focussing on within-firm productivity changes of surviving plants thus introduces a "dynamic" selection bias.

The previous literature has dealt with this identification issue mostly in structural models with (strong) parametric assumptions on firm productivity processes and the entry and exit of firms (e.g. Clementi and Palazzo 2016; Garcia-Macia, Hsieh, and Klenow 2019; Ottonello and Winberry 2023). In this short paper, I study the non-parametric identification of heterogeneous productivity growth in the presence of selection. Non-parametric identification allows to discipline drivers of economic growth more robustly and clarifies the empirical variation that identifies them.

Specifically, I consider a production side where firm productivity can be separated into an endogenous drift component and a - potentially highly persistent - productivity shock that drives variation in productivity across similar firms. As I show further below, this framework nests many endogenous and exogenous growth

[1] For one, there seems to be a consensus in the literature that the large variations in income per capita across countries are mostly accounted for by differences in total factor productivity (e.g. see Jones 2016), which are primarily driven by productivity differences and by misallocation of resources. Furthermore, productivity growth is the centerpiece of most modern growth theories (Acemoglu 2008; Jones 2022).

models in the literature. I show separate identification of the two components without making functional form assumptions on the arbitrarily time-varying drift component, the productivity shock process nor the firm entry and exit processes that drive endogenous selection. The constructive identification proof in this paper clarifies how a weighted within-plant productivity change identifies the endogenous drift component, with appropriate weights chosen such that the "dynamic" selection bias exactly cancels out.

Technically, identification comes from two key identifying restrictions. The first restriction is that idiosyncratic productivity shocks follow the same underlying general first-order, ergodic Markov process across firms and time. This allows for flexible forms of error dependence (e.g. it is always very unlikely to move from the bottom to the top of the distribution, but more likely to move from the top to the bottom), but requires this dependence to stay fixed over time. Importantly, the setup still allows for different endogenous responses of entry, exit and productivity trend growth to productivity shocks over time. Stated differently, the framework treats the productivity shock as coming from a flexible exogenous process (which cannot be affected by policy), while allowing policy (and other aggregate changes) to flexibly affect entry, exit and the endogenous drift.

The second set of identifying restrictions concerns firm exit. While the setup allows fully flexible entry, identification requires that (1) firms' exit decisions are not based on future productivity shock realizations and that (2) there is a form of common support in firms' exit decisions.[2] Together, these two assumptions guarantee that one can always correct for any dynamic selection bias by reweighting the distribution of surviving firms' productivities to account for the productivities of firms that exited. The common support restriction ensures that there are always surviving firms with similar productivity as exiting firms, an assumption that can be tested and which finds strong empirical support in firm- and plant-level data. Importantly, firms' endogenous exit decisions can still be based on productivity, on other observables, on future expectations as well as unobservables.

I start with the basic setup where the trend component is the same across all firms. This nests standard neoclassical growth models that feature exogenous aggregate productivity growth and firm selection (e.g. Clementi and Palazzo 2016; Luttmer 2007) as well as endogenous growth models with a common productivity growth component (e.g. Romer 1990). In this setup, identification of the common time-varying productivity growth trend comes from the insight that all changes in the distribution of within-firm changes in productivity can be attributed to changes in the common trend component as long as the distribution of firms over the idiosyncratic productivity shocks is at the stationary distribution. While the initial distribution over

[2]The first assumption only ensures that we can directly condition on firm productivity. If exit decisions are partly based on signals about future productivity, other common selection-correction steps can be taken. I do not explore this issue further here.

idiosyncratic productivity shocks may be far from the stationary distribution and selection prevents it from reaching or getting closer to the stationary distribution, the stationary distribution can be identified if the panel is long enough and the key identifying restrictions on selection hold. Specifically, the restrictions on selection ensure that whenever a firm selectively exits based on idiosyncratic productivity between t and $t+1$, one can find a matching surviving firm with similar overall productivity (and hence similar idiosyncratic productivity). In practice, this leads to constructing a stationary distribution of idiosyncratic productivity from a synthetic panel of units over time and then enforcing the weights of this distribution to solve for changes in the aggregate trend component.

The paper then extends the basic setup in two important directions. First, I show how one can allow for different levels of productivity and different productivity trend growth across observed groups of firms. This allows for heterogeneity across industries and other fixed firm characteristics. Importantly, this also allows to identify cohort effects, generalizing the setup such that it nests an important strand of the (endogenous) growth literature where growth is driven by imitation and learning (e.g. Asturias et al. forthcoming; Luttmer 2007; Sampson 2016). Identification of such group-specific effects relies on constructing group-specific stationary distributions.[3] Second, I discuss how the setup can be extended to allow for heterogeneous productivity drift across the distribution. Intuitively, the above identification argument allows to identify any changes in the entire productivity distribution as long as the distribution of idiosyncratic productivity is at the stationary distribution.

The structure of the paper is as follows: I start with a brief overview of the basic setup and discuss which (endogenous) growth frameworks are and are not nested by this setup. After showing identification and estimation in this context, I extend the basic setup in the two dimensions discussed above. The last section concludes and discusses how the identification arguments in this paper extend to other contexts such as price and demand dynamics.

Basic setup

Firms i at time t have productivity: $Y_{it} = Z_t exp(\varepsilon_{it})$. Z_t captures the underlying (endogenous) drift in productivity that - in the basic setup - is assumed to be shared across all firms and that I call aggregate technology throughout. We only observe firms that are producing over time and given that firms enter and exit, the productivity panel is generally unbalanced. ε_{it} is an exogenous, idiosyncratic firm-specific

[3] Importantly, the non-parametric identification in this paper identifies cohort effects not from productivity differences of cohorts at entry, because there is potentially selection at entry. They are rather identified from weighted productivity differences after a cohort is sufficiently long observed to be able to construct a stationary distribution. This means that cohort effects for the most recent cohorts are not generally identified without stronger assumptions on entry.

productivity process and one can think of changes in ε_{it} as knowledge shocks from the unpredictability of R&D, learning shocks in the production process, shocks in the production process itself or any other idiosyncratic exogenous component that affects productivity. This exogenous productivity shock process is allowed to follow a general first-order, ergodic Markov process: $\varepsilon_{it}(\varepsilon_{it-1})$. This means that the probability and magnitude of different shocks can vary systematically depending on the level of productivity of the firm. For example, shocks in the production process may be less volatile for highly productive firms as uncorrelated production shocks decrease in importance as the firm scales up.

Which growth models are nested by this basic setup and which are not? It is helpful to distinguish which models are nested on the basic production side and which are nested by the identification assumptions. On the production side, the setup is linked to models where firm productivity drives firm dynamics (Lucas 1978). Identification is most closely linked to discrete-time firm dynamics models in the tradition of Hopenhayn (1992), but as I show in the Appendix, in principle, the setup can also nest the production side for continuous-time firm dynamics models in the tradition of Klette & Kortum (2004) and Aghion & Howitt (1992) such as Acemoglu et al (2018). The key economic restriction in the basic setup is that the (potentially endogenous) trend Z_t is common across all firms. This nests standard firm dynamics models with neoclassical growth. At the business cycle level, the setup nests standard real business cycle models with aggregate uncertainty (e.g. Clementi & Palazzo 2016) and the identification in this paper allows to separately identify aggregate TFP shocks from selection effects. The setup also nests models of long-run growth where Z_t is exogenous (e.g. Midrigan & Xu 2014). It also nests endogenous growth models that give a shared Z_t across firms as in classic endogenous growth models (e.g. Romer 1990), as well as endogenous growth models where Gibrat's Law holds, that is, trend growth is independent of firm size. The latter holds in many endogenous growth models where Z_t is common across all firms, but can vary with changes in aggregates such as prices, taxes and R&D subsidies, such as Bento & Restuccia (2017) and many models of creative destruction (e.g. Klette & Kortum 2004, Atkeson & Burstein 2010, Peters 2020).

The production side of the basic setup does not nest endogenous growth models where endogenous productivity improvements are heterogeneous across firms. A recent example is Ottonello & Winberry (2023) where costs of investments in productivity improvements do scale with firm-level productivity, but do not necessarily scale exactly such that Gibrat's law holds and productivity growth is constant across firms. The basic setup also does not nest models where Z_t varies across groups or cohorts as in various growth models where entrants' productivity increases over time, such as endogenous growth models of imitation and learning (e.g. Alvarez, Buera, and Lucas 2013; Lucas and Moll 2014; Perla and Tonetti 2014; Buera and Oberfield 2016; Sampson 2016) and exogenous growth variants in this spirit (e.g. Luttmer 2007, Asturias et al 2023).

In the extensions, I show how group- or cohort-specific trend growth is relatively easily dealt with, while more general heterogeneous productivity growth is more difficult.

Throughout, the log-additive form is going to be more useful, so that:

$$y_{it} = z_t + \varepsilon_{it}(\varepsilon_{it-1})$$

In the following, we are interested in identifying the path of z_t (up to a normalization of $z_0 = 0$). For this, let us first note the two main reasons why this is a difficult problem. Both reasons can be exemplified by looking at changes in average y_{it} over time:

$$\frac{1}{N_t}\sum_{i\in\mathcal{N}_t} y_{it} - \frac{1}{N_{t-1}}\sum_{i\in\mathcal{N}_{t-1}} y_{it-1} = \underbrace{z_t - z_{t-1}}_{\Delta z} + \underbrace{\frac{1}{N^S_{t,t-1}}\sum_{i\in\mathcal{N}^S_{t,t-1}} \Delta s_{it}}_{\text{Survivor } \Delta s} + \underbrace{\frac{1}{N^E_t}\sum_{i\in\mathcal{N}^E_t} s_{it}}_{\text{Entry } \bar{s}} - \underbrace{\frac{1}{N^X_t}\sum_{i\in\mathcal{N}^X_t} s_{it-1}}_{\text{Exit } \bar{s}}$$

Changes in average y_{it} over time only identify changes in z under the special case that average changes in s among survivors as well as changes in average s between exiting and entering plants exactly cancel out. The first main problem is that we expect changes in average productivity among entering and exiting plants to be different from zero, since plant exit and entry likely depend on s_{it} and z_t and because the productivity pool of entrants may be very different from the productivity pool of incumbent plants. For example, if there is selection on entry and exit, we expect entering plants to be more productive than exiting plants and these terms not to cancel out. One way to indirectly test this, is to compare y_{it} among plants that exit next period and y_{it} among plants that entered this period.

Given that we observe who enters and who exits, we can also directly focus on surviving plants. The focus on surviving plants gives the simpler problem:

$$\frac{1}{N^S_{t,t-1}}\sum_{i\in\mathcal{N}^S_{t,t-1}} \Delta y_{it} = \underbrace{z_t - z_{t-1}}_{\Delta z} + \underbrace{\frac{1}{N^S_{t,t-1}}\sum_{i\in\mathcal{N}^S_{t,t-1}} \Delta s_{it}}_{\text{Avg mean reversion of survivors}}$$

However, this still leaves a bias term that gives the average mean reversion in productivity across survivors. Whether this term vanishes or is different from zero depends on where the distribution of survivors is and how selected the group of survivors are. In general, this bias term is not zero because entering plants that survive are not necessarily drawn from the ergodic distribution of productivity s (and hence may see either positive or negative mean reversion) and surviving plants may see negative mean reversion over time as long as they are positively selected on productivity.

In the following, I discuss identification of z and then propose an estimator for the time path of z.

Identification

Proposition 1 (Main identification result). *Under the following four assumptions:*

1. *(**Common first-order stationary Markov process**) s follows the same general first-order, stationary ergodic Markov process for all i & t.*
2. *(**Selective exit**). The decision to exit after period t can flexibly depend on observables and unobservables X_{it} as well as productivity s_{it}, but may not depend on future productivity s_{it+1}. Specifically,*

$$\mathbb{P}(exit) = f(X_{it}, s_{it}, z_t) \quad with \quad \mathbb{P}_t(exit) \perp\!\!\!\perp s_{i,t+1}|s_{i,t}$$

3. *(**No complete exit over s**) $\mathbb{P}_t(exit|s_{it}) < 1\, \forall s \in Supp(s)$*
4. *(**Connected support in s**) For each period t, there exists at least a subset of the support of s in that period which is fully contained in the support of all s in all future periods. Formally: $\forall t, \exists S_t \subset Supp(s_{it})$ for which $S_t \subset \cup_{\tau > t} Supp(s_{i\tau})$.*

the path $z_t\, \forall t$ is identified given some normalization z_τ for some $\tau \in [0, T]$ and $\max t \equiv T \to \infty$.

Proof. To already convey the idea of a suitable estimator for the time path of z_t, let us proof Proposition 1 constructively. Identification proceeds sequentially in two fundamental steps. In the first step, I show identification of the density of the stationary distribution of s, which is identified for $t \to \infty$. In the second step, the density of the stationary distribution is used to identify the path of z_t backwards by starting at some final time T. The density of the stationary distribution is key because it can be used to construct weights under which a weighted difference Δy_{it} exactly identifies Δz_t. Specifically, there exist weights ω_s such that $\sum_{i \in \mathcal{N}^S_{T+1,T}} \omega(s_{iT})(s_{iT+1} - s_{iT}) = 0$ (where $\sum_i \omega_s(s_i) = 1$). These weights recover the stationary distribution of s. Denote by $f^{SS}(s)$ the density of the stationary distribution at s and by $f_t(s)$ the density of the distribution of s at time t. Assuming that this distribution shares the support of the stationary distribution, we have:[4]

$$\lim_{N \to \infty} \sum_{i \in \mathcal{N}^S_{t+1,t}}^{N} \frac{f^{SS}(s_{it})}{f_t(s_{it})} \Big(log(s_{it+1}) - log(s_{it}) \Big) = 0$$

[4]Potentially add regularity conditions here. E.g. maybe need that sample size increase maps to strictly monotonic increase in mass everywhere?)

The weights are thus defined by $\omega_s(s_{it}) \equiv \frac{f^{SS}(s_{it})}{f_t(s_{it})}$ and are a function of the unknown density function of the stationary distribution of s. To identify the density $f^{SS}(s)$, start with the distribution of plants at t_0 over known y_{i0}. The idea is to follow survivors (as they follow the process for s), while replacing exiting plants with plants that stay in the panel that have similar y_{it}. More formally, denote the initial set of plants by $\mathcal{N}_0$ where each plant is given a uniform weight $\tilde{\omega}_{i0} = \frac{1}{N_0}$. We are interested in updating $\mathcal{N}$. For this, pass on the weight of each surviving plant and redistribute the weight of each plant that exits to close plants around them.[5] This gives $\mathcal{N}_1$. Updating in this way allows to eventually pass on weight to plants that have entered the economy, even if they have entered in an arbitrarily selective way. As $t \to \infty$, surviving plants will eventually populate the entire support of s and this procedure gives a synthetic sample $\mathcal{N}_\infty$ with weights $\tilde{\omega}_{i\infty}(s_{i\infty})$ that directly identify the density $f^{SS}(s)$.

The second step of the proof takes the identified density $f^{SS}(s)$ and works backwards from time T. Normalizing the final value z_T, one can show that z_{T-1} solves a fixed point problem. Specifically:

$$\sum_{i \in \mathcal{N}^S_{T,T-1}} \omega_{\hat{s}_{T-1}(z_{T-1})}(y_{iT} - y_{iT-1}) = z_T - z_{T-1} + \sum_{i \in \mathcal{N}^S_{T,T-1}} \omega_{\hat{s}_{T-1}(z_{T-1})}(s_{iT} - \hat{s}_{iT-1}(z_{T-1})) = z_T - z_{T-1}$$

where the last equality holds only if the guess z_{T-1} is correct. It thus gives a nonlinear equation in z_{T-1} (since the weights and the right-hand side depend on z_{T-1}). (Give conditions under which this has a unique solution). One can iterate on this procedure to identify the path of z_t backwards. At any point in time $t < T-1$, one can also alternatively guess z_{T-1} and instead of using weights at all, estimate the bias term $\sum_{i \in \mathcal{N}^S_{T,T-1}}(s_{iT} - \hat{s}_{iT-1}(z_{T-1}))$ directly using future survivors with similar s. This alternative relaxes the assumption of a common support with the stationary distribution and instead only requires that we can build a sample with similar survivors - requiring a much weaker connected support.

Estimation

Estimation proceeds along the lines of the constructive identification proof. In the first step, one sequentially builds the synthetic panel with weights $\omega_s(s_{it})$ (which sum to 1 in each year). In principle, one can use any standard matching estimator for passing on the weight for exiting plants. Below, I show that a Kernel matching estimator works well, because matching is only based on one variable and the Kernel estimator distributes the weight widely across multiple observations, reducing variance.[6]

One can then estimate $f^{SS}(s)$ using observed s in the last period T and constructed weights $\hat{\omega}_s(s_{iT})$. Any

[5] As $N \to \infty$ and the assumption that exiting probabilities are always strictly lower than one, there always exists a plant that is arbitrarily close to an exiting plant. Under regularity conditions on the process for s, it is sufficient to be close in a symmetric way, that is one can for example use a Kernel. Make this more explicit.

[6] Note that one can readily match based on further variables to minimize the risk of model misspecification.

standard density estimator such as a Kernel density estimator works here. To reduce variance, one can also estimate $f^{SS}(s)$ on the last x periods (where x is at the discretion of the researcher). In general, for any fixed T, the bias on the estimated weights is increasing in the persistence of the process as well as in the distance of the initial distribution from the stationary distribution. That is, for large T and low persistence, one can use more periods in the end to estimate $f^{SS}(s)$.[7] Once the density is estimated, one can then proceed in sequentially estimating the path z_t. For each period t and for each guess of z_t, this means one has to estimate $f_t(s_{it}(\hat{z}_t))$. Again, any standard density estimator works here. One can then construct the weights according to: $\omega_{st}(s_{it}) \equiv \frac{f^{SS}(s_{it})}{f_t(s_{it})}$. Alternatively, one can choose not to use weights and instead directly estimate the bias from mean reversion. In that case, one can again use any kind of matching estimator to match plants in t with productivity $s(\hat{z}_t)$ to future survivors with similar s. The variance in the bias estimate reduces with the number of matched plants such that one to many matches are recommended. As before, a Kernel-based matching estimator is a natural choice here. In either the approach with weights or with an estimated bias term, one then finds z_t that solves the fixed point problem, requiring a standard root finder. I return to the practicalities of estimation and Monte Carlo evidence further below.

Extensions

In the following, I discuss a few important extensions to the previous setup and estimation that allow to study more general and empirically relevant cases.

More general forms of endogeneous growth

This subsection extends the previous setup to allow for an endogenous growth process z that differentially affects the entire distribution y. The key restriction that is required is that z keeps the relative distribution of units unchanged, that is, $z(y_{-1})$ is monotonic in y_{-1}. The identification idea is the same as before: in the presence of an ergodic Markov shock process s, knowing the stationary distribution of s allows to identify changes in the entire distribution of y that is due to a general growth process z. Specifically, at the stationary distribution of s, the quantile change in z is given by:

$$\Delta z_q \equiv z_t|_q - z_{t-1}|_q = \Delta y_{it}|_q$$

(Any feedback on this part is greatly appreciated!!)

[7] A formal treatment of optimally solving this trade-off is beyond the scope of this short note.

Group-specific growth rates

The previous setup can easily be extended to allow for group-specific differences in the level and growth of z_t. That is, $y_{igt} = z_{gt} + \varepsilon_{it}(\varepsilon_{it-1})$. Such groups could be different industries or fixed types of firms or more generally prices that differ across groups of products or industries. In principle, one can simply apply the estimator separately to different groups in the data. However, this does not allow to compare the level of groups over time. The benefit of the setup is that such level comparisons are actually possible. To see how, note that at the stationary distribution of ε_{it}, differences in the distributions of y_{igt} across groups capture only differences in the level of z_{gt} across groups. Hence, one can construct synthetic stationary distributions of firms i within each group at time T and infer differences in the level of z_{gT} from average differences in y_{igT} enforcing these group-specific weights.

Once level differences are identified, one can obtain group-specific growth rates from applying the previous estimator for each group separately. For efficiency reasons, one may also want to exploit information across groups. In the case of using future firms to estimate the mean reversion bias, one can then also match firms across groups since future differential growth rates are known and one can directly match on ε_{it}.

Cohort effects and changes in entrants

Another important issue is dealing with cohort effects, which can be seen as a special case of group-specific growth rates. In the case of firms, new cohorts of entrants may enter with different productivity. General differences in the process of ε_{it} across cohorts are generally hard to deal with as long as one is not willing to simply separately estimate processes for different subgroups in the data. However, one can make further progress by only assuming that new cohorts c enter with a different level in productivity, such that: $y_{ict} = z_t + z_c + \varepsilon_{ict}(\varepsilon_{ict-1})$. Similar to group-specific growth rates, one can build cohort-specific synthetic stationary distributions and then identify z_c from differences in y_{icT} enforcing the distribution weights. Once z_c is identified, one can then identify the entire time path z_t from the full sample using differences in $\tilde{y}_{ict} \equiv y_{ict} - z_c$. One can also make further assumptions on z_c, such as monotonicity in c or a constant trend: $z_c = \alpha * t$. This ensures that z_c is also identified for young cohorts for which synthetic stationary distributions at T cannot be constructed.

An important remark is that the estimation of cohort effects as proposed here provides considerably more flexibility than is commonly assumed in the literature. The reason is that the approach makes no restrictions on changes in the distribution of entering plants over time. The approach allows each cohort of entrants to be arbitrarily selected and also arbitrarily differ in the degree of selection across cohorts. If one is willing to make stronger assumptions on the initial distribution of entrants, one can also directly identify z_c from differences in the entrant distributions at entrance, using information on the aggregate z_t.

Conclusion

Acemoglu, Daron. 2008. *Introduction to Modern Economic Growth.* Princeton University Press.

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. 2018. "Innovation, Reallocation, and Growth." *American Economic Review* 108 (11): 3450–91.

Asturias, Jose, Sewon Hur, Timothy J. Kehoe, and Kim J. Ruhl. forthcoming. "Firm Entry and Exit and Aggregate Growth." *American Economic Journal: Macroeconomics*, forthcoming. https://doi.org/10.1257/mac.20200376.

Clementi, Gian Luca, and Berardino Palazzo. 2016. "Entry, Exit, Firm Dynamics, and Aggregate Fluctuations." *American Economic Journal: Macroeconomics* 8 (3): 1–41.

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow. 2019. "How Destructive Is Innovation?" *Econometrica* 87 (5): 1507–41.

Jones, Charles I. 2016. "The Facts of Economic Growth." In *Handbook of Macroeconomics*, 2:3–69. Elsevier.

———. 2022. "The Past and Future of Economic Growth: A Semi-Endogenous Perspective." *Annual Review of Economics* 14: 125–52.

Luttmer, Erzo GJ. 2007. "Selection, Growth, and the Size Distribution of Firms." *The Quarterly Journal of Economics* 122 (3): 1103–44.

Ottonello, Pablo, and Thomas Winberry. 2023. "Investment, Innovation, and Financial Frictions."

Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy* 98: S71–102.

Sampson, Thomas. 2016. "Dynamic Selection: An Idea Flows Theory of Entry, Trade, and Growth." *The Quarterly Journal of Economics* 131 (1): 315–80.