

Moore's Law and Economic Growth*

Pablo D. Azar
Federal Reserve Bank of New York

September 25, 2022

Abstract

Over the past 60 years, semiconductor sizes have decreased by 50% every 18 months, a trend known as Moore's Law. This paper studies the physical channel through which Moore's Law affects GDP growth. A new model incorporates physical constraints on firms' production functions and allows for new types of spillovers from changes in the physical characteristics of products. I use the model and a new dataset to estimate the effect of the electronic miniaturization channel on productivity growth. The results show that between 11.74% and 18.63% of productivity growth in the 1960-2019 period can be attributed to electronic miniaturization.

1 Introduction

Developing new combinations of existing materials is a key driver of economic growth. The discovery of these new combinations is often driven by improvements in the physical properties of the materials being combined. For example, modern computers and electronics only exist because semiconductor sizes have been decreasing by a factor of 50% every 18 months since the 1960s, a fact known as Moore's Law. In spite

*The views expressed in this paper are those of the author and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. I am grateful to Nathan Godin, Sergio Olivas and Rebecca Reubenstein for excellent research assistance, and thank Daron Acemoglu, Enghin Atalay, Marco Cipriani, Adam Copeland, Arnaud Costinot, Marco Del Negro, Julian Di Giovanni, Chad Jones, Gustavo Joaquim, Nick Krasney, Jetson Leder-Luis, Michael Lee, Maxim Pinkovskiy, Harrison Pugh, Alp Simsek, Cory Smith, Noam Tanner for their helpful comments. Correspondent Address: 33 Liberty Street, New York, NY, 10045. E-mail: pablo.azar@ny.frb.org.

of this, most economic models of innovation do not take into account improvements in materials' physical properties when modeling technological progress.

To capture how improvements in the properties of materials lead to new technological innovations, this paper builds a model where firms have physical constraints on their production functions, and where these constraints dictate which inputs can and cannot be combined. This captures real-world examples of physical limitations in production, such as engines needing to be above a certain efficiency threshold, airplanes having to satisfy strict weight limits, and electronics needing to be small enough to fit into office equipment, cars and industrial machines.

As the physical properties of materials improve, industries' constraints are relaxed and new productivity-enhancing combinations are developed. I use this observation to develop a dynamic version of the model where the physical properties of materials change over time, allowing for new feasible combinations of inputs. I apply this dynamic model to study how Moore's Law affects economic growth. As semiconductors shrink, electronic components can be incorporated into exponentially many new combinations that can be used to produce goods and services. Each different combination of input electronic components corresponds to a different "recipe" to make a machine, and each recipe corresponds to a random productivity draw from a thick-tailed distribution. Through this process, Moore's law leads to exponentially many new ways to provide goods and services, yielding exponential growth in productivity. Conversely, I show that if the individual electronic components do not shrink beyond a positive lower bound, then no growth in productivity is possible.

Using the predictions of the model, and a new dataset on product weights, I estimate how the *physical channel* of Moore's Law affected productivity growth. More concretely, I answer the following counterfactual question: *How much lower would aggregate productivity be if electronics were just as powerful as today, but as large in size as they were in the 1960s?* I focus on the special case where the combinatorial constraints are Knapsack constraints limiting the weight of materials that can be used in a manufactured product such as a car or an industrial machine. A structural estimation exercise shows that electronic miniaturization leads to an annualized productivity increase between 0.07% and 0.12% in the entire 1960-2019 period, accounting for 11.74% to 18.63% of all TFP growth. In the second half of this period (1990-2019), when computers and electronics were adopted by a large number of industries and households, the productivity increase from miniaturization ranged

between 0.12% and 0.20%, accounting for 18.91% to 30.03% of all TFP growth.¹

Empirically, the closest point of comparison to this paper are the results by Jorgenson and Stiroh (1999,2000) and Jorgenson (2005), and the more recent work by the Bureau of Economic Analysis (Barefoot et al. 2018). Jorgenson and Stiroh show that information technology (including computers, software and communications equipment) accounted for a majority of TFP growth in the 1973-2002 period.² The Bureau of Economic Analysis reports that, for the 2006-2016 period, the digital economy accounted for 28% of real GDP Growth. An important distinction between existing results and those in this paper is that I estimate only the *physical channel* of Moore’s Law, and not any other gains in productivity due to semiconductors becoming faster or memory becoming larger. Nevertheless, the results in Sections 5 show that a substantial fraction of productivity growth is due to miniaturization.

1.1 Making Combinatorial Growth More Realistic

The theoretical model in this paper is motivated by Romer (1993)’s observation that, as ideas arrive to the market, they can be combined with each other to yield exponential growth in the number of existing products. Weitzman (1998) made this intuition formal by proposing a model of recombinant growth, where new ideas are generated by researchers who combine existing ideas. Acemoglu and Azar (2020) and Jones (2021) built on Weitzman’s work by developing models where growth is driven by the ability of firms to choose from an exponentially increasing number of sets of input suppliers. Each combination is a “recipe” with a different level of productivity, and firms choose the cost-minimizing combination. As new inputs arrive, the number of recipes increases exponentially, and economic output grows.

One issue with the existing combinatorial growth models is that they allow for arbitrary combinations of inputs. In reality, many combinations of inputs are not feasible or non-sensical.³ My model addresses this objection head-on by incorporating


¹The range in estimates depends on whether one wants to capture productivity spillovers using Hulten’s Theorem. The lower bound assumes that Hulten’s Theorem does not apply (e.g. because production functions are very rigid), while the upper bound applies Hulten’s theorem to obtain productivity spillovers.

²Table 8 in Jorgenson (2005) shows that information technology contributed 68% of all TFP growth in the 1973-1989 period, 88% of all TFP growth in the 1989-1995 period, and around 66% of all TFP growth in the 1995-2002 period.

³An early objection attributed to Akerlof by Jones (2021), states that “Yes the number of possible combinations is huge, but aren’t most of them like chicken ice cream!”

combinatorial constraints restricting which sets of inputs can be feasibly combined.

Figure 1: Knapsack Constraint in Automobile Manufacturing



Note: When manufacturing a car, the total sum of component weights cannot exceed the weight of the car, which in this figure is 4000 pounds. Small electronic components, such as modern GPS devices (0.5 lbs), cameras (0.06 lbs) and Electronic Control Units (2 lbs) all fit comfortably inside the car. However, state-of-the-art supercomputers (5500 lbs) cannot be incorporated into the car without violating the weight constraint.

As a simple example of these combinatorial constraints, consider the production a car, illustrated in Figure 1. Each unit produced requires an engine, wheels, tires, a chassis, as well as fuel, exhaust, cooling, lubrication and electrical systems. In addition, there are optional components such as GPS navigation systems, rear and front-view cameras, and, more recently, specialized computers for self-driving.⁴ Different combinations of these systems will lead to different quality cars. The quality-adjusted production function can be represented using a menu of Leontief functions, indexed by the set S of components used in production:

$$Y = A(S) \min\left\{L, \min_{j \in S} \left\{ \frac{X_j}{a_j} \right\} \right\} \quad (1)$$

where $A(S)$ represents the quality of a car that uses a set of inputs S , and a_j is the number of units of component j that need to be used to produce one car.

Of course, not every set of inputs S is feasible. For example, if the set S does not include an engine, or a steering system, then the car cannot be produced. This can be expressed more formally by saying that there's a set E of essential inputs such that $E \subset S$ for every feasible set S of components that car manufacturers could choose. Furthermore, and crucial to the model introduced in this paper, there are

⁴Other components, such as anti-lock breaking systems (ABS), used to be optional but are now mandatory in the United States and other jurisdictions.

physical constraints on the elements of the set S . The simplest such constraint is a *weight constraint*: a car that is designed to weigh at most 4000 pounds has to have components whose weights add up to less than 4000 pounds. More formally, if θ_j represents the weight of the j^{th} component in pounds, the set S of input components must satisfy the Knapsack constraint

$$\sum_{j \in S} a_j \theta_j \leq 4000. \quad (2)$$

Because of this weight constraint, a state-of-the-art supercomputer which weighs 5500 pounds cannot be used as a component in a car. This fact links advances in computer miniaturization to advances in car quality. As transistors have shrunk over the last few decades, new productivity-enhancing components such as Electronic Control Units and GPS navigation have become increasingly available to be used in production.

Besides providing a more realistic framework for combinatorial growth, the introduction of physical constraints allows for new and interesting comparative statics. A change in one product’s physical characteristics can trigger “physical spillovers”, which affect the physical properties of materials that are downstream in the supply chain. In Section 5, I illustrate these physical spillovers with an empirical exercise that shows how changes in the weight of semiconductors affect the weight of computers and electronics. As semiconductors shrink by 50% every 18 months, computers and electronics also shrink, and become easier to embed into manufactured products. I use the heterogeneity in product weights and thresholds induced by these spillovers to estimate the effect of electronic miniaturization on aggregate productivity.

1.2 Further Related Work

Production Networks and Distortions This paper belongs within the input-output network literature started by Leontief. In particular, it relies on the models of input-output linkages proposed by Long and Plosser (1983) and analyzed further by Ciccone(2002), Gabaix (2011), Jones (2011), Acemoglu et. al. (2012), Acemoglu et. al (2017), Bartleme and Gorodnichenko (2015), Bigio and La’o (2020), Baqaee (2018), Fadinger, Ghiglino and Teteryatnikova (2018), Liu (2017), Baqaee and Farhi (2019a, 2019b), Caliendo et. al (2018) and Liu and Tsyvinski (2021).

The main modeling contribution of this paper is to augment the Long and Plosser model by having physical constraints which distort firms’ production sets, leading to endogenous adoption of suppliers. The main empirical contribution is to use this augmented model to estimate how changes in the physical characteristics of computers and electronic devices increased manufacturing productivity in the United States.

Because the input-output weights in the model are endogenous, this paper is closely related to work by Jones (2013), Bigio and La’O (2020), Liu (2018), Fadinger et. al. (2018), and Caliendo et. al. (2018) who analyze models of production networks with distortions. This paper is also related to work by Carvalho et. al. (2016) and Baqaee and Farhi (2019a), who study models of endogenous input-output networks where firms’ production functions are not Cobb-Douglas. In all of these models, the input-output structure of the economy is endogenous. The economic intuition is that—in the presence of distortions such as markups, taxes and subsidies—the allocation of goods in the economy may be inefficient. Furthermore, when the production function is not Cobb-Douglas, increases in productivity in one sector may not increase aggregate TFP as much as they would in a non-distorted economy.

In my model, distortions arise from the combinatorial constraints on production, and affect allocative efficiency and the propagation of shocks. For example, if industry j has a weight constraint that binds from above, it will substitute away from heavy inputs into lighter inputs. If a heavy machine in industry i becomes more productive, this increase in productivity may not propagate to industry j because industry j ’s weight constraint will prevent it from demanding more units of that machine.

Endogenous Production Networks The literature on production networks and distortions considers changes in the input-output network at the intensive margin. However, there is a recent and growing literature that analyzes changes in the production network at the extensive margin. Atalay et. al. (2011) and Carvalho and Voigtlander (2015) give rule-based⁵ models of network formation. More recently, Lim (2018), Oberfield (2018), Acemoglu and Azar (2020), Acemoglu and Tahbaz-Salehi (2020), Taschereau-Dumouchel (2020) and Elliott, Golub and Leduc (2021) consider models where the structure of the network is determined by individual firms

⁵For example, Atalay et. al. (2011) use a preferential attachment model.

maximizing profit.

Lim (2018) builds a model where the structure of production networks is determined by the tradeoff between the benefits of new firm-to-firm linkages, and the costs of maintaining these relationships. These costs are offset by the benefits of the new relationships that are formed. In Oberfield’s (2018) model, each firm bargains with a set of suppliers and chooses one input, leading suppliers with very high productivity to become “superstars”, supplying many firms at once. Acemoglu and Azar (2020) present a model where firms choose arbitrary sets of suppliers, where each possible set of suppliers induces a different production technology. They use this framework to construct an endogenous growth model, where the source of growth is the exponentially increasing number of combinations of suppliers that a firm can use in its production function. Acemoglu and Tahbaz-Salehi (2020) develop a framework where firms purchase inputs from “customized” suppliers and use bargaining to reach agreement on how to split the surplus generated by their relationship. They apply this model to understand how firm failures can cascade throughout the supply chain and amplify recessions. Taschereau-Dumouchel’s (2020) model studies the cascading effect of firms deciding to enter or exit the market. An exiting firm will have negative spillovers on its clients and suppliers, leading multiple firms to exit the network simultaneously. Elliott, Golub and Leduc’s (2021) model shows how—when firm-to-firm linkages can fail—arbitrarily small shocks to productivity can cause a substantial decrease in output.

In my model, changes in the parameters of the model can lead to a change in the input-output network at the extensive margin. A sudden change in the size of a product (say, because of technological advancement) can lead to a change in the structure of the input-output network and a jump in output. For example, desktop computers replaced mainframes in the 1980s, laptops replaced desktops in the 1990s and 2000s, and tablets replaced laptops in the 2010s. In all these changes, the shift in demand toward lower-weight products was made possible by technological improvements.

Endogenous Growth Another close point of comparison for this paper is Weitzman (1998) together with the follow-up work of Auerswald, Kauffman, Lobo, and Shell (2000) and Ghiglino (2012). The main difference with these models of recombinant growth is that they assume that there is an idea-generating function that

grows exponentially when new ideas are obtained by combining existing ideas. In contrast, I assume that materials arrive more slowly, linearly with time instead of exponentially. Firms can use combinations of products in production, with each combination yielding a different productivity draw. As there is an exponentially growing number of combinations—under the right distributional assumptions—productivity grows exponentially as new materials and new combinations arrive.

Physical Properties and International Trade In the international trade literature, Hummels and Skiba (2004) show that products’ physical features, such as their weight, affect trading costs. They use this observation to argue for the existence of the Alchian-Allen effect, where high quality varieties within an industry are more likely to be exported than low quality varieties. Evans and Harrigan (2003), Harrigan (2005), Baldwin and Harrigan (2007) and Harrigan and Deng (2010) expand on this observation and show that heavier goods are traded between physically close countries, while lighter goods are traded by physically distant countries. Similarly, lighter goods are much more likely to be traded by air than by sea.

Other papers in the international trade literature study how trading frictions affect supply chains. Chaney (2014) uses firm-level data from French exports to develop a model of supply chain formation with informational frictions. Antràs and Chor (2013) develop a model with contracting frictions to study vertical integration along the supply chain. Antràs, Fort, and Tintelnot (2017) build a model where countries choose a combinatorial set of other countries to import from. In contrast with the first two papers—which deal with informational and contracting barriers to supply chain formation—the friction affecting supply chain formation in this paper comes from physical constraints which prevent the adoption of some suppliers. The model of Antràs, Fort, and Tintelnot (2017) is closer to mine, with the main differences being that they use supermodularity of the profit function to ensure tractability. This supermodularity property may not apply when firms have combinatorial constraints on which sets of inputs are feasible. To obtain tractability, I assume that productivity draws for different combinations of suppliers are drawn from independent random distributions.

1.3 Roadmap

The rest of the paper is organized as follows. Section 2 shows the baseline static model of production with combinatorial constraints. Section 3 shows that a generically unique equilibrium exists, characterizes the equilibrium under standard assumptions on productivity distributions, and derive comparative statics showing how changes in physical properties affect output. Section 4 develops a dynamic model where improvements in products’ physical characteristics drives GDP growth. Section 5 shows the results of an empirical exercise, illustrating how shrinking semiconductors affected US productivity during the 1960-2019 period. Section 6, discusses potential future work, and concludes. The Appendices contain proofs along with additional discussion of the algorithms used to compute the number of feasible combinations.

2 A Static Model

This section presents a tractable general equilibrium model that captures how changes in the physical properties of primary goods allows them to be combined in new ways, expanding the production possibilities frontier.

In the model, there is a finite number N of primary industries, which one can think of as producing materials or components that have physical properties. Changes in the physical properties of one primary good have effects on the physical properties of other primary goods as well. In particular, this type of spillover effect captures the notion that smaller semiconductors have led to smaller electronic components and smaller electronic machines in general.

Firms in the final industry can choose the subset of primary industries’ goods that they use as inputs. Different subsets S of primary inputs represent the many different “recipes” that can be used to produce the final good. Each subset S further corresponds to an independent random productivity draw $A(S)$, so that new combinations of inputs—such as replacing mechanical buttons with touchscreens—may lead to a higher quality version of the final good.

A novel aspect of the model is that firms in the final industry cannot choose the set S of inputs in an arbitrary way. Instead, S is restricted by constraints that may depend on the physical properties of the primary goods contained in S .

2.1 Market Structure

Primary Industries Let $\mathcal{N} = \{1, \dots, N\}$ be the set of primary industries. Each primary industry $i \in \mathcal{N}$ has a strictly quasi-concave, continuous and increasing production function with constant returns to scale $Y_i = F_i((X_{ij})_{j=1}^N, L_i)$, where X_{ij} is industry i 's demand for good j and L_i is industry i 's demand for labor. Good i has a vector of physical characteristics $\theta_i \in \mathbb{R}^p$, which is determined by a continuously differentiable function $\theta_j = \theta_j(\zeta)$. The input ζ is a vector of fundamental properties which can affect any primary industry.

In addition to the standard assumptions of quasi-concavity, continuity, and monotonicity, I make the assumption that labor is essential in the production of every primary good, so that $F_i(X, 0) = 0$ for all X . This assumption prevents equilibria where production is fully automated.

Example 1 (Leontief Production Functions and Size Spillovers). *Consider an economy where primary industry i has a Leontief production function $Y_i = A_i \min\{L_i, \min\{\frac{X_{ij}}{\alpha_{ij}}\}\}$. With some abuse of notation, we can write $\alpha_{ij} = 0$ when industry i does not use industry j 's good as an input.*

In this example, each primary good has only one physical property, its size, denoted by $\theta_i \in \mathbb{R}_{\geq 0}$. Since producing one unit of good i requires α_{ij} units of good j , the size of one unit of good i will be given by $\theta_i = \sum_{j=0}^N \alpha_{ij}\theta_j + \zeta_i$, where $\zeta_i \geq 0$ is an idiosyncratic factor affecting the size of good i . In matrix form, this equation is given by

$$\theta = \alpha\theta + \zeta \tag{3}$$

As long as $(I - \alpha)^{-1}$ exists, we can write

$$\theta = (I - \alpha)^{-1}\zeta, \tag{4}$$

which is a continuously differentiable function of ζ . For convenience of notation, we denote the Leontief inverse matrix $(I - \alpha)^{-1}$ by \mathcal{L} , and write $\theta = \mathcal{L}\zeta$. The intuition behind this specification is that if the idiosyncratic weight ζ_j of good j decreases by $\Delta\zeta_j$, then the weight θ_i of industry i 's good will decrease, through network spillover effects, by $\mathcal{L}_{ij}\Delta\zeta_j$.

The Final Industry The final industry produces a single good using primary industries' goods as inputs. The production technology is given as a menu of production functions, indexed by the subset $S \subset \mathcal{N}$ of primary inputs used in production. More concretely, given a set of inputs S , the final industry has a strictly quasi-concave, continuous and increasing production function with constant returns to scale $Y_f(S) = A(S)F(S, (X_{fi})_{i \in S})$, where $A(S) \in \mathbb{R}_{>0}$ is a Hicks-Neutral productivity shifter that depends on S and $(X_{fi})_{i \in S}$ is a vector of demands for primary goods.

Firms in the final industry cannot use arbitrary sets S of primary inputs. Instead, they are bound by a constraint

$$G((\theta_i)_{i \in S}, S) \leq \tau \tag{5}$$

where G is a vector-valued function and τ is a vector of thresholds. I denote by $\mathcal{F}(\theta, \tau)$ the collection of all feasible input subsets $S \subset \mathcal{N}$ for which the constraint (5) holds.⁶

Example 2 (Spanning Tree Constraints). *A very prominent example of a real world constraint is the construction of a spanning tree in a network. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph where the vertices \mathcal{V} represent cities, and edges \mathcal{E} represent connections between cities. These connections could be roads, telecommunication links or electric cables. A spanning tree is a subset of edges $S \subset \mathcal{E}$ satisfying the following two conditions:*

- **Spanning Condition.** *For every city $v \in \mathcal{V}$, there is at least one edge $e \in S$ such that v is an endpoint of e .*
- **Tree Condition.** *No set of edges in $e \in S$ form a cycle.*


⁶Throughout this paper, I use the combinatorial constraint given by equation (5). There is a more general constraint which allows the feasibility of set S to vary with the demands $(X_{fi})_{i \in S}$ of primary goods in the set S . This constraint is given by

$$G(Y, (X_{fi})_{i \in S}, (\theta_i)_{i \in S}, S) \leq \tau Y \tag{6}$$

where G is a vector-valued function that has constant returns to scale in $(Y, (X_{fi})_{i \in S})$ and is quasi-convex as a function of $(Y, (X_{fi})_{i \in S})$. Since $G(Y, (X_{fi})_{i \in S}, (\theta_i)_{i \in S}, S)$ has constant returns to scale, the general constraint (6) is equivalent to $G(1, (\frac{X_{fi}}{Y})_{i \in S}, (\theta_i)_{i \in S}, S) \leq \tau$. When the production function F is Leontief, the intermediate demands per unit of output $\frac{X_{fi}}{Y}$ are fixed, and (6) reduces to (5).

The spanning condition ensures that every city is connected by the tree S . The tree condition ensures that S has no cycles, preventing waste in many applications. For example, in electricity transmission, sending energy through a cycle $v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_n \rightarrow v_1$ would be less efficient than not using that cycle in the transmission path.

Figure 2: Illustration of Spanning Tree Constraint from Example 4



Note: This figure illustrates a spanning tree in a graph $G = (V, E)$. The seven vertices represent different cities. The edges represent possible routes between pairs of cities, and each edge e has a price P_e . The selected edges in red represent the minimum-cost spanning tree. That is, they are a cost-minimizing subset of edges which visit every city in the network, while not having any cycles.

One can model these kind of applications more concretely as a production problem with combinatorial constraints. The primary goods are the edges $e \in \mathcal{E}$, each of which may have a cost P_e . The final good is produced using a constant returns to scale function $Y_f(S) = F((X_e)_{e \in S})$ subject to the combinatorial constraint that S is a tree.⁷ For instance, if $F((X_e)_{e \in S}) = \min_{e \in S} (X_e)$, then industries in the final firm solve the minimum-spanning tree problem $\min_{S \text{ is a tree}} \sum_{e \in S} P_e$.

Households There is a representative household which consumes C units of the final good, and has a strictly increasing and concave utility function $U(C)$. The household supplies one unit of labor inelastically, and has a budget constraint $P_f C \leq W$, where P_f is the price of the final good and W is the wage paid to labor. Throughout, I assume that the wage is the numeraire, so that its price $W = 1$.

⁷For the purposes of this paper, we only need to know that $\mathcal{F} = \{S \subset \mathcal{E} : S \text{ is a tree}\}$ is a well-defined collection of subsets. This collection can be defined in a standard way, via an inequality such as the one given in Equation (5). The interested reader is referred to <http://www.columbia.edu/~cs2035/courses/ieor6614.S16/mst-1p.pdf> for an integer programming formulation in the form of Equation (5).

2.2 Equilibrium

In this subsection, I give a definition of equilibrium. Because final industry firms can choose their set of inputs, their overall production technologies may be non-convex. Thus, in some situations, it will be more appropriate to state firms' optimization objectives in terms of cost-minimization instead of profit maximization.⁸

Definition 1 (Marginal Cost Functions). *Let P_i denote the price of industry i 's good. The marginal cost functions of primary industries, and the final industry are defined as follows:*

1. For a primary industry i ,

$$K_i(P_1, \dots, P_N) = \min_{L_i, (X_{ij})_{j=1}^N} \sum_{j=1}^N P_j X_{ij} + L_i, \text{ subject to } F_i((X_{ij})_{j=1}^N, L_i) = 1. \quad (7)$$

2. For the final industry,

$$K_f(P_1, \dots, P_N) = \min_{S \in \mathcal{F}(\theta, \tau), (X_{fi})_{i \in S}} \sum_{i \in S} P_i X_{fi} \quad (8)$$

$$\text{Subject to: } A(S)F(S, (X_{fi})_{i \in S}) = 1 \text{ and } G((\theta_i)_{i \in S}, S) \leq \tau. \quad (9)$$

Definition 2. *An equilibrium is a tuple $\mathcal{E} = (P_1^*, \dots, P_N^*, P_f^*, C^*, L^*, Y^*, X^*, S^*)$ such that*

1. Markets are competitive

$$P_i^* = K_i(P_1^*, \dots, P_N^*) \text{ for all } i \in \mathcal{N} \text{ and } P_f^* = K_f(P_1^*, \dots, P_N^*). \quad (10)$$

2. Firms in the final industry choose input sets to minimize marginal costs

$$S^* \in \arg \min_{S \in \mathcal{F}(\theta, \tau)} \min_{(X_{fi})_{i \in S}} \sum_{i \in S} P_i^* X_i^* \text{ subject to the constraints in (13)}. \quad (11)$$

⁸More concretely, a firm in the final industry with a constant returns to scale production function will obtain zero profits by choosing any feasible set of inputs S and choosing $(X_{fi})_{i \in S}$ to maximize profits. This *does not mean* that all feasible sets S are equally likely to be used in equilibrium. Instead, firms which choose S to minimize marginal costs will be the only ones that will receive positive demand in equilibrium, since they can charge the lowest prices and undercut all competitors.

3. Given S^* and P^* , firms choose X^* and L^* to minimize marginal costs.
4. Households choose C^* to maximize utility subject to their budget constraint.
5. Markets clear, so that

$$Y_i^* = F_i((X_{ij}^*)_{j=1}^N, L_i^*) = \sum_{j=1}^N X_{ji}^* + X_{fi}^* \quad \forall i \in \mathcal{N} \quad (\text{Primary Market Clearing})$$

$$Y_f^* = A(S^*)F(S^*, (X_{fi}^*)_{i \in S}) = C^* \quad (\text{Final Market Clearing})$$

$$\sum_{i=1}^N L_i^* = 1. \quad (\text{Labor Market Clearing})$$

3 Equilibrium Analysis and Comparative Statics

In the model presented in Section 2, firms in the final industry face a menu of mutually exclusive production technologies, indexed by the set S of primary inputs that they choose. Thus, production technologies are non-convex. In spite of this non-convexity, an equilibrium exists under very mild necessary and sufficient conditions. Furthermore, the equilibrium is generically unique and efficient.

Existence and Uniqueness It is immediate that, for an equilibrium to exist with competitive markets, prices have to equal marginal costs, and Equation (10) has to hold. The first Theorem in this section shows that, if Equation (10) has a solution, then an equilibrium exists.⁹

Theorem 1. *An equilibrium exists if and only if there exists a vector of positive prices $P_1, \dots, P_N > 0$ such that Equation (10) holds.*

Uniqueness Under mild conditions, the equilibrium will be generically unique. First, I show that if the marginal cost functions are strictly concave, the equilibrium price vector is always unique.

Proposition 1. *If the primary industry cost functions $(K_i(P))_{i=1}^N$ are strictly concave as a function of prices, then the equilibrium price vector is unique if it exists.*

⁹All proofs in this section follow standard techniques, and are deferred to Appendix A.

Proposition 1 holds, for example, when the primary industry production functions are CES with non-zero elasticity of substitution. For the special case where production functions are Leontief, the cost function $K_j(P)$ is linear in prices, and not strictly concave. Nevertheless, when the Leontief coefficients add up to less than 1, there is a unique price vector.

Proposition 2. *Let $F_i(A_i, L_i, X) = \min\{A_i L_i, \min_{j \in \mathcal{N}} \{\frac{X_{ij}}{\alpha_{ij}}\}\}$. If the matrix $(I - \alpha)$ is invertible and all entries in $(I - \alpha)^{-1}$ are non-negative, then there exists a unique equilibrium price vector P^* .*

Even when prices are unique, equilibrium quantities may not be. This is because there may exist two different input sets S^*, S^{**} which both minimize the final industry's costs. If this happens, the final industry may have two different primary input demands X^*, X^{**} corresponding to the two different input sets. I show below that this only happens for a subset of productivity parameters $(A(S))_{S \in \mathcal{F}}$ that has measure zero. More formally, I define an equilibrium to be generically unique as follows.

Definition 3 (Generically Unique Equilibrium). *The equilibrium is generically unique if the set*

$$\mathcal{A} = \{(A(S))_{S \in \mathcal{F}} : \text{There exist at least two distinct equilibria } \mathcal{E}, \mathcal{E}'\}$$

has Lebesgue measure zero in $\mathbb{R}^{|\mathcal{F}|}$.

Theorem 2. *If the equilibrium price vector P^* is unique, then the equilibrium \mathcal{E} is generically unique.*

3.1 Efficiency

Because each intermediate industry faces a non-convex production function, the welfare theorems no longer apply. Thus, one may ask whether a social planner may be able to produce an allocation of goods that is welfare-improving over the equilibrium allocation. I show below that this does not happen. The underlying economic intuition is that firms in the final industry will choose a set S which minimize their marginal costs. If they do not, then another firm will undercut them. These cost

savings are passed down to households. To simplify the [proof](#), I make the additional assumption that production and utility functions are differentiable.

Theorem 3. *Assume that all production functions are continuously differentiable in labor and intermediate inputs, and that the household utility function is continuously differentiable in consumption. Then the equilibrium \mathcal{E} is Pareto Efficient.*

3.2 A Tractable Formula for Expected Output

To derive analytical equilibrium aggregates, I make the following standard assumption.

Assumption 1. *For all sets $S \in \mathcal{F}$, the productivity term $A(S) = \bar{A}(S)\phi(S)$, where $\bar{A}(S)$ is a deterministic component that may depend on the set S , and $\phi(S)$ is a random variable drawn independently from a Frechet distribution with CDF $\Psi(x) = e^{-x^{-\kappa}}$.*

This assumption is used by Acemoglu and Azar (2020) on their model of endogenous production networks, and is inspired by the use of Frechet productivity draws in Kortum (1997) and Eaton and Kortum (1998) to microfound gravity equations in international trade models. To give a tractable formula for aggregate output, I define what the cost of intermediate products would be if all productivity terms were deterministic (that is, if $A(S) = \bar{A}(S)$, and $\phi(S) = 1$).

Definition 4 (Deterministic Cost Function). *Given an input set $S \in \mathcal{F}$, a vector P of primary prices, and a threshold τ , the deterministic cost function is given by*

$$\bar{K}_f(S, P, \tau) = \min_{(X_{fi})_{i \in S}} \sum_{i \in S} P_i X_{fi} \quad (12)$$

$$\text{Subject to: } \bar{A}(S)F(S, (X_{fi})_{i \in S}) = 1. \quad (13)$$

I can now characterize the equilibrium output as a function of deterministic cost functions and the feasible collection $\mathcal{F}(\theta, \tau)$.

Theorem 4 (Static Output Characterization). *Suppose that Assumption 1 holds. Then the output of the final industry is a Frechet Random variable with shape parameter κ and scale parameter $(\sum_{S \in \mathcal{F}(\theta, \tau)} \bar{K}_f(S, P, \tau)^{-\kappa})^{\frac{1}{\kappa}}$.*

Corollary 1. *Suppose that Assumption 1 holds. Then, expected log-output is given by*

$$\mathbb{E}[\log Y_f] = \frac{1}{\kappa} \log\left(\sum_{S \in \mathcal{F}(\theta, \tau)} \bar{K}(S, P, \tau)^{-\kappa} \right) + \frac{\gamma}{\kappa} \quad (14)$$

where γ is the Euler-Mascheroni constant.

3.3 Comparative Statics

Using equation (14), one can derive comparative statics for expected log-output. Because of the combinatorial nature of production, these comparative statics will be discontinuous: a small change in the physical properties of a product may lead to a large change in aggregate GDP. This is a reflection of a realistic pattern, where productivity growth can often be drastic.

To state the discontinuous comparative statics results, I use Dirac's Delta function $\delta(\cdot)$, which can be informally understood as a distribution with an infinite point mass at 0, and zero mass everywhere else on the real line. More formally, we have the following definition.

Definition 5. *Dirac's Delta function is a linear functional operating on the space of smooth distributions with compact support, defined by $\langle \delta, \nu \rangle = \int_{-\infty}^{\infty} \delta(\tau) \nu(\tau) d\tau = \nu(0)$.*

Lemma 1. *Let $\delta(\cdot)$ be Dirac's Delta function. Then,*

$$\frac{\partial}{\partial \theta_i} \sum_{S \in \mathcal{F}(\theta, \tau)} \bar{K}(S, P, \tau)^{-\kappa} = - \sum_{S \subset \mathcal{N}} \delta(\tau - G(\theta, S)) \bar{K}(S, P, \tau)^{-\kappa} \frac{\partial G}{\partial \theta_i} \quad (15)$$

Applying the chain rule to equation (14) and applying Lemma 1, one can derive the following first-order comparative static.

Proposition 3.

$$\frac{\partial \mathbb{E}[\log Y_f]}{\partial \theta_i} = - \frac{1}{\kappa} \frac{\sum_{S \subset \mathcal{N}} \delta(\tau - G(\theta, S)) \bar{K}(S, P, \tau)^{-\kappa} \frac{\partial G}{\partial \theta_i}}{\sum_{S \in \mathcal{F}(\theta, \tau)} \bar{K}(S, P, \tau)^{-\kappa}} \quad (16)$$

Proposition 3 gives the expected change in log-output given a change in the physical properties of one good. We can see that when θ_i increases, the effect on expected log-output depends on how this increase affects the function G . If $\frac{\partial G}{\partial \theta_i}$ is positive (as in a knpsack constraint), then log-output will decrease because the constraint $G(\theta, S) \leq \tau$ gets closer to being binding. On the other hand, if $\frac{\partial G}{\partial \theta_i}$ is negative, then expected output will increase, because the constraint $G(\theta, S) \leq \tau$ becomes more relaxed. While this is intuitive, it allows us to capture general physical constraints, where changes in the physical property θ may make the constraint tighter or looser, depending on what we are modeling.

Making the comparative statics continuous with a continuum of industries. We can think of our model in Section 2 as representing the supply chain of one final industry. With only one industry that has a combinatorial constraint, it is natural for comparative statics to be discontinuous because technological progress is drastic: either a new feasible recipe S is discovered, or it is not.

In practice, there may be multiple industries, with different thresholds τ —something that is reflected in the data in Section 5. To capture this heterogeneity in thresholds, I make the following assumption.

Assumption 2. *There is a continuum of industries with different thresholds. The mass of industries with threshold τ is given by a measurable function $\nu(\tau)$, which has support in $[0, +\infty)$.*

Under this assumption, I derive a continuous analogue of Proposition 4.

Proposition 4. *Suppose Assumption 2 holds. Then*

$$\frac{\partial \mathbb{E}[\log Y_f]}{\partial \theta_i} = -\frac{1}{\kappa} \frac{\sum_{S \subset \mathcal{N}} \nu(G(\theta, S)) \bar{K}(S, P, \tau)^{-\kappa} \frac{\partial G}{\partial \theta_i}}{\sum_{S \in \mathcal{F}(\theta, \tau)} \bar{K}(S, P, \tau)^{-\kappa}} \quad (17)$$

Proposition 4 adds another layer of intuition. If the mass $\nu(G(\theta, S))$ of industries with threshold $G(\theta, S)$ is large, then the effect that a change in θ will have on expected output will be larger. The reason for this is immediate: an industry with threshold $\tau = G(\theta, S)$ will have a binding combinatorial constraint for a new recipe S . When the combinatorial constraint binds, a small change in θ will make the new recipe S feasible, leading to increased productivity.

3.4 Spillovers of Physical Properties

In this subsection, I give a quantitative formula for how changes in the physical properties of one primary good propagate through the economy and have an effect on aggregate output. As in Section 2, the vector θ of physical properties is a differentiable function $\theta(\zeta)$ of some fundamental properties ζ_1, \dots, ζ_K .¹⁰ The following result is an immediate corollary from Proposition 4 and the chain rule.

Corollary 2.

$$\frac{\partial \mathbb{E}[\log Y_f]}{\partial \zeta_j} = -\frac{1}{\kappa} \sum_{i=1}^N \frac{\sum_{S \subset \mathcal{N}} \nu(G(\theta, S)) \bar{K}(S, P, \tau)^{-\kappa} \frac{\partial G}{\partial \theta_i} \frac{\partial \theta_i}{\partial \zeta_j}}{\sum_{S \in \mathcal{F}(\theta, \tau)} \bar{K}(S, P, \tau)^{-\kappa}} \quad (18)$$

This result is very general, and applies to many areas outside miniaturization. To give a concrete example, room temperature superconductors that can be produced reliably would have impacts across the energy, transportation, defense, industrial and medical sectors (Johns et. al. 1990). The conductivity, temperature, and size of superconducting inputs could be modeled as property vectors $\theta(\zeta)$ which may depend on the properties ζ of the materials used to create these inputs. The impact of superconductors across industries could be modeled through a constraint function $G(\theta, S)$, which would specify size, temperature and conductivity constraints of different industrial machines using superconductors.

4 A Dynamic Model With Empirical Implications

In this section, I extend the model to allow for multiple periods and to allow for changes in the physical characteristics $\theta(t)$ of primary inputs over time. In this general dynamic setting, I show that changes to the physical characteristics of primary goods can lead to long-run growth. In the more specific setting of electronic miniaturization, I will show a converse: long run growth is impossible without miniaturization.

One qualitatively novel aspect of this model is that changes in the physical properties of materials can allow for the organic introduction of new products. In contrast, most existing models of economic growth have an abstract notion of *research* which

¹⁰Here I emphasize that the number of fundamental properties K may be different from the number of primary industries N , or the dimension p of the property vector θ .

leads to the arrival of new products.¹¹ Using combinatorial constraints, I can give a more concrete notion of what this research entails. Some primary products may not be feasibly used at all—i.e. because they are too heavy to be used in the final good, as computers were in the 1930s and 1940s.¹² Only once their physical characteristics change, do those products “enter the market,” and they can be used as intermediate goods in the production of the final good.

Market Structure The dynamic model is exactly the same as the static model, with two exceptions:

1. there is a countable number of time periods indexed by $t \in \mathbb{N}$; and
2. there are a countable number of primary goods indexed by $i \in \mathbb{N}$.

Let $\mathcal{F}(t) = \{S \subset \mathbb{N} : G(\theta(t), S) \leq \tau\}$ denote the set of all feasible combinations at time t , and let $\mathcal{N}(t) = \{i \in \mathbb{N} : \text{there exists } S \in \mathcal{F}(t) \text{ such that } i \in S\}$ denote the set of all primary goods which can be used in at least one feasible input combination. I call this the set of *usable primary inputs* at time t . Note that $|\mathcal{F}(t)| \leq 2^{|\mathcal{N}(t)|}$. To ensure that output at any time t is finite, I make the following assumption.

Assumption 3. *For every time $t \in \mathbb{N}$, the number of usable primary inputs $|\mathcal{N}(t)|$ is finite.*

Assumption 3 allows the model to have an infinite number of products, almost all of which have not entered the market yet. However, at any point in time, only a finite number of primary goods can be used in production of the final good. Quantum computers are a good example of a good which exists today, but cannot be feasibly used in production. While they exist in very specific laboratory conditions, it would

¹¹Romer (1990)’s model assumes that there’s an infinite number of varieties in the world, and that they can be unlocked by investing in research. Aghion and Howitt (1992)’s model assumes that there’s an infinite number of quality improving ideas, which themselves are unlocked by research. In both of these models, a positive research intensity yields an exponentially increasing flow of new varieties or ideas, which itself leads to economic growth. Acemoglu and Azar (2020) take a different approach, allowing only one new product or idea to arrive at each period of time. In their model, as in this one, exponential growth comes from the exponential combination of products. However, this kind of growth still requires a new product to arrive periodically (in an exogenous fashion), so that the number of combinations increases exponentially.

¹²Another good example of this phenomenon is steam engines. James Watt worked on steam engines for 15 years before the invention of high-precision machining tools made steam engines energy-efficient and commercially feasible.

require significant advances in engineering and physics to make them practical to use in *any* application.

The rest of the section proceeds as follows. In subsection 4.1, I show a result that illustrates a linear relationship between growth and the number of feasible combinations $\log |\mathcal{F}(t)|$, which can be taken to the data. In subsection 4.2, I relax the assumptions in the simplified example of subsection 4.1, and show an asymptotic linear relationship between expected log-output and $\log |\mathcal{F}(t)|$. The results in subsections 4.1 and 4.2 allow for arbitrary combinatorial constraints, not just Knapsack ones. In subsection 4.3, I focus on the context of electronic miniaturization with Knapsack constraints, and give an example illustrating how new goods can become usable over time and drive growth via miniaturization. In subsection 4.4, I show that long-run growth is impossible if primary goods do not shrink beyond a given threshold. Subsection 4.5 show how—even when there are a finite number of combinations and long-run growth is impossible—the short-run growth patterns reflect the data. The results in subsection 4.6 show how the growth rate of the economy can be affected by different combinatorial constraints.

4.1 A Simplified Model

In this subsection, I show a simplified model that illustrates the log-linear relationship between output and the number of feasible combinations. Each primary good is produced using only labor, so that $Y_i = L_i$. The final good is produced using a menu of Leontief production functions

$$Y_f(S) = A(S) \min_{i \in S} (X_i)_{i \in S} \quad (19)$$

subject to the combinatorial constraint

$$\sum_{i \in S} \theta_i(S)(t) < 1. \quad (20)$$

The marginal cost of good i is $K_i = 1$, and the marginal cost of the final good is $K_f(S) = \frac{|S|}{A(S)}$. Since markets are competitive and the household's budget constraint binds, final output is equal to $Y_f(S) = \frac{1}{K_f(S)} = \frac{A(S)}{|S|}$.¹³ To further simplify, I use the

¹³More concretely, if S^* is the set chosen by firms in the final industry, then $P_f = K_f(S^*)$. The

following instantiation of Assumption 1.

Assumption 4. *The productivity term $A(S)$ is equal to $\phi(S)|S|$, where $\{\phi(S)\}$ is a family of i.i.d. Frechet random variables with shape parameter κ and scale parameter 1.*

Under Assumption 4, we have $Y_f(S) = \phi(S)$, where ϕ is a Frechet random variable. Letting $\mathcal{F}(t)$ denote the collection of all feasible sets at time t , the optimal S is given by $S^*(t) = \arg \max_{S \in \mathcal{F}(t)} \phi(S)$. Using Corollary 1, write

$$\mathbb{E}[\log Y_f(S^*(t))] = \frac{1}{\kappa} \log |\mathcal{F}(t)| + \frac{\gamma}{\kappa}.$$

Let $g(t) = \Delta \mathbb{E}[\log Y_f(S^*(t))]$. The following Theorem is immediate from the above analysis

Theorem 5. *Suppose Assumptions 3 and 4 hold. Then the expected growth rate at time t is given by*

$$g(t) = \frac{1}{\kappa} \Delta \log |\mathcal{F}(t)|. \tag{21}$$

Theorem 5 holds for arbitrary combinatorial constraints $G(\theta, S) < \tau$, and does not require these constraints to be Knapsack constraints. Even with the simplified production structure where all primary goods are produced using only labor, it gives a very strong empirical connection between log-output and the log-number of feasible combinations $\log |\mathcal{F}(t)|$. In Section 5, I estimate equation (21) and use the estimated parameter κ to simulate the effect of Moore's Law on economic growth.

4.2 A General Dynamic Model with Asymptotic Empirical Implications

The simplified model of subsection 4.1 yields a very precise empirical implication, but does not allow for more general production structures. I now show an asymptotic equivalent of equation (21) that holds under much more general conditions. Using Corollary, 1, we can write expected log-output at time t for a general production

household's budget constraint is $P_f Y_f = 1$, yielding $Y_f = \frac{1}{P_f} = \frac{1}{K_f(S^*)}$.

structure as

$$\mathbb{E}[\log Y_f(t)] = \frac{1}{\kappa} \log\left(\sum_{S \in \mathcal{F}(t)} \bar{K}(S, P, \tau)^{-\kappa}\right) + \frac{\gamma}{\kappa}.$$

To gain some tractability, I need to ensure that growth is driven by the fact that new combinations allow new random productivity draws, and not by the fact that newly arriving sets become exponentially cheaper over time. Similarly, newly arriving sets S cannot become so expensive that they are impossible to use. To capture this more formally, I make the following assumption on how the deterministic costs of different input combinations change over time.

Assumption 5. *Let $P(t)$ be the equilibrium vector of primary good prices at time t . There exist functions $K_u(t), K_\ell(t)$ such that*

$$K_\ell(t) \leq \bar{K}(S, P(t), \tau) \leq K_u(t) \text{ for all } t, S \in \mathcal{F}(t) \quad (22)$$

$$\lim_{t \rightarrow \infty} \frac{\log K_\ell(t)}{t} = \lim_{t \rightarrow \infty} \frac{\log K_u(t)}{t} = 0. \quad (23)$$

The economic content of Assumption 5 is straightforward: growth isn't driven by the fact that any one combination becomes exponentially cheap over time, nor is it hampered by the possibility that new combinations become prohibitively expensive to use. While new primary goods arrive over time and new combinations S may be unlocked, the prices $\bar{K}(S, P(t), \tau)$ of these new combinations don't increase or decrease exponentially over time.

Assumption 5 is not onerous. In fact, it is satisfied under very mild conditions for CES functions and very general specifications of the deterministic component of the productivity term $\bar{A}(S)$, as shown in the following proposition.

Proposition 5. *Consider the family of constant elasticity of substitution production functions with elasticity $\sigma < 1$:*

$$F_i(X_i, L_i) = \left[\left(1 - \sum_{j=1}^{|\mathcal{N}(t)|} \alpha_{ij}\right) L_i^{\frac{\sigma-1}{\sigma}} + \sum_{j=1}^{|\mathcal{N}(t)|} \alpha_{ij} X_{ij}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \quad (24)$$

where $\sum_{i=1}^{\infty} \alpha_{ij}^\sigma < \chi$ for some constant $\chi < 1$, and for all $i \in \mathbb{N}$. Furthermore, the

final good production function is given by

$$Y_f(S, X) = \phi(S)\bar{A}(S)\left(\sum_{i \in S} X_i^{\frac{\rho-1}{\rho}}\right)^{\frac{\rho}{\rho-1}}, \quad (25)$$

where $\phi(S)$ is a Frechet random variable and $\bar{A}(S)$ satisfies $t^{-\nu} \leq \bar{A}(S) \leq t^\nu$ for some constant ν and all $S \in \mathcal{F}(t)$. Then Assumption 5 holds.

According to proposition 5, production functions can be arbitrary CES functions, and the deterministic productivity terms $\bar{A}(S)$ can oscillate through a very broad range $[t^{-\nu}, t^\nu]$ over time.

Assumption 5 and Corollary 1 yield an asymptotic version of Equation (21) which holds on average across all time periods, instead of pointwise at every time period. More concretely, we have the following Theorem.

Theorem 6. *Suppose Assumptions 1, 3, and 5 hold, and assume further that $\lim_{t \rightarrow \infty} \frac{\log |\mathcal{F}(t)|}{t} = D$ for some constant $D > 0$. Then*

$$\lim_{t \rightarrow \infty} \frac{\mathbb{E}[\log Y_f(t)]}{t} = \frac{D}{\kappa}. \quad (26)$$

Theorem 6 characterizes the average growth rate $\frac{\mathbb{E}[\log Y_f(t)]}{t}$ in terms of the asymptotic average number of feasible combinations at time t . The assumption that $\lim_{t \rightarrow \infty} \frac{\log |\mathcal{F}(t)|}{t} = D$ guarantees that new primary goods (and new feasible combinations) arrive at a steady average rate over time. For example, if there were no combinatorial constraints and D goods arrived on average each period, then $|\mathcal{F}(t)| \approx 2^{Dt}$, and $\log |\mathcal{F}(t)| \approx \log 2 \cdot Dt$. More generally—when there are combinatorial constraints— D measures the rate at which new input combinations become feasible. If D was zero, then there wouldn't be enough new combinations to generate growth. Similarly, if $\frac{\log |\mathcal{F}(t)|}{t}$ oscillated with time and had no limit, then the growth rate sequence $g(t)$ would also oscillate.

4.3 Physical Characteristics and Product Arrival

In this subsection, I give an example of how changing physical characteristics of products can be an engine for product arrival. We can think of the final product as a computer or a phone, and the primary products as different electronic components,

such as CPUs, GPUs, microphones, cameras and gyroscopes. The combinatorial constraint is a Knapsack constraint

$$\sum_{i \in S} \theta_i(t) < 1 \quad (27)$$

with a threshold τ for the final industry is normalized to 1. At time t , the size of product i is given by $\theta_i(t) = 2^{i-t-1}$. This formula reflects the fact that—while almost all products are too large to use in production at any given time t —they shrink over time so that every product can be feasibly used in production in the long run. Table 1 shows the changing size of the first 4 products over the first 3 time periods.

t	$\theta_1(t)$	$\theta_2(t)$	$\theta_3(t)$	$\theta_4(t)$	$ \mathcal{F}(t) $
1	0.5	1	2	4	1
2	0.25	0.5	1	2	3
3	0.125	0.25	0.5	1	7

Table 1: Changing product sizes over time for the example in Subsection 4.3. The last column shows the number of feasible combinations at each period t . In period 1, only good 1 can be used in production

At any given time t , only the first t primary goods can be used in the production of the final good, since $\theta_i(t) = 2^{i-t-1} \geq 1$ for all $i > t$. Furthermore, the sizes of the first t goods are $\frac{1}{2^t}, \frac{1}{2^{t-1}}, \dots, \frac{1}{2}$. Since $\sum_{i=1}^t 2^{i-t-1} < \sum_{n=1}^{\infty} 2^{-n} = 1$, any non-empty combination of the first t goods can be feasibly used to produce the final good. These two facts combined imply that $|\mathcal{F}(t)| = 2^t - 1$. The growth rate at time t is $g(t) = \frac{1}{\kappa} \log(2^t - 1) - \log(2^{t-1} - 1)$. As t grows large, the exponential term 2^t is much larger than 1. Using Theorem 6, we can approximate the asymptotic growth rate by $g^* \approx \frac{\log 2}{\kappa}$.

4.4 No Growth Without Miniaturization

The previous subsection characterized the growth rate when all electronic component sizes shrank over time. In this subsection, I show a converse result: if the combinatorial constraint takes a Knapsack form $\sum_{i \in S} \theta_i(t) \leq \tau$, and the sizes $\theta_i(t)$ are bounded below by a constant θ_{lower} , then there can be no long-run growth from the combination of electronic components.

Theorem 7. *Suppose Assumptions 1 and 5 hold. Assume further that $|\mathcal{N}(t)| \leq t^\nu$ for some constant $\nu > 0$, and that there exists a lower bound $\theta_{\text{lower}} > 0$ such that $\theta_i \geq \theta_\ell$ for all $i \in \{1, 2, 3, \dots\}$. Then $g^* = 0$.*

The intuition behind Theorem 7 is that—if primary goods do not shrink over time—then any input combination S used to produce the final good will only have a bounded number of items. Even if the potential number $|\mathcal{N}(t)|$ of items that can be combined grows to infinity, the number of combinations will grow at a subexponential rate, and asymptotic growth will be 0. The assumption that $|\mathcal{N}(t)| \leq t^\nu$ prevents growth from being driven by an exponential number of new primary goods arriving, rather than an exponential number of combinations.

4.5 Short-run Growth Waves with a Finite Number of Primary Goods

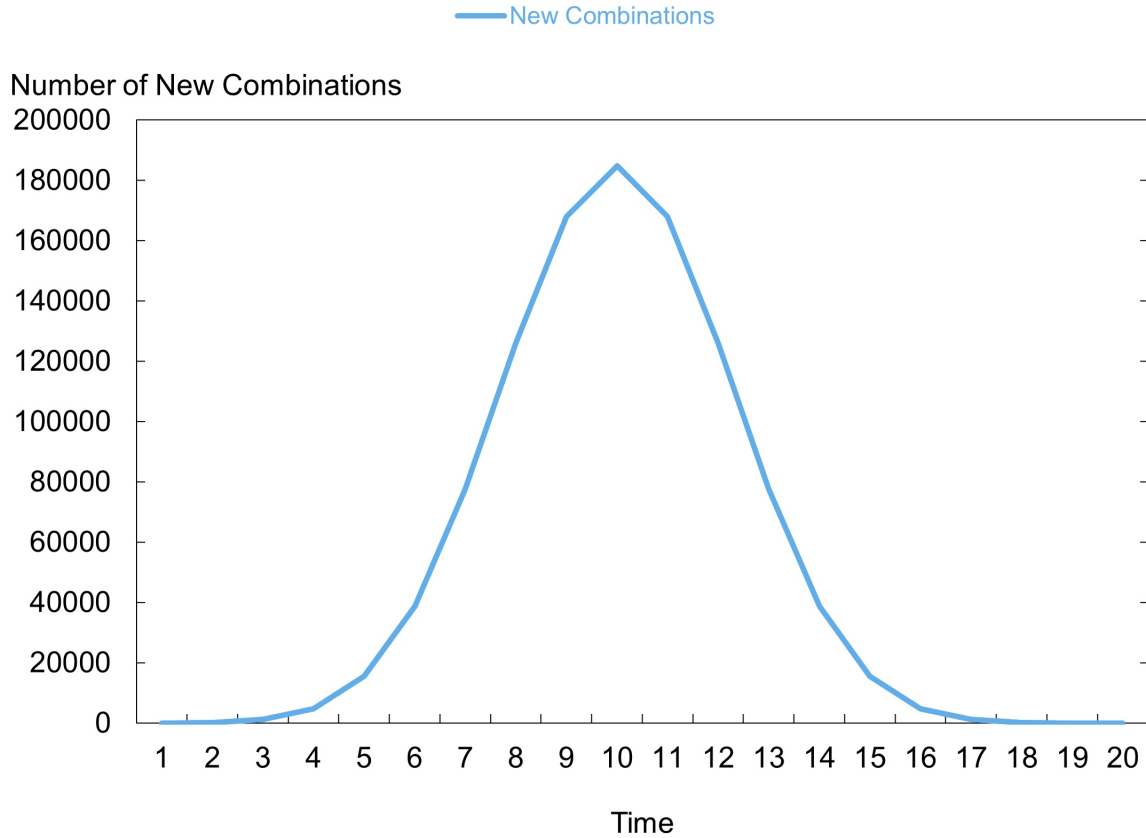
Throughout this section, I have assumed that there is an infinite number of primary goods, which may arrive slowly over time. If the number of primary goods is finite, then there will exist bounds such that $|\mathcal{N}(t)| \leq \bar{N}$ and $|\mathcal{F}(t)| \leq \bar{F}$, and $g^* = \lim_{t \rightarrow \infty} \frac{\log |\mathcal{F}(t)|}{t} = 0$. Nevertheless, even with a finite number of primary goods, there can be short-run growth. Furthermore, this short-run growth can exhibit interesting wave patterns, where productivity undergoes a drastic phase transition where productivity increases nearly overnight. This reflects our intuition that computers and electronics were specialized tools in the 60s, 70s and 80s, but then were adopted almost immediately overnight in the 90s and early 2000s by most industries. This intuition is captured theoretically by the results in this subsection, and is also reflected in the structural estimation exercise of Section 5.

As in Subsection 4.3, I normalize the final good threshold $\tau = 1$. There is a finite number of products N that does not change over time. Each product $i \in \{1, \dots, N\}$ has a size $\theta_i(t) = \frac{1}{t}$.¹⁴ At time $t = 1$, the number of feasible combinations is $\binom{N}{1}$. At time $t = 2$, combining pairs of inputs becomes feasible, so the number of feasible combinations is $\binom{N}{1} + \binom{N}{2}$. At time t , the number of feasible combinations is $|\mathcal{F}(t)| = \sum_{i=1}^t \binom{N}{i}$.

¹⁴Note here that the product sizes are not declining exponentially. This captures something we do observe in the data, where computers and other electronic products which partially depend on semiconductors, but which also partially depend on plastics, metals and other fixed-weight inputs, will decline in size over time at a sub-exponential rate.

As shown in Figure 3, $|\mathcal{F}(t)|$ grows slowly at first, increases sharply around $t = \frac{N}{2}$, and then quickly flattens out. This is consistent with the economic intuition of a drastic phase transition, and is reflected in the data in Figure 4.

Figure 3: Number of New Combinations with Finite Number of Goods



Note: This figure illustrates the example in subsection 4.5, where there are $N = 20$ goods and the size of all goods at time t is $\theta_i(t) = \frac{1}{t}$. The number of new combinations at time t is $\binom{N}{t}$. The number of new combinations over time follows a binomial bell curve, which is reflected in the data in Figure 4

4.6 Spanning Tree Constraints

So far, subsections 4.1 and 4.2 have illustrated results for general combinatorial constraints at a very abstract level, and subsections 4.3, 4.4 and 4.5 have focused on the electronic miniaturization context. In this subsection, I show how the growth

rate depends on the combinatorial structure of the economy in the spanning tree context of Example 2. For example, we can think of the telecommunications industry connecting cities across the US. As the number of cities grows, the number of ways $|\mathcal{F}(t)|$ in which the network can be connected also grows, and productivity increases.

More concretely, consider a network of cities $\mathcal{G}(t) = (\mathcal{V}(t), \mathcal{E}(t))$ that grows over time. The combinatorial constraint is a spanning tree constraint, as in Example 2. Every time period, a new city is founded, so that $|\mathcal{V}(t)| = t$. The degree of each city in the network—that is, the number of links connecting it to other cities—is a constant d that does not change with time. Alon (1990) shows that the number of spanning trees in $\mathcal{G}(t)$ is equal to $|\mathcal{F}(t)| = \{d[1 - \sigma(\mathcal{G}(t))]\}^t$, where $0 \leq \sigma(\mathcal{G}(t)) \leq \epsilon(d)$, and $\lim_{d \rightarrow \infty} \epsilon(d) = 0$.

The time t growth rate is $g(t, d) = \Delta \log |\mathcal{F}(t)| = t \log d + t \log(1 - \sigma(\mathcal{G}(t))) - (t - 1) \log d - (t - 1) \log(1 - \sigma(\mathcal{G}(t - 1))) = \log d + t \log(1 - \sigma(\mathcal{G}(t))) - (t - 1) \log(1 - \sigma(\mathcal{G}(t - 1)))$. Since $0 \leq \sigma(\mathcal{G}(t)) \leq \epsilon(d)$, the following inequality is immediate

Theorem 8. *For any fixed degree d , the growth rate at time t satisfies the bound*

$$|g(t, d) - \log d| \leq \epsilon(d). \tag{28}$$

Theorem 8 states that the growth rate of the economy hovers around $\log d$, with the error term $\epsilon(d)$ shrinking as d grows large. Thus, for very large d , we have that $g^* \approx \log d$. This is just one example of many, relying on existing combinatorics results, that shows how the growth rate depends on the combinatorial structure of the economy.

5 Structural Empirical Analysis

In this Section, I leverage the theoretical results of Section 4 to estimate how the physical channel of Moore’s Law affected productivity growth. Assuming production functions have a Leontief functional form, I estimate that between 11.74% of all productivity growth in the 1960-2019 period can be attributed to electronic miniaturization. This estimate—which does not include productivity spillovers in the supply chain captured by Hulten’s Theorem—can also be interpreted as a robust lower bound on the effect of Moore’s Law on productivity growth, which applies to

all CES functions. If instead we assume that production functions take the Cobb-Douglas functional form, then 18.63% of all productivity growth in the 1960-2019 period can be attributed to the physical channel of Moore’s Law.

A structural estimation of the physical channel of Moore’s Law on GDP growth must take into account a handful of factors. First, semiconductor weights in the simulation must decline by a factor of 2 every 18 months. Second, it must account for the physical spillover effects that smaller semiconductors have had on the weight of computers and electronics. Third, it must compute the larger number of feasible input combinations for manufacturing and service industries as semiconductors, computers and electronics become smaller over time. Finally—if Hulten’s Theorem applies—such a simulation has to include the productivity spillover effects that more productive machines have on the rest of the economy.

The rest of this section describes this structural estimation. Subsection 5.1 describes the data, including a novel dataset of product weights. Subsection 5.2 shows how to compute weight spillover effects from semiconductors to other electronic components and computers. Subsection 5.3 estimates the coefficient $\frac{1}{\kappa}$ in Equation (21), to obtain the elasticity of productivity with respect to the number of feasible combinations.¹⁵ Finally, Subsection 5.4 puts everything together, and shows an estimate over time of how changes in electronic component weights increased productivity—both under the assumption that production functions are Leontief, and under the assumption that production functions are Cobb-Douglas.

5.1 Data

To estimate the effect of electronic miniaturization on productivity, I use the most recent detail-level BEA input-output table, corresponding to the year 2012, as well as a new dataset of products’ weights, measured in 2019. The weights are computed using IHS Markit’s PIERS database of Schedule B information for imports into the United States. This dataset contains shipment-level information on products, including 6-digit harmonized-system (HS) codes, price paid, volume (in Twenty-Foot Equivalent Units, or TEUs),¹⁶ weight (in pounds), units and quantities. I use a

¹⁵Appendix B describes the dynamic programming algorithm to compute the number of feasible combinations for every industry i at every time t .

¹⁶A TEU is the volume of a standard 20-foot cargo container, which is approximately 1172 cubic feet or 33 cubic meters. While I observe the size, measured in TEUs, I rely on weight observations

crosswalk to merge the datasets and compute the median weight for every tradeable BEA industry.¹⁷ Observations where the units are not boxes, packages or containers are dropped. The remaining observations represent 87.1% of the data.

Table 2 shows the top 5 and bottom 5 industries separately ranked by four metrics. Panel A shows industries ranked by volume, measured in TEUs. We can see that the largest products are trucks, buses, conveyors, and machine tools.¹⁸ The smallest products are parts of other transportation equipment, parts of dental equipment, parts of telephones, ammunition and lamps. Panel B shows the top and bottom 5 industries ranked by weight. Again, the heaviest products are trucks and buses, followed now by metal-forming machine tools and rolling mill machinery. The lightest products are missile and aerospace parts, phone parts, dental equipment parts, toys, and lamps. Panel C shows products ranked by density, with fabricated metal products and parts of transportation equipment and ammunition being the most dense, and missiles, space vehicles, trucks and aircraft being the least dense. Panel D shows the industries ranked by price per pound. Advanced industries such as the aerospace and X-ray apparatus industries are the most expensive per pound. The cheapest products per pound are structural metal products, ornamental metal work, and bolts, nuts, and screws.

I use this data to compute the number of possible input combinations $|\mathcal{F}(\theta, \tau_i)|$ for each industry in the machinery manufacturing and service sectors. I categorize industries that have BEA codes starting with 333 (Machinery), 335 (Electronic Equipment, Appliances and Components), 336 (Transportation Equipment), 337 (Furniture) and 339 (Miscellaneous Manufacturing) as industries that produce machines. Industries whose BEA code starts with 4 are either retail or transportation. Industries with an initial BEA digit of 5 provide services.¹⁹ Following the model,

because weight observations are less noisy since they are precisely measured with a scale, as opposed to estimated as fractions of a container.

¹⁷Tradeable BEA industries are those which have an HS counterpart in the PIERS data, and represent 191 out of 391 industries. Farming, Fishing, Oil and Gas Production, Mining and Manufacturing are included. Services are excluded and assigned a weight of 0.

¹⁸Except for truck and bus bodies, all products are smaller in volume than one 20-foot container. While it may seem that the volume of truck and bus bodies (12 TEUs) is too large, the average semi truck is around 8,262 cubic feet (72 feet long, 8.5 feet wide, 13.5 feet tall) or 7.05 TEUs which is in the ballpark of the data. The weight of an average semi truck (without any cargo) is around 35,000 pounds, which is also in the ballpark of the data, as Panel B shows. For these reasons, I do not consider the size and weight of truck and bus bodies to be an error in the data.

¹⁹It is important to note that this excludes computers and electronics manufacturing (BEA code 334). When mapping the data to the model in Section 2, it is helpful to think of computers

Table 2: Top and Bottom Industries by Physical Characteristics

Panel A: Manufacturing Industries Ranked by Volume (TEUs)			
Top 5 Industries		Bottom 5 Industries	
Truck and bus bodies	12.000	Other transportation equipment, nspf, and parts, nspf	0.001
Truck trailers	1.000	Dental equipment, supplies, and parts, nspf	0.001
Travel trailers and campers	0.670	Small arms ammunition, nspf	0.001
Machine tools, metal-forming, and parts, nspf	0.220	Electric lamps	0.002
Conveyors and conveying equipment, and parts, nspf	0.208	Telephone and telegraph apparatus, and parts, nspf	0.002
Panel B: Manufacturing Industries Ranked by Weight (Lbs)			
Top 5 Industries		Bottom 5 Industries	
Truck and bus bodies	31,378.010	Guided missiles and space vehicles, and parts, nspf	9.645
Truck trailers	7,607.040	Telephone and telegraph apparatus, and parts, nspf	10.907
Travel trailers and campers	3,179.060	Dolls and stuffed toy animals	11.352
Machine tools, metal-forming, and parts, nspf	2,601.450	Electric lamps	12.138
Rolling mill machinery, and parts, nspf	2,416.508	Dental equipment, supplies, and parts, nspf	13.553
Panel C: Manufacturing Industries Ranked by Density (Lbs/TEUs)			
Top 5 Industries		Bottom 5 Industries	
Other transportation equipment, nspf, and parts, nspf	36,788.140	Guided missiles and space vehicles, and parts, nspf	857.333
Fabricated plate work	28,377.130	Truck and bus bodies	2,614.834
Structural metal parts, nspf	28,208.570	Aircraft	3,306.925
Fabricated structural metal products, nspf	25,487.860	Aircraft equipment, nspf	3,748.971
Small arms ammunition, nspf	23,185.870	Travel trailers and campers	4,744.866
Panel D: Manufacturing Industries Ranked by Price (\$/Lbs)			
Top 5 Industries		Bottom 5 Industries	
Guided missiles and space vehicles, and parts, nspf	259.448	Structural metal parts, nspf	0.617
Aircraft	117.643	Architectural and ornamental metal work, nspf	0.797
Missile and rocket engines	90.623	Fabricated structural metal products, nspf	1.008
Aircraft equipment, nspf	68.309	Truck trailers	1.102
X-ray apparatus and tubes and related irradiation apparatus	38.773	Bolts, nuts, screws, rivets, and washers	1.367

Note: This table shows how different industries are ranked according to different physical features of their median product. Panel A ranks industries by volume. Panel B ranks industries by weight. Panel C ranks industries by density. Panel D ranks industries by Price-Per-Pound. The data is obtained from Schedule B reports of imports into the United States, which provide shipment level details on product quantities, weights, volumes, units and prices.

each machine-producing industry has a weight constraint of the form

$$\sum_{j \in S_i} \alpha_{ij} \theta_j \leq \tau_i. \quad (29)$$

I use industry j 's share of industry i 's revenue $\frac{P_j X_{ij}}{P_i Y_i}$, observed in the 2012 input-output matrix, as a proxy for α_{ij} . For each industry j available in the PIERS data, and each shipment σ , I observe the weight $\theta_{j\sigma}$ of the shipment. I compute θ_j by taking the median over all shipments σ for industry j .

For machine-producing industries, I compute the threshold τ_i from the dataset. Because thresholds represent the largest plausible weight a machine could take, I compute $\tau_i = \max_{\sigma} \theta_{i\sigma}$, so that the threshold weight is the maximum observed weight over all shipments. The service sector is not present in the trade data, so the threshold must be assigned in a different way. I choose $\tau_i = 30lbs$ for industries in the service sector, representing the maximum weight an adult service worker could carry comfortably in a knapsack throughout an entire day.

I use the NBER-CES dataset on manufacturing productivity to obtain a measure of Total Factor Productivity (TFP) for manufacturing industries. This dataset spans the years 1958-2018. For non-manufacturing industries, I use the BEA-BLS productivity dataset which spans the years 1987-2019.

5.2 Computing Weight Spillovers and the Number of Feasible Combinations

Weight Spillovers. To implement this simulation, I begin by taking the weight $\theta_s(2019)$ of semiconductors in 2019 as exogenous. In this simulation, the weight of semiconductors is an idiosyncratic factor, and in particular does not depend on the weight of any other manufactured products. To match the notation in Section 2, I write $\theta_s(t) = \zeta_s(t)$, where the ζ notation emphasizes that the weight of semiconduc-

and electronics as primary inputs which can change in weight, and machines as final goods which can take different combinations of primary inputs, but which have fixed thresholds. The reason for this distinction is that—while the possible inputs that go into a machine might change—the ultimate weight of the machine, and its corresponding knapsack threshold will not. As a concrete example, the weight of a hydraulic press have not varied substantially over time, while the weight of computers and electronics have changed dramatically. The effect of shrinking electronics on hydraulic presses has not been to decrease the weight of machines, but rather to improve their productivity by allowing electronics to be embedded into the machine.

tors is idiosyncratic.

For every year t between 1960 and 2019, I use Moore’s Law to extrapolate the weight $\zeta_s(t) = 2^{\frac{1}{1.5} \cdot (2019-t)} \zeta_s(2019)$ of semiconductors in year t going back to 1960. I denote the change in weight of semiconductors by $\Delta\zeta_s(t) = \zeta_s(t) - \zeta_s(2019) = (2^{\frac{1}{1.5} \cdot (2019-t)} - 1) \zeta_s(2019)$.

As a second step, I use the input-output matrix α to compute the weight spillover effects of semiconductors on computers and electronics for year t . Applying the formula

$$\Delta\theta(t) = (I - \alpha^{-1})\Delta\zeta$$

from Equation (4). To isolate the effect of semiconductor miniaturization on computer and electronics’ prices, I set $\Delta\zeta = (0, 0, \dots, 0, \Delta\theta_s(t), 0, \dots, 0)'$, so that the only idiosyncratic changes in weight are coming from semiconductor miniaturization.

However, applying Equation (4) directly would yield an overestimate of the effect of miniaturization for two reasons. First, this formula would imply that a change in weight of semiconductors would change the weight of every other product, not just computers and electronics. Second, this formula does not treat semiconductor weight as exogenous, as we assumed in the previous step. Indeed, there is a feedback loop through which smaller semiconductors would imply smaller machines, which would themselves lead to even smaller semiconductors beyond the shrinking dictated by Moore’s Law.

To avoid these pitfalls in the second step, I use the sub-matrix α_E whose rows and columns correspond only to electronic and computer manufacturing industries, including semiconductor manufacturing.²⁰ Furthermore, to ensure that the weight of semiconductors remains exogenous, I set the row of α_E corresponding to semiconductors to 0, so that no other products’ weight can influence the exogenously set weight of semiconductors. With this specification, I can compute the effect the spillover effect of semiconductor miniaturization on the weights of computers and electronics as

$$\Delta\theta_E(t) = (I - \alpha_E)^{-1}\Delta\zeta(t) \tag{30}$$

where $\theta_E(t)$ is a vector representing the weight of computers and electronics industries

²⁰Electronic and computer manufacturing industries are those whose BEA code starts with the 3-digit prefix 334. The semiconductor manufacturing industry has BEA code 334413.

at time t .

Computing the number of feasible combinations. To compute $|\mathcal{F}(\theta, \tau_i)|$, I use a dynamic programming algorithm described in Appendix B. The results are illustrated in Figure 4, which shows how Moore’s Law affected the number of feasible combinations for the average manufacturing and non-manufacturing industries. The figure clearly shows that manufacturing industries were early adopters: one can place more computers in heavy machinery, even if these computers are very large as they were in the 60s and 70s. The figure also captures the PC revolution of the 1980s, with an increasing adoption of new input combinations in non-manufacturing industries, peaking in the late 1990s and having a small bump in the early 2010s with the introduction of smartphones and tablets.

5.3 Estimating The Productivity Parameter $\frac{1}{\kappa}$

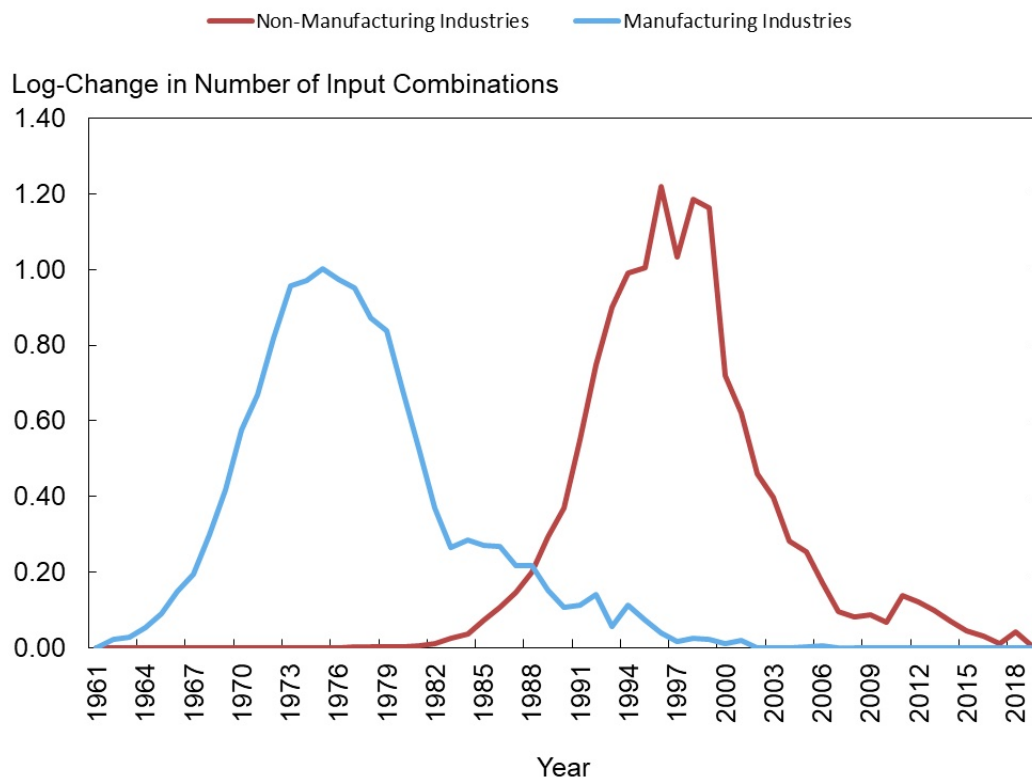
As we go back in time and the weight of semiconductors grows exponentially, some machines which are feasible to build in 2019—such as cars with assisted navigation—become infeasible. At the industry level, this means that machine-producing industries become less productive because they have fewer feasible input combinations that they can turn into products. Equation (21) tells us that the expected log-output of a machine-producing industry i is given by $\mathbb{E}[\log Y_i] = \frac{1}{\kappa} \log |\mathcal{F}(\theta(t), \tau_i)|$. The model in this paper abstracts away from capital, and normalizes labor to one unit, supplied inelastically. Thus, within the context of the model, Equation (21) measures *expected output-per-capita* as a function of the number of combinations. Since output-per-capita growth is driven by productivity growth, it is standard to use TFP as the left-hand-side variable in this equation. Thus, to compute $\frac{1}{\kappa}$, I estimate a regression of the form

$$\Delta \log A_{i,t} = \frac{1}{\kappa} \Delta \log |\mathcal{F}_{i,t}| + \beta \Delta \log A_{i,t-1} + \gamma_i + \delta_t + \epsilon_{i,t} \quad (31)$$

where $\frac{1}{\kappa} \Delta \log |\mathcal{F}_{i,t}|$ captures the effect of the change in the number of feasible combinations on productivity, $\beta \Delta \log A_{i,t-1}$ captures autoregressions in productivity, γ_i is an industry-specific fixed effect, and δ_t captures time trends in productivity growth.²¹

²¹It is important to note that γ_i captures industry-level differences in baseline growth rates, rather than industry-level differences in productivity levels. That is, an industry with higher γ_i

Figure 4: Average Log-Change in Number Of Feasible Input Combinations for Manufacturing and Non-Manufacturing Industries



Note: This figure illustrates how the number of feasible combinations changed over time for manufacturing and non-manufacturing industries. The blue line represents the average value of $\Delta \log |\mathcal{F}_{i,t}|$ for machine manufacturing industries, while the red line represents this value for non-manufacturing industries. We can see that the effect of Moore's Law on manufacturing peaked in the mid 1970s, while the effect on non-manufacturing industries peaked in the late 1990s.

Table 3 captures the summary statistics for the main variables $\Delta \log A_{i,t}$ and $\Delta \log |\mathcal{F}_{i,t}|$ in Equation (31). Table 4 shows that the Top 5 changes in $|\mathcal{F}_{i,t}|$ happened across varied service sector industries—such as Insurance providers, Newspaper Publishers, Entertainment, and Photography—in 1999 and 1998. This is consistent with these industries’ transition from analogue distribution methods, such as paper and film, to digital distribution methods, which was enabled by the exponential increase in hard-drive storage capabilities during the late 1990s.

Table 3: Summary statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
$\Delta \log \mathcal{F}_{i,t} $	0.235	0.4	0	1.903	10089
$\Delta \log A_{i,t}$	0.004	0.046	-0.399	0.543	6314

Note: This table shows summary statistics for the main variables in the regression equation (31).

Table 4: Top 5 changes in $|\mathcal{F}_{i,t}|$

Year	Industry	$\Delta \log \mathcal{F}_{i,t} $
1999	Direct Life Insurance Carriers	1.902954
1999	Motion Picture and Video Industries	1.848053
1999	Newspaper Publishers	1.733246
1999	Software Publishers	1.713959
1998	Photographic Services	1.683655

Note: This table shows the Top 5 Year-Industry pairs sorted by the change in the number of combinations due to Moore’s Law. We can see that the largest changes are concentrated in various service industries (Financial, Entertainment, Publishing) and happen in the late 1990s.

Table 5 shows the estimates for this regression under different specifications. The table has three panels. Panel A considers all industries. Panel B restricts the regression only to non-manufacturing industries, and Panel C restricts the regression to only machine manufacturing industries, excluding computers and electronics.²² Each Panel has 6 columns, representing different regression specifications. The first will have higher average productivity growth each year.

²²Here, machine manufacturing industries are those with a NAICS code starting with 33, except computers and electronics (those starting with 334). Computers and electronics are excluded because the threshold in Equation (5) changes over time, making the calculation of the number of combinations inapplicable to them as described above in this section. Non-manufacturing industries are those with NAICS codes starting with 4,5,6,7 or 8.

column does not have industry-level fixed effects for growth rates (so γ_i is omitted). The second column includes industry level fixed effects. The third and fourth column repeat these specifications for the period 1960-1989. The fifth and sixth column repeat these specifications for the period 1990-2019.²³

The results in the table show that the coefficient κ^{-1} is significantly positive, and hovers around $\kappa^{-1} \approx 0.004$ when estimated on all industries or manufacturing industries. This implies that a 1% increase in the number of input combinations yields an approximate 0.004% increase in industry productivity. The size of the coefficient is stronger for non-manufacturing industries, hovering around 0.01 for most specifications. For manufacturing industries, results are significant for the 1960-1990 period, but not after.

5.4 Computing the Effect of Moore’s Law on Aggregate TFP

The last step is to compute the change in aggregate TFP given the estimated changes $\Delta \log A_{i,t}$ in the productivity of machine-producing industries. One way to do this is to assume that all production functions in the economy, together with the household utility functions, are Cobb-Douglas and the Domar weights don’t change.²⁴ In this case, Hulten’s Theorem applies exactly, and we can compute

$$\Delta \log TFP_t = \sum_{i=1}^N D_{i,t} \Delta \log A_{i,t} \quad (32)$$

where D_i is industry i ’s Domar weight. Using this formula, I compute the effect of electronic miniaturization on annual TFP growth for each year in the 1960 to 2019 period, which is the blue line in Figure 5.

The assumption that production functions are Cobb-Douglas may not always apply, as discussed in Appendix C. While for short time horizons Hulten’s Theorem and Equation (32) are still good approximations,²⁵ for long horizons the quality of the approximation degrades. When production functions are arbitrary, one can provide

²³Here it is important to note that—for non-manufacturing industries—we only have data going back to 1987. Nevertheless, the results for the latter period are significant, even with a limited timespan.

²⁴For this calculation, I use the Domar weights from 2019. The implication of this simplifying assumption is that the consumption shares and the input-output matrix don’t change over time.

²⁵This is because the economy is efficient, as proved in Theorem 3.

Table 5: Regression Estimates for κ^{-1}

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: All Industries						
κ^{-1}	0.004*** (0.002)	0.004** (0.002)	0.006** (0.002)	0.004* (0.002)	0.004* (0.002)	0.004 (0.003)
Observations	6162	6162	1928	1928	4234	4234
Panel B: Non-Manufacturing Industries						
κ^{-1}	0.011** (0.005)	0.011** (0.005)	0.026** (0.011)	0.006 (0.037)	0.010* (0.005)	0.010* (0.006)
Observations	2914	2914	188	188	2726	2726
Panel C: Manufacturing Industries						
κ^{-1}	0.003 (0.002)	0.003 (0.002)	0.004* (0.002)	0.004* (0.002)	-0.017 (0.011)	-0.017 (0.011)
Observations	3248	3248	1740	1740	1508	1508
Period	1960-2019	1960-2019	1960-1989	1960-1989	1990-2019	1990-2019
Growth Fixed Effects	No	Yes	No	Yes	No	Yes

Note: This table shows different estimates of the coefficient $\frac{1}{\kappa}$ in Equation (31). Panel A considers all industries. Panel B restricts the regression to retail, transport and service industries, and Panel C restricts the regression only to manufacturing. Each Panel has 6 columns, representing different regression specifications. The first column does not have industry-level fixed effects for growth rates. The second column includes industry level fixed effects. The third and fourth column repeat these specifications for the period 1960-1989. The fifth and sixth column repeat these specifications for the period 1990-2019.

a lower bound on the estimate $\Delta \log TFP_t$ by ignoring the effect of productivity spillovers. That is, we can assume that a change in $\Delta \log A_{i,t}$ only affects consumers of industry i 's good, but not other industries. If utility functions are CES, and β_i is the share of consumer spending on industry i 's good, then we have

$$\Delta \log GDP \geq \sum_{i=1}^N \beta_i \Delta \log A_i. \quad (33)$$

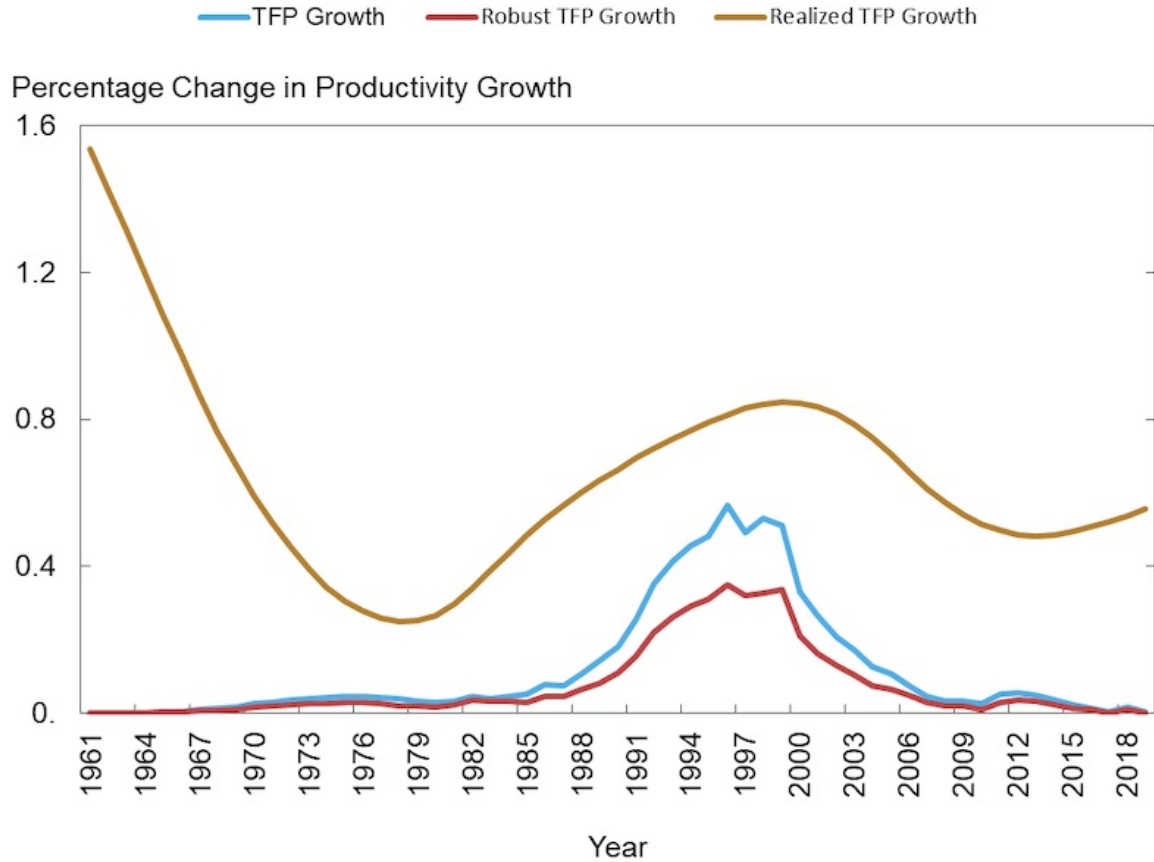
I compute this lower bound for each year in the 1960 to 2019 period, and plot it as a red line in Figure 5.

Figure 5 shows the simulated yearly changes in productivity obtained from electronic miniaturization. As mentioned above, the blue line uses the Cobb-Douglas specification, while the red line uses the robust specification. A yellow line shows the HP-filtered trend in realized TFP, obtained from FRED. We can see from the figure that the change in productivity is not uniform, but comes in waves, with small peaks in the late 1970s and early 2010s, and a large wave from the 1980s to the late 2000s. The largest effect is in the late 1990s, when computers and electronics were introduced in nearly every industry. The effects are computed using 0.004 as the estimate for κ^{-1} .

The annualized growth rate due to electronic miniaturization in the Cobb-Douglas specification is 0.12%, while the corresponding robust growth rate is 0.07%. If we focus on the years 1990-2019, when the IT revolution was in full force, we obtain that the annualized growth rate due to miniaturization for the Cobb-Douglas specification is 0.20%, while the robust specification yields an annualized 0.12% growth rate.

For comparison purposes, the annualized TFP growth rate during the whole 1960-2019 period—obtained from FRED—was 0.64%, and the annualized growth rate in the 1990-2019 period was 0.66%. Thus, in the Cobb-Douglas specification, electronic miniaturization accounts for 18.63% of all TFP growth in the entire period, and 30.03% of TFP growth in the 1990-2019 period. In the robust specification, electronic miniaturization accounts for 11.74% of all TFP growth in the entire period, and 18.63% of TFP growth in the 1990-2019 period.

Figure 5: Productivity Growth Attributable to Electronic Miniaturization by Year in Structural Simulation



Note: This figure illustrates the effects of electronic miniaturization on TFP growth over the 1960-2019 period. The blue line represents an estimate for the amount of growth attributable to miniaturization assuming that production functions are Cobb-Douglas. The red line presents a robust estimate that does not assume a functional form, but which shuts down all spillover effects from productivity. The yellow line represents the HP-Filtered realized TFP trend. We can observe from this figure that the effect of miniaturization comes in waves, small effects in the late 1970s and early 2010s, and a large wave from the 1980s to the late 2000s, which peaked in the late 1990s.

6 Conclusion

Over the past few decades, transistors have been getting exponentially smaller and cheaper over time, leading to significant changes across industries and countries. The effect of cheaper transistors on GDP has been widely studied. This paper is the first attempt to build a model where the miniaturization of electronic components leads to increases in aggregate productivity.

To study the effect of smaller electronics on GDP, I have introduced a new model that incorporates physical constraints into the profit-maximization problem of the firm. Even though the constraints are discrete, one can still derive tractable formulas for aggregate output, growth rates, and compute simulated counterfactuals.

My model is flexible enough to allow for arbitrary combinatorial constraints instead of the baseline weight constraint. This opens the door to more general models that capture physical constraints in production. In future work, I hope to explore these more general constraints from an empirical point of view.

I have left unexplored the competitive aspects of semiconductor innovation, and the races between firms to develop ever-shrinking transistors. In future work, I hope to develop the model further to include an oligopolistic market structure and an endogenous growth model where firms invest in research and compete to develop smaller varieties within an industry. Combining this model with micro-level data on semiconductor production would yield new insights on how overcoming physical constraints in production leads to economic growth.

References

- [1] Acemoglu, D., Autor, D. and Patterson C. (2017). Bottlenecks: Sectoral Imbalances and the US Productivity Slowdown. *Manuscript*.
- [2] Acemoglu D., and Azar ,P. (2020). Endogenous Production Networks. *Econometrica*, 88(1), 33-82.
- [3] Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5), 1977-2016

- [4] Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2017). Microeconomic origins of macroeconomic tail risks. *The American Economic Review*, 107(1), 54-108.
- [5] Acemoglu, D., and Tahbaz-Salehi, A. (2020). Firms, failures, and fluctuations: the macroeconomics of supply chain disruptions *National Bureau of Economic Research Working Paper*. (No. w27565).
- [6] Aghion, P., and Howitt, P. (1992). A model of Growth through Creative Destruction. *Econometrica*, 60(2), 323-351.
- [7] Alon, N. (1990). The number of spanning trees in regular graphs. *Random Structures & Algorithms*, 1: 175-181.
- [8] Antràs, P., and Chor, D. (2013). Organizing the global value chain. *Econometrica*, 81(6), 2127-2204.
- [9] Antràs, P. Fort T. C. and Tintelnot (2017). The Margins of Global Sourcing: Theory and Evidence from U.S. Firms, *American Economic Review*, 107(9), 2514-64.
- [10] Auerswald, P., Kauffman, S., Lobo, J., & Shell, K. (2000). The production recipes approach to modeling technological innovation: An application to learning by doing. *Journal of Economic Dynamics and Control*, 24(3), 389-450.
- [11] Azar P. D.(2021b) A Microfoundation for Cobb-Douglas Production Functions with Discrete Size Constraints. *Manuscript*
- [12] Baqaee, D. R. 2018: Cascading Failures in Production Networks, *Econometrica*, 86 (5), 1819–1838.
- [13] Baqaee, D. R., and Farhi, E. (2019a). The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem . *Econometrica*, 87 (4), 1155–1203.
- [14] Baqaee, D.R., and Farhi, E. (2019b). Productivity and Misallocation in General Equilibrium. *The Quarterly Journal of Economics*, 135 (1), 105–163.
- [15] Bartelme, D., and Gorodnichenko Y. (2015) Linkages and Economic Development. *National Bureau of Economic Research Working Paper No. 21251*.

- [16] Bigio, S., and La'O, J. (2020). Distortions in Production Networks. *The Quarterly Journal of Economics*, 135 (4), 2187–2253.
- [17] Boehm, J. and Oberfield, E. (2018). Misallocation in the Market for Inputs: Enforcement and the Organization of Production, NBER Working Paper 24937, August 2018.
- [18] Caliendo, L., Parro, F., and Tsyvinski, A. (2017). Distortions and the Structure of the World Economy. *National Bureau of Economic Research Working Paper No. w23332*.
- [19] Chaney T. (2014). The Network Structure of International Trade. *American Economic Review*, 104(11), 3600-3634.
- [20] Eaton, J., and Kortum, S. (2001). Technology, trade, and growth: A unified framework. *European economic review*, 45(4), 742-755.
- [21] Elliott M., Golub B. and Leduc M.V. (2021). Supply Network Formation and Fragility. *Manuscript*
- [22] Fadinger, H., Ghiglino C. and Teteryatnikova M. (2016) Income differences and input-output structure. *CEPR Discussion Paper No. 11547*.
- [23] Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica*, 79(3), 733-772.
- [24] Ghiglino, C. (2012). Random walk to innovation: Why productivity follows a power law. *Journal of Economic Theory*, 147(2), 713-737.
- [25] Grossman, G. M., and Helpman, E. (1991). Quality ladders in the theory of growth. *The Review of Economic Studies*, 58(1), 43-61.
- [26] Grossman, G. M., and Helpman, E. (1992). Innovation and growth in the global economy. *MIT press*.
- [27] Hulten, C. R. (1978). Growth accounting with intermediate inputs. *The Review of Economic Studies*, 511–518.
- [28] Jones, C. I. (1995). R & D-based models of economic growth. *Journal of political Economy*, 103(4), 759-784.

- [29] Jones, C. I. (2005). The shape of production functions and the direction of technical change. *The Quarterly Journal of Economics* 120(2), 517-549.
- [30] Jones, C. I. (2011). Intermediate goods and weak links in the theory of economic development. *American Economic Journal: Macroeconomics* 3(2), 1-28.
- [31] Jones, C. I. (2021). Recipes and Economic Growth: A Combinatorial March Down an Exponential Tail. *National Bureau of Economic Research Working Paper (No. w28340)*.
- [32] Jorgenson, D.W. (2005). Accounting for Growth in the Information Age. *Handbook of Economic Growth* 1 (A), 743-815.
- [33] Jorgenson, D.W., Stiroh, K.J. (1999). Information technology and growth. *American Economic Review* 89 (2), 109–115.
- [34] Jorgenson, D.W., Stiroh, K.J. (2000). U.S. economic growth at the industry level. *American Economic Review* 90 (2), 161–167.
- [35] Kennan, J. (2001). Uniqueness of Positive Fixed Points for Increasing Concave Functions on \mathbb{R}^n : An Elementary Result. *Review of Economic Dynamics*, 4(4), 893–899.
- [36] Lim, K. (2018). Endogenous Production Networks and the Business Cycle. *Manuscript*.
- [37] Liu, E. (2019). Industrial policies and economic development. *The Quarterly Journal of Economics* 139(4).
- [38] Liu, E. and Tsyvinski, A. (2020). Dynamical Structure and Spectral Properties of Input-Output Networks *National Bureau of Economic Research Working Paper (No. w28178)*.
- [39] Long, J. B., and Plosser, C. I. (1983). Real business cycles. *Journal of political Economy*, 91(1), 39-69. Chicago
- [40] Olea, J.L.M. and Pflueger, C. (2013). A robust test for weak instruments. *Journal of Business & Economic Statistics*, 31(3), pp.358-369.

- [41] Oberfield, E. (2018). A theory of input–output architecture. *Econometrica*, 86(2), 559-589.
- [42] Romer, P. M. (1990). Endogenous technological change. *Journal of political Economy*, 98(5, Part 2), S71-S102.
- [43] Simonovska, I., and Waugh, M. E. (2014). The elasticity of trade: Estimates and evidence. *Journal of international Economics*, 92(1), 34-50.
- [44] Taschereau-Dumouchel, M. (2020). Cascades and fluctuations in an economy with an endogenous production network. *Manuscript*.
- [45] Weitzman, M. L. (1998). Recombinant growth. *The Quarterly Journal of Economics*, 113(2), 331-360.

Online Appendix

A Proofs from Section 3

A.1 Proof of Theorems 1, 2 and 3

Proof of Theorem 1. It is immediate that Equation (10) is necessary for an equilibrium to exist. To prove sufficiency, assume that we are given prices P_1, \dots, P_N which satisfy Equation (10). The rest of the proof shows how to construct a feasible allocation that will satisfy all the equilibrium conditions.

Final Industry Price Because markets are competitive, the final industry price is given by its marginal cost $P_f = K_f(P_1, \dots, P_N)$.

Final Output and Consumption Since the utility function of the household is strictly increasing, the budget constraint binds, and $P_f C = 1$. If the equilibrium exists, the final good market clearing condition implies that $Y = C = \frac{1}{P_f}$.

Final Industry Demands Let $S \in \mathcal{F}$ be a set that minimizes the final industry's marginal cost.²⁶ Since the final industry's production function is strictly quasi-concave in $(X_{fi})_{i \in S}$, the final industry's demands for primary industry products are completely determined given Y_f and S .

Primary Industry Output and Demands To determine the output of primary industries, write primary good i 's market clearing condition as

$$P_i Y_i = P_i X_{fi} + \sum_{j=1}^N P_i X_{ji}.$$

Multiply and divide the term $P_i X_{ji}$ by $P_j Y_j$, so that the market clearing condition becomes $P_i Y_i = P_i X_{fi} + \sum_{j=1}^N \frac{P_i X_{ji}}{P_j Y_j} P_j Y_j$. Let $\hat{Y}_i = P_i Y_i$, $\hat{X}_i = P_i X_{fi}$, and $\hat{\alpha}_{ij} = \frac{P_i X_{ji}}{P_j Y_j}$.

²⁶Since \mathcal{F} is finite, such a set always exists.

Since production functions have constant returns to scale, the ratio $\frac{X_{ji}}{Y_j}$ is completely determined by primary good firms' cost-minimization problem. The revenue share $\hat{\alpha}_{ij}$ can be computed from this ratio and the given vector of prices.

Write the market clearing condition in matrix form as $\hat{Y} = \hat{X} + \hat{\alpha}\hat{Y}$. Since labor is an essential input to every primary industry, industry j 's payments to all of its non-labor suppliers $\sum_{i=1}^N P_i X_{ji}$ are strictly less than its total revenue $P_j Y_j$. This implies that the matrix of revenue shares $\hat{\alpha}$ satisfies $\sum_{i=1}^N \hat{\alpha}_{ij} < 1$, which itself implies that $(I - \hat{\alpha})^{-1}$ exists. Therefore, the vector of primary industry revenues is given by

$$\hat{Y} = (I - \hat{\alpha})^{-1} \hat{X}.$$

Checking the Equilibrium Conditions The constructed allocation satisfies the primary and final good market clearing conditions by construction. In addition, the sets S and the demands for primary inputs X are chosen to minimize firms' marginal costs. Let $\frac{L_i}{Y_i}$ be the per-unit labor used to produce one unit of good i . This ratio follows—like input demands—from primary industries' cost minimization problem. Finally, set $L_i = \frac{L_i}{Y_i} \cdot Y_i$, and note that the labor market clearing condition $\sum_{i=1}^N L_i = 1$ follows from Walras' Law. Thus, all market clearing, firm optimization and household optimization conditions are satisfied, and the constructed allocation is an equilibrium. ■

Proof of Proposition 1. The primary price vector, if it exists, satisfies the system of equations

$$P_j = \sum_{k=1}^K \alpha_{jk} P_k + \frac{1}{A_j}.$$

Letting $B = (\frac{1}{A_1}, \dots, \frac{1}{A_J})'$, write this in matrix form as

$$(I - \alpha)P = B.$$

Since $(I - \alpha)^{-1}$ exists and all of its coefficients are non-negative, the unique price vector is given by $P = (I - \alpha)^{-1}B$. ■

Proof of Proposition 1. Recall, from the proof of Theorem 1, that if there exists a vector of primary prices (P_1, \dots, P_N) which satisfies the equations $P_i = K_i(P_1, \dots, P_N)$,

then there exists a unique equilibrium price vector P^* such that $P_i^* = P_i$. Thus, it suffices to show that if $K_i(\cdot)$ is strictly concave, then there exists at most one vector of primary prices that satisfies equation (10).

I use the following Lemma from Kennan (1999):

Lemma 2 (Kennan). *If $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is an increasing, strictly concave function such that*

1. $f_i(0, \dots, 0) \geq 0$,
2. *there exists $a = (a_1, \dots, a_n)$ such that $a_i > 0$ and $f_i(a) > a_i$ for all $i \in \{1, \dots, n\}$, and*
3. *there exists $b = (b_1, \dots, b_n)$ such that $b_i > a_i$ and $f_i(b) < b_i$ for all $i \in \{1, \dots, n\}$,*

then there exists a unique positive vector x such that $f(x) = x$.

In the statement of Lemma 2, define $f : \mathbb{R}^N \rightarrow \mathbb{R}^N$ such that $f_i(P) = K_i(P)$. Because equation (10) holds, f has a fixed point. I use Lemma 2 to prove that this fixed point is unique. I proceed to show that all the conditions of the Lemma hold. First, note that, by assumption f is strictly concave.

I now show that $f(0) > 0$. Recall that $K_i(P) = \min_{X_i, L_i} \sum_{j=1}^{P_j} X_{ij} + L_i$ subject to $F_i(X_i, L_i) = 1$, and that labor is an essential factor of production, so that the optimal L_i must be greater than 0. This means that $K_i(0) > 0$, since—even if prices were zero—the wage would still be equal to 1 (since labor is the numeraire), and a non-zero quantity of labor will be demanded. Thus, the condition $f_i(0, \dots, 0)$ in Lemma 2 is satisfied.

I now show that there exists $a > 0$ such that $f(a) > a$. This follows from the continuity of f . Let $\{a(n)\}_{n=1}^{\infty}$ be a sequence of positive vectors $a(n)$ such that $\lim_{n \rightarrow \infty} a(n) = 0$. Since f is continuous, $\lim_{n \rightarrow \infty} f(a(n)) = f(0) > 0 = \lim_{n \rightarrow \infty} a(n)$. Thus, for n large enough, $f(a(n)) > a(n)$.

Finally, I show that there exists $b > a$ such that $f(b) < b$. Let P be a fixed point of f (which exists by assumption). Then $f(P) = P$. Since f is strictly concave, $f(2P) < \frac{1}{2}f(P) + \frac{1}{2}f(0) = f(P)$. Thus, by setting $b = 2P$, the last condition of Lemma 2 is satisfied.

I conclude that f has a unique fixed point. Therefore, there exists only one vector P_1^*, \dots, P_N^* of primary prices such that equation (10) is satisfied, and a unique

equilibrium price vector P^* . ■

Proof of Theorem 2. Let $\mathcal{A} = \{A \in \mathbb{R}^{|\mathcal{F}|} : |\arg \min_S K(S, P^*, A)| \geq 2\}$ be the set of productivity parameters which admit more than one cost-minimizing set for the final industry. Let S^*, S^{**} be two different possible input sets, and define $\mathcal{A}(S^*, S^{**}) = \{A : K(S^*, P^*, A(S^*)) = K(S^{**}, P^*, A(S^{**}))\}$. The following set inclusion holds

$$\mathcal{A} \subset \bigcup_{S^*, S^{**}} \mathcal{A}_i(S^*, S^{**}).$$

Since the right-hand side of this set inclusion is a countable union of sets, it suffices to show that each $\mathcal{A}(S^*, S^{**})$ has measure zero for each pair S^*, S^{**} . For any such pair, define the function

$$\Delta(A, S^*, S^{**}) = K(S^*, P^*, A(S^*)) - K(S^{**}, P^*, A(S^{**}))$$

and note that $\mathcal{A}(S^*, S^{**})$ is exactly the set of productivity parameters A for which $\Delta(A, S^*, S^{**}) = 0$. Since K is strictly increasing in A , we have $A \in \mathcal{A}(S^*, S^{**})$ if and only if $A(S^*) = A(S^{**})$. But this implies that $\mathcal{A}(S^*, S^{**})$ is an $|\mathcal{F}| - 1$ dimensional subset of a $|\mathcal{F}|$ dimensional space, so it must have measure zero. Since \mathcal{A} is a subset of a countable union of measure zero sets, it also must have measure zero.

The above argument shows that the set of cost-minimizing set of inputs chosen by the final industry is generically unique. Given this unique set S^* , and the fact that the final industry's production function $Y_f(S^*) = A(S^*)F(S^*, (X_{fi})_{i \in S^*})$ is strictly quasi-concave as a function of $(X_{fi})_{i \in S^*}$ implies that the final industry's demands for primary inputs are uniquely determined. Finally, given $(X_{fi})_{i \in S^*}$ and the fact that primary industry production functions are strictly quasi-concave, we conclude that primary industry demands $(X_{ij})_{i,j=1}^N, (L_i)_{i=1}^N$ for industry and labor are uniquely determined. Thus, the equilibrium is generically unique. ■

Proof of Theorem 3. To start the proof, assume that the set S of final industry inputs is fixed. Given S , welfare is given by $U(S) = \max_{C,X,Y,L} U(C)$ subject to the the feasibility constraints ([Primary Market Clearing](#), [Final Market Clearing](#), [Labor Market Clearing](#)). Let $\chi_i(S), \chi_f(S), \chi_L(S)$ be the respective Lagrange multipliers for the primary good, final good and labor feasibility constraints. Because

production and utility functions are differentiable, the social planner's problem has the following first-order conditions:

$$\frac{\partial F_i}{\partial X_{ij}} = \frac{\chi_j(S)}{\chi_i(S)}, \quad \frac{\partial F_i}{\partial L_i} = \frac{\chi_L(S)}{\chi_i(S)}, \quad \frac{\partial F}{\partial X_{fi}} = \frac{\chi_i(S)}{\chi_f(S)} \quad \text{and} \quad \frac{dU}{dC} = \chi_f(S).$$

Given S , one can ask what the equilibrium allocation $(C(S), X(S), Y(S), L(S))$ and prices $P(S)$ would be if final industry firms took S as their fixed set of suppliers. This would be the same as the welfare-maximizing allocation taking S as given. To see this, set $\chi_i(S) = P_i(S)$, $\chi_f(S) = P_f(S)$ and $\chi_L(S) = W(S)$. Then the first-order conditions

$$\frac{\partial F_i}{\partial X_{ij}} = \frac{P_j(S)}{P_i(S)}, \quad \frac{\partial F_i}{\partial L_i} = \frac{W(S)}{P_i(S)}, \quad \frac{\partial F}{\partial X_{fi}} = \frac{P_i(S)}{P_f(S)} \quad \text{and} \quad \frac{dU}{dC} = P_f(S).$$

match the first-order conditions of the firm's problems and the household's problem in equilibrium. Thus, if we assign the Lagrange multipliers to match equilibrium prices and take the equilibrium input sets S as given, any corresponding equilibrium allocation $(C(S), X(S), Y(S), L(S))$ will maximize welfare.

Finally, note that since firms in the final industry choose S^* to minimize $P_f(S)$, and the utility is strictly concave, we have that

$$\frac{dU}{dC}|_{C=C(S^*)} = P_f(S^*) \leq P_f(S) = \frac{dU}{dC}|_{C=C(S)} \implies U(C(S)) \leq U(C(S^*)).$$

The market clearing condition for the final good yields $C(S) = \frac{W(S)}{P_f(S)}$. Using the above equations, this implies that $C(S) = \frac{\partial F_i}{\partial L_i} \cdot \frac{\partial F}{\partial X_{fi}} = \frac{\chi_L(S)}{\chi_f(S)}$. ■

A.2 Proof of Theorem 4

I begin by recalling some properties of the Frechet distribution. I use these properties to solve for equilibrium quantities and prices.

Properties of the Frechet Distribution In this subsection, I recall the definition of a Frechet distribution's shape and scale parameters and show how these parameters

change when under multiplication by a constant, exponentiation, and the maximum operator.

Definition 6. Given a scale parameter $s > 0$ and a shape parameter $\kappa > 0$, a Frechet distribution with scale s and shape κ is given by the CDF $\Psi(x) = e^{-\left(\frac{x}{s}\right)^{-\kappa}}$.

Lemma 3. Let $\gamma, c > 0$, and let X be a random variable drawn from a Frechet distribution scale parameter s and shape parameter κ . Then cX^γ is a random variable drawn from a Frechet distribution with scale parameter cs^γ and shape parameter $\frac{\kappa}{\gamma}$.

Proof.

$$\Phi(x) = \text{Prob}[cX^\gamma \leq x] = \text{Prob}\left[X \leq \left(\frac{x}{c}\right)^{\frac{1}{\gamma}}\right] = \Psi\left(\left(\frac{x}{c}\right)^{\frac{1}{\gamma}}\right) = e^{-\left(\frac{\left(\frac{x}{c}\right)^{\frac{1}{\gamma}}}{s}\right)^{-\kappa}} = e^{-\left(\frac{x}{cs^\gamma}\right)^{-\frac{\kappa}{\gamma}}}$$

■

Lemma 4. Let X_1, \dots, X_n be Frechet random variables drawn from independent distributions with shape parameter κ and scale parameters s_1, \dots, s_n , respectively. Then $\max(X_1, \dots, X_n)$ is drawn from a Frechet distribution with scale parameter $\left(\sum_{i=1}^n s_i^\kappa\right)^{\frac{1}{\kappa}}$ and shape parameter κ .

Proof. Let $\Psi_i(x) = e^{-\left(\frac{x}{s_i}\right)^{-\kappa}}$ be the CDF of X_i . Then

$$\text{Prob}[\max(X_1, \dots, X_n) \leq x] = \prod_{i=1}^n \Psi_i(x) = \prod_{i=1}^n e^{-\left(\frac{x}{s_i}\right)^{-\kappa}} = e^{-\left(x\left(\sum_{i=1}^n s_i^\kappa\right)^{-\frac{1}{\kappa}}\right)^{-\kappa}}$$

■

Lemma 5. If U is a standard exponential distribution with CDF $\Phi(x) = 1 - e^{-x}$ and $X = sU^{-\frac{1}{\kappa}}$, then X is a Frechet distribution with scale parameter s and shape parameter κ .

Proof.

$$\text{Prob}[X \leq x] = \text{Prob}[sU^{-\frac{1}{\kappa}} \leq x] = \text{Prob}\left[U \geq \left(\frac{x}{s}\right)^{-\kappa}\right] = e^{-\left(\frac{x}{s}\right)^{-\kappa}}$$

where the sign change in $\text{Prob}[sU^{-\frac{1}{\kappa}} \leq x] = \text{Prob}\left[U \geq \left(\frac{x}{s}\right)^{-\kappa}\right]$ is justified because $f(x) = \left(\frac{x}{s}\right)^{-\kappa}$ is a decreasing function. ■

Proof of Theorem 4. Since the final good market is competitive, firms will choose the set of inputs S which minimizes the cost function $K(S, P, \tau) = \frac{1}{\phi(S)} \bar{K}_f(S, P, \tau)$. This is equivalent to maximizing $\frac{1}{K(S, P, \tau)} = \phi(S) \bar{K}_f(S, P, \tau)^{-1}$. By Assumption 1, $\phi(S)$ is a Frechet random variable with scale parameter 1 and shape parameter κ . By Lemma 3, $\frac{1}{K(S, P, \tau)}$ is a Frechet random variable with shape parameter κ and scale parameter $\bar{K}_f(S, P, \tau)^{-1}$. By Lemma 4, $\max_{S \in \mathcal{F}(\theta, \tau)} \frac{1}{K(S, P, \tau)}$ is a Frechet random variable with shape parameter κ and scale parameter $(\sum_{S \in \mathcal{F}} \bar{K}_f(S, P, \tau)^{-\kappa})^{\frac{1}{\kappa}}$. Since markets are competitive, $P_f = (\max_{S \in \mathcal{F}(\theta, \tau)} \frac{1}{K(S, P, \tau)})^{-1} = (\sum_{S \in \mathcal{F}} \bar{K}_f(S, P, \tau)^{-\kappa})^{-\frac{1}{\kappa}}$. We can solve for Y_f from the household's budget constraint $P_f Y_f = W = 1$, so that $Y_f = \frac{1}{P_f}$, which is a Frechet random variable with shape parameter κ and scale parameter $(\sum_{S \in \mathcal{F}} \bar{K}_f(S, P, \tau)^{-\kappa})^{\frac{1}{\kappa}}$. ■

A.3 Proofs from Subsection 3.3

Proof of Lemma 1. Let $H(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases}$ and recall that $\delta(x) = \frac{dH}{dx}$. Write

$$\sum_{S \in \mathcal{F}(\theta, \tau)} \bar{K}(S, P, \tau)^{-\kappa} = \sum_{S \subset \mathcal{N}} \bar{K}(S, P, \tau)^{-\kappa} H(\tau - G(\theta, S)).$$

Using the chain rule and linearity of differentiation, one obtains the desired result

$$\frac{\partial}{\partial \theta_i} \sum_{S \in \mathcal{F}(\theta, \tau)} \bar{K}(S, P, \tau)^{-\kappa} = - \sum_{S \subset \mathcal{N}} \delta(\tau - G(\theta, S)) \bar{K}(S, P, \tau)^{-\kappa} \frac{\partial G}{\partial \theta_i}.$$

■

Proof of Proposition 3. This follows from applying the chain rule to equation (14) and applying Lemma 1. ■

Proof of Proposition 4. This follows from Proposition 3 and the definition of the Delta function. ■

A.4 Proofs From Section 4

Proof of Proposition 5. Let $N(t) = |\mathcal{N}(t)|$ be the number of useful primary goods. The cost function for industry i is

$$K_i(P) = \left((1 - \sum_{j=1}^{N(t)} \alpha_{ij})^\sigma + \sum_{j=1}^{N(t)} \alpha_{ij}^\sigma P_j^{1-\sigma} \right)^{\frac{1}{1-\sigma}}.$$

Since markets are competitive, prices are equal to marginal costs, so that $P_i^* = \left((1 - \sum_{j=1}^{N(t)} \alpha_{ij})^\sigma + \sum_{j=1}^{N(t)} \alpha_{ij}^\sigma P_j^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$. It is convenient to raise both sides in the previous equation to the $1-\sigma$ power to obtain the following system of linear equations in $Q^*(t) = ((P_1^*)^{1-\sigma}, \dots, (P_{N(t)}^*)^{1-\sigma})$:

$$Q_i^* = \left((1 - \sum_{j=1}^{N(t)} \alpha_{ij})^\sigma + \sum_{j=1}^{N(t)} \alpha_{ij}^\sigma Q_j^* \right)^{\frac{1}{1-\sigma}}.$$

The solution to this set of equations can be written as

$$Q^*(t) = (I - \mathcal{A}(t))^{-1} \mathcal{B}(t)$$

where $\mathcal{A}(t)$ is an $N(t) \times N(t)$ whose $(i, j)^{th}$ entry is α_{ij}^σ , and $\mathcal{B}(t)$ is an $N(t) \times 1$ vector whose i^{th} entry is $(1 - \sum_{j=1}^{N(t)} \alpha_{ij})^\sigma$.

We now use the fact that $\sum_{j=1}^{\infty} \alpha_{ij}^\sigma < \chi < 1$. The first implication of this fact is that $\|\mathcal{A}(t)\|_\infty =_{def} \max_{1 \leq i \leq N(t)} \sum_{j=1}^{N(t)} |\mathcal{A}_{ij}(t)| < \chi < 1$. The second implication is that the Leontief inverse $\mathcal{L}(t) = (I - \mathcal{A}(t))^{-1}$ always exists, and satisfies $\|\mathcal{L}(t)\|_\infty \leq \sum_{p=0}^{\infty} \|\mathcal{A}(t)\|_\infty^p \leq \frac{1}{1-\chi}$. The third implication is that the entries in the vector $\mathcal{B}(t)$ are uniformly bounded below by $1 - \chi$, since

$$\mathcal{B}_i(t) = \left(1 - \sum_{j=1}^{N(t)} \alpha_{ij} \right)^\sigma > 1 - \sum_{j=1}^{N(t)} \alpha_{ij} > \left(1 - \sum_{j=1}^{N(t)} \alpha_{ij}^\sigma \right) \geq 1 - \chi.$$

The first fact yields $Q^*(t) = (I - \mathcal{A}(t))^{-1}\mathcal{B}(t) > I \begin{pmatrix} 1 - \chi \\ 1 - \chi \\ \dots \\ 1 - \chi \\ \dots \\ 1 - \chi \end{pmatrix} = \begin{pmatrix} 1 - \chi \\ 1 - \chi \\ \dots \\ 1 - \chi \\ \dots \\ 1 - \chi \end{pmatrix}$. The

second fact yields

$$\|Q^*(t)\|_\infty \leq \frac{\chi}{1 - \chi}.$$

So far, I have shown that the equilibrium vector $Q^*(t) = ((P_1^*)^{1-\sigma}, \dots, (P_{N(t)}^*)^{1-\sigma})$ is bounded above and below by constants. Raising all components to the $\frac{1}{1-\sigma}$ power, we obtain the uniform bound on the equilibrium price vector:

$$(1 - \chi)^{\frac{1}{1-\sigma}} \leq P_i^*(t) \leq \left(\frac{\chi}{1 - \chi}\right)^{\frac{1}{1-\sigma}} \text{ for all } i, t. \quad (34)$$

Since the final good has a CES production function with elasticity of substitution ρ , the corresponding deterministic cost function is

$$\bar{K}(S, P(t)) = \frac{1}{\bar{A}(S)} \left(\sum_{j \in S} P_j^{1-\rho} \right)^{\frac{1}{1-\rho}} \quad (35)$$

Combining equations (34) and (35) yields the upper and lower bounds

$$\frac{|S|}{\bar{A}(S)} (1 - \chi)^{\frac{1}{1-\sigma}} \leq \bar{K}(S, P(t)) \leq \frac{|S|}{\bar{A}(S)} \left(\frac{\chi}{1 - \chi}\right)^{\frac{1}{1-\sigma}}. \quad (36)$$

Let $K_\ell(t) = \min_{S \in \mathcal{F}(t)} \frac{|S|}{\bar{A}(S)} (1 - \chi)^{\frac{1}{1-\sigma}}$, and $K_u(t) = \max_{S \in \mathcal{F}(t)} \frac{|S|}{\bar{A}(S)} \left(\frac{\chi}{1 - \chi}\right)^{\frac{1}{1-\sigma}}$. Taking logarithms and limits as $t \rightarrow \infty$, one obtains the following limit for the lower bound $K_\ell(t)$:

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{\log K_\ell(t)}{t} &= \lim_{t \rightarrow \infty} \frac{\min_{S \in \mathcal{F}(t)} \log\left(\frac{|S|}{\bar{A}(S)}\right)}{t} + \frac{\log(1 - \chi)^{\frac{1}{1-\sigma}}}{t} = \\ &= \lim_{t \rightarrow \infty} \frac{\min_{S \in \mathcal{F}(t)} \log\left(\frac{|S|}{\bar{A}(S)}\right)}{t}. \end{aligned} \quad (37)$$

Analogously, one can obtain the following limit for the upper bound $K_u(t)$:

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{\log K_u(t)}{t} &= \lim_{t \rightarrow \infty} \frac{\max_{S \in \mathcal{F}(t)} \log\left(\frac{|S|}{A(S)}\right)}{t} + \frac{\left(\frac{\chi}{1-\chi}\right)^{\frac{1}{1-\sigma}}}{t} = \\ &= \lim_{t \rightarrow \infty} \frac{\max_{S \in \mathcal{F}(t)} \log\left(\frac{|S|}{A(S)}\right)}{t}. \end{aligned} \quad (38)$$

To show that both of these limits go to 0, one can use the fact that $t^{-\nu} \leq \frac{|S|}{A(S)} \leq t^\nu$ for all $S \in \mathcal{F}(t)$. This yields:

$$\lim_{t \rightarrow \infty} \frac{-\nu \log t}{t} \leq \lim_{t \rightarrow \infty} \frac{\min_{S \in \mathcal{F}(t)} \log\left(\frac{|S|}{A(S)}\right)}{t} \leq \lim_{t \rightarrow \infty} \frac{\max_{S \in \mathcal{F}(t)} \log\left(\frac{|S|}{A(S)}\right)}{t} = \lim_{t \rightarrow \infty} \frac{\nu \log t}{t}.$$

Since $\lim_{t \rightarrow \infty} \frac{-\nu \log t}{t} = \lim_{t \rightarrow \infty} \frac{\nu \log t}{t} = 0$, the squeeze theorem implies $\lim_{t \rightarrow \infty} \frac{\log K_\ell(t)}{t} = \lim_{t \rightarrow \infty} \frac{\log K_u(t)}{t} = 0$. ■

Proof of Theorem 6. Write expected log-output at time t as $\mathbb{E}[\log Y_f(t)] = \frac{1}{\kappa} \log\left(\sum_{S \in \mathcal{F}(t)} \bar{K}(S, P, \tau)^{-\kappa}\right) + \frac{\gamma}{\kappa}$. Using Assumption 5, we can give upper and lower bounds for expected output:

$$\frac{1}{\kappa}(\log |\mathcal{F}(t)| + \log K_\ell(t)) + \frac{\gamma}{\kappa} \leq \mathbb{E}[\log Y_f(t)] \leq \frac{1}{\kappa}(\log |\mathcal{F}(t)| + \log K_u(t)) + \frac{\gamma}{\kappa}.$$

Dividing by t , taking limits, and using the fact that $\lim_{t \rightarrow \infty} \frac{\log K_\ell(t)}{t} = \lim_{t \rightarrow \infty} \frac{\log K_u(t)}{t} = 0$ yields the desired result

$$\lim_{t \rightarrow \infty} \frac{\mathbb{E}[\log Y_f(t)]}{t} = \frac{D}{\kappa}.$$

■

Proof of Theorem 7. From Theorem 6, we can write $g^* = \lim_{t \rightarrow \infty} \frac{\log |\mathcal{F}(t)|}{t}$. Since $\theta_i \geq \theta_\ell$ for all i , each feasible set can combine at most $q = \lceil \frac{\tau}{\theta_{lower}} \rceil$ inputs. Thus $|\mathcal{F}(t)| \leq \sum_{k=0}^q \binom{N(t)}{k}$. If $N(t) < \frac{q}{2}$ for all t , then $|\mathcal{F}(t)|$ is bounded for all t and $\lim_{t \rightarrow \infty} \frac{\log |\mathcal{F}(t)|}{t} = 0$. Thus, we can focus on the case where $N(t) > \frac{q}{2}$ for t large enough. In this case, the binomial coefficients are increasing in q . This yields an upper bound $\sum_{k=0}^q \binom{N(t)}{k} \leq q \binom{N(t)}{q}$. The binomial coefficient $\binom{N(t)}{q}$ is upper-bounded by the formula $\left(\frac{e^{N(t)}}{q}\right)^q$. Combining this series of upper bounds yields

$$\log |\mathcal{F}(t)| \leq q(1 + \log N(t) - q).$$

Since $N(t) \leq t^\nu$, $\log N(t) < \nu \log t$. Dividing by t and taking limits as t goes to infinity, we obtain $g^* = 0$. ■

B Algorithms for Computing the Number of Feasible Combinations with Size Constraints

In this Appendix, I give a recursive formula to compute $|\mathcal{F}(\theta, \tau)| = |\{S : \sum_{j \in S} \theta_j \leq \tau\}|$, the number of feasible combinations under a size constraint. Without loss of generality, I assume that all θ_j and τ are rational numbers with p digits of precision. Multiplying all of them by 10^p , I can further assume that they are all integers in $\{0, 1, \dots, M\}$ for some large integer M .

For any $j \in \mathcal{J}$ and $m \in \{0, \dots, M\}$, define $\mathcal{C}(j, m) = |\{S \subset \{1, \dots, j\} : \sum_{j \in S} \theta_j \leq m\}|$ to be the number of feasible sets which only contain the first j industries and whose total size is less than or equal to m . For any $j \in \mathcal{J}$, let $N_0(j) = |\{j' : 1 \leq j' \leq j \text{ and } \theta_{j'} = 0\}|$ be the number of industries whose index is less than equal to j and which have size $\theta_{j'} = 0$. Note that $\mathcal{C}(J, M) = \mathcal{C}$ is the quantity that we want to compute, and that $\mathcal{C}(j, 0) = 2^{N_0(j)}$ for any $j \in \mathcal{J}$. Furthermore, note that $\mathcal{C}(1, m) = 2$ if $\theta_1 \leq m$ and $\mathcal{C}(1, m) = 1$ if $\theta_1 > m$. Using $\{\mathcal{C}(j, 0), \mathcal{C}(1, m)\}_{j \leq J, m \leq M}$ as the base cases, one can use the recursive formula

$$\mathcal{C}(j, m) = \mathcal{C}(j - 1, m) + \mathcal{C}(j - 1, m - \theta_j) \quad (39)$$

to build up the dynamic programming table all the way up to $\mathcal{C}(J, M) = \mathcal{C}$.

The recursive formula (39) is justified because there are two kinds of sets $S \subset \{1, \dots, j\}$ which satisfy the constraint $\sum_{j' \in S} \theta_{j'} \leq m$. The first kind is those sets which do not contain j . The number of such sets which do not contain j is $\mathcal{C}(j - 1, m)$. The second type of set is those that contain j . One can write each of these sets in a unique way as $S = \{j\} \cup S'$ where $S' \subset \{1, \dots, j - 1\}$ and $\sum_{j' \in S'} \theta_{j'} \leq m - \theta_j$. There are exactly $\mathcal{C}(j - 1, m - \theta_j)$ of these sets, which justifies formula (39).

This dynamic programming algorithm runs in time $O(J \times M)$. This is tractable when M is small, but quickly becomes intractable when M is very large. Dyer (2003) gives a tractable approximation algorithm whose running time does not depend on M , and which computes \mathcal{C} with arbitrary precision.

Both the exact and approximate counting algorithms can be generalized to situations where the feasible sets S have to satisfy multiple knapsack constraints.²⁷ These running time of these algorithms increases exponentially with the number of constraints (see Dyer (2003) for more details on these generalizations.)

C A note on the interpretation of the Cobb Douglas Assumption

In Section 5, I use the assumption that production functions are Cobb-Douglas. This assumption is used to plot the blue line in Figure 5, and needs to be examined carefully. More specifically, I am assuming that the production function of industry i is $Y_i(S_i) = A_i(S_i) \prod_{j \in S_i} X_{ij}^{\alpha_{ij}} L_i^{1 - \sum_{j \in S_i} \alpha_{ij}}$, with the corresponding weight constraint being $\sum_{j \in S_i} \alpha_{ij} \theta_j \leq \tau_i$. Here the weight constraint only on the set S_i of chosen inputs, and does not depend on the intermediate demands X_{ij}, Y_i, L_i of the firm. More importantly, not all of the 2^N possible input sets will satisfy this weight constraint. For Leontief production functions, the constraint $\sum_{j \in S_i} \alpha_{ij} \theta_j \leq \tau_i$ has a natural interpretation, given in Examples 1 and 2: the coefficient $\alpha_{ij} = \frac{X_{ij}}{Y_i}$ corresponds to the number of units of good j needed to make one unit of good i , and the linear combination $\sum_{j \in S_i} \alpha_{ij} \theta_j$ represents the sum of weights of all “ingredients” in the production function. For Cobb-Douglas production functions, we can think of the different inputs $j \in S_i$ as component systems in a machine (e.g. the fuel system, engine system and exhaust system in a car). The components always have the same size $\alpha_{ij} \theta_j$, but firms in industry i can choose different levels of quality X_{ij} to produce output of quality Y_i . In this way, we can have a Cobb-Douglas production function where input demands X_{ij} can vary continuously, but the combinatorial constraint is discrete.²⁸ Another justification for Cobb-Douglas production functions follows a line of argument from Jones (2005), who argues that Cobb-Douglas production functions arise organically at the *aggregate level* when individual firms in an industry have a large menu of Leontief production functions. Jones (2005) proves

²⁷These arise naturally, for example, if there are both size and weight constraints.

²⁸Note that if we assumed instead that firms face the combinatorial constraint $\sum_{j \in S_i} \frac{X_{ij}}{Y_i} \theta_j \leq \tau_i$, then any set may become feasible by substituting away from heavy inputs and towards lighter inputs or labor (essentially, the firm can satisfy the constraint for any set S_i by setting X_{ij} to be very small for all intermediate inputs and L_i very large.) Thus, miniaturization in this setting would not lead to TFP gains.

this result for production functions that use only capital and labor as inputs. Azar (2021b) extends this result to show that, if firms have a menu of Leontief production functions with exogenous combinatorial constraints, then the aggregate production function will be approximated by a Cobb-Douglas production function with the same exogenous combinatorial constraint.