Innovation-driven growth in a multi-country world*

Till Gross§ Paul Klein¶

Carleton University Stockholm University

First draft: May 24, 2022
This version: May 22, 2024

Abstract

We develop a multi-country model of endogenous growth through innovation. The key feature of the model is that some ideas are globally applicable, while others are of local use only. Each country consists of a number of locations. There are innovation spillovers across locations and therefore across country borders. We argue that this model is both inherently plausible and consistent with an important set of growth facts. For instance, by computing a transition, we show that the model is capable of replicating a protracted decline in measured research productivity in the rich part of the world.

JEL Classification: E20, F43, O30, O40
Keywords: Innovation, endogenous growth, technological change, open economy, international spillovers

---

*We thank participants at the 2022 edition of the Vienna Macro Café, at the Carleton University and Bank of Canada brownbag series, at the 2023 Oslo Macro Conference, at the Paris School of Economics macro seminar, at the 2023 meeting of the Canadian Macro Study Group, at CERGE-EI, and at the University of Glasgow for helpful comments and suggestions. This research was supported in part by the Social Sciences and Humanities Research Council of Canada. Klein thanks *Jan Wallanders och Tom Hedelius stiftelse* for their generous financial support. Grant number: P22-0114.

§Mailing address: Department of Economics, Carleton University, C-870 Loeb Building, 1125 Colonel By Drive, Ottawa, ON K1S 5B6, Canada. Email: `till.gross@carleton.ca`.

¶Mailing Address: Department of Economics, Stockholm University, SE-10691 Stockholm, Sweden. Email: `paul.klein@su.se`.

1 Introduction

Endogenous growth models in the tradition of Romer (1990), where the ultimate source of growth is an expanding set of varieties, are appealing on several grounds. They are inherently plausible, intuitive, and tractable. Both balanced growth paths and the transitions to them can be characterized and computed in a reasonably straightforward manner. As a result, these models lend themselves well to growth accounting and to policy analysis. However, they appear to be incompatible with salient growth facts, especially when it comes to cross-country comparisons.

The purpose of the present paper is to construct and analyze a new multi-country endogenous growth model in the tradition of Romer (1990), but which departs from the original model in crucial ways. We argue that our model is consistent with an important set of growth facts, but still tractable enough to be used for growth accounting and policy analysis. We also find it inherently plausible.

Our main point of departure is to assume that the world consists of many distinct locations in which people live, work, save, and innovate, i.e. produce ideas for new varieties of goods. Some of these ideas are of universal applicability, but, crucially, some are useful only in the particular location where they were invented. We think of a *country* as simply a set of such locations. As we will show, assuming innovation spillovers across locations, and assuming also that jurisdictional boundaries do not constitute barriers to these spillovers, it does not matter how borders between countries are drawn. What matters are the people and technologies in the various locations, as well as the connections between the locations.

The main alternatives to product innovation models in the style of Romer (1990) are A-K models as in Rebelo (1991), where factor-accumulation drives economic growth; semi-endogenous growth models in the style of Jones (1995a), where population growth leads to productivity growth; and Schumpeterian models, where the ultimate source of growth is process innovation rather than product innovation – a canonical reference, among countless others, is Aghion and Howitt (1992). We do not argue that the Romer (1990) tradition is superior to these alternatives. Our contribution is to develop a suitably modified model in this tradition that is broadly in line with an important set of growth facts, inherently plausible, and tractable.

The main implications of our model may be summarized as follows:

(i) **Scale effects at the country level**: The population size of a country has no effect on the level or the growth rate of productivity if that country is fully integrated into the world economy with respect to the sharing of ideas. For partially integrated countries, larger ones will (ceteris paribus) enjoy a higher *level* (but not long-run growth rate) of productivity.[1]

(ii) **Convergence in growth rates**: All (fully or partially) integrated countries' productivity growth will converge to a common rate; this rate is determined by all countries' characteristics, and larger countries have a larger impact.[2]

(iii) **Non-convergence in levels**: A country's productivity depends on its characteristics; for instance, a country with a (permanently) higher research effort per capita will, other things being equal, enjoy a (permanently) higher level of productivity.

(iv) **Scale effects at the global level**: The productivity growth rate of the world as a whole depends positively, but not linearly, on global population size. Indeed for any reasonable parameter values, the relationship is strictly concave.

(v) **Catchup**: A relatively poor country that is open to the flow of ideas from abroad tends to grow faster than an already rich country.

(vi) **Protracted decline in measured research productivity**: The model can easily be calibrated in such a way as to replicate a sustained decline in measured research productivity in the initially richer part of the world.

As we discuss in Section 2, there are several models in the existing literature that share some of these properties, but we are not aware of any that shares all of them. Meanwhile, we argue here that these implications are broadly in line with the facts.

[1]We give a precise definition of what it means to be *fully integrated into the world economy with respect to the sharing of ideas* in Section 3 below. In Section 4.2, we define what it means to be only partially integrated.

[2]In Section 4 we spell out the precise sense in which larger countries have a larger impact.

A more detailed discussion of the empirical evidence in found in Appendix D, available online.[3] Here we just briefly note the following.

On point (i), there is no empirical association between the size and either the level or the growth rate of GDP per capita. Alesina et al. (2005) take a more detailed look at the data and find that size matters only to the extent that an economy is not sufficiently open. In other words, open economies do not suffer from being small, but less open ones do. On point (ii), there is of course no corresponding fact, given that we are unlikely to be on a balanced-growth path. Indeed, the transitional dynamics of our model can be very protracted, and so growth rates may differ across countries for a long time, as indeed they do in the data. Meanwhile, the long-run convergence of growth rates is very desirable from an applied standpoint, because it enables us to compute both transitions and balanced-growth paths. Point (iii) is desirable because it allows for cross-country growth accounting, and is of course broadly in line with the evidence as well. For instance, the gap between U.S. and Canadian real GDP per capita has been about 20 percent for many decades now.

Point (iv) is perhaps a bit controversial in that the scale effect shared by many innovation-driven endogenous growth models has been heavily criticized for being counterfactual. Nevertheless, we would argue that the evidence in Galor (2010), who documents that global growth in GDP per capita really only took off once the global population had reached 1.5 billion, is consistent with a mild scale effect operating at the global level. We may note in this context that our model implies that the long-run growth rate is a strictly concave (as opposed to linear) function of the size of the world's population. This means that the scale effect is not as stark as in the Romer (1990) model and in fact quite small in our baseline scenario. Moreover, removing the (global) scale effect comes at a steep price in terms of plausibility. Specifically, our model implies that the distribution of population across jurisdictions matters for growth as soon as there is any relevant heterogeneity across jurisdictions. (For instance, if one country has a more productive innovation sector than another.) Models without a global scale effect lack this very plausible implication.

Meanwhile, as far as point (v) is concerned, there are many examples of relatively

[3] All appendices are found in http://paulklein.ca/newsite/research/growthopenappendix.pdf.

poor countries and regions tending to catch up with (though not necessarily to converge to) their relatively rich peers. This is especially so with countries and regions that are highly integrated, such as Canadian provinces, U.S. states and regions inside the European Union. It is also evident in the case of countries that break their isolation and join the world economy such as Vietnam or China in the aftermath of the reforms of the 1980s.

Regarding point (vi), an interesting feature of our model is that it can quite naturally generate a scenario where measured research productivity declines for a very long time in a technologically advanced country (or set of countries). As emphasized by Jones (1995b) and, more recently, by Bloom et al. (2020), ever more resources in the United States have been poured into research and development, with meagre or no apparent effects on the growth rate of GDP per capita. As these authors quite rightly point out, this is not consistent with a balanced growth path of any innovation-driven endogenous growth model in the tradition of Romer (1990). Nevertheless, what we show is that it *is* consistent with a suitably defined *protracted transition* to a balanced-growth path. A persistent decline in measured research productivity can happen on such a path for two distinct reasons in our model.

The first mechanism through which measured research productivity may decline along a transition is related to our assumption that new ideas are invented on the basis of old ones. This is the assumption that leads to the implication of convergence in growth rates. We will get to the mathematical details of the model in Section 3, but for now it may be helpful to use a metaphor. Imagine two skaters pulling a sled, each connected to the sled by a bungee cord. Evidently, the speed of both skaters depends in some sense on the efforts of both. Imagine now that one of the skaters starts skating more quickly. At first, she races ahead and a gap opens up between the skaters. This is what happens in our model too. Suppose, for instance, that a country or set of countries (but not the whole world) suddenly becomes better at coming up with new useful ideas. As a result, those countries, like one of our skaters, quickly take the lead. After a while, however, the leading countries find that they are so far ahead that the stock of existing ideas in the rest of the world is so far behind that it is not very helpful for coming up with new ideas. At this point, research productivity declines in the leading countries, just as our leading skater slows down as the bungee cord becomes taut and the quicker skater is encumbered by having to pull a larger fraction of the weight of the sled.

The second mechanism is related to a congestion effect in research. Imagine a scenario where a large country joins the world economy, and hence joins the world stage by developing ideas that can be adopted elsewhere. This is helpful for the growth of the world economy as a whole, but it also implies that inventors in technologically advanced countries (say, the United States) face more competition. What this means intuitively is that ideas invented in the United States may to some extent be scooped by similar ideas developed in China, so that U.S. based inventors produce fewer and less valuable patents for the same effort.

Our main innovation is that we allow for two types of growth-generating ideas: ones with a narrow range of applicability (we call these ideas "local") and ideas with a universal range of applicability (we call these ideas "global"). Global ideas are universal; such an idea invented in one place can be applied anywhere else. Local ideas are the opposite; such an idea can only be applied in the location it was invented. When developing new ideas, any inventor stands on the shoulders of giants: their productivity depends on the world's total stock of global ideas as well as the ideas invented in the inventor's own location. In this context, we think of a location as an entity smaller than a country but still including many people.[4]

We think that in reality there are truly global ideas, such as computers, that can be used by humans anywhere. Other ideas are regional, such as drip irrigation which is useful mainly in arid climates. Still other ideas are truly local in their applicability, such as types of food that cater to local tastes or solutions to the problems that arise when building a road or a tunnel through a particular landscape. Each bridge is unique and custom-made for the particular local conditions, even though general principles of bridge building apply universally. Overall, ideas are probably best described by a continuum of applicability ranges, from uniquely local to fully universal; in this paper, we focus on the two extremes (leaving out all ideas of intermediate applicability) to keep the model tractable.

The distinction between global and local ideas is not only plausible in itself, but

[4]A related but quite distinct framework is that of McGrattan and Prescott (2009, 2010), where growth is not endogenous, but where the level of output depends on intangible or "technology" capital which is non-rival across locations. They find that the degree of openness to foreign multinationals' technology capital is an important determinant of the level of output.

allows us to reconcile convergence in growth rates across jurisdictions with indefinitely persisting differences in the level of productivity resulting from different rates of innovation. In other words, thanks to our distinction between local and global ideas, our model implies that, even among countries that are fully integrated into the world economy with respect to the flow of ideas, more innovative countries will enjoy a higher *level* of output per capita. At the same time, countries that only differ in size will share the same productivity. Moreover, our model implies that all countries will eventually *grow* at the same rate.

Another feature of our model is that some intermediate goods are provided competitively and others monopolistically. This coexistence of monopolistically and competitively supplied intermediate goods arises from the assumption that patents expire after one model period (20 years). Thus "new" (just invented) varieties are supplied monopolistically, and "old" (previously invented) varieties are supplied competitively. We view this as a realistic assumption, but it also keeps the model more tractable, and makes it useful for policy analysis.

We focus not only on balanced growth paths, but also on transitions. Given that transitions turn out to be potentially very protracted, this is crucial for exploring the model's implications and for comparing them with data. These protracted transitions enable us to shed new light on some important empirical phenomena, notably catch-up growth and declining research productivity. We emphasize in this context that we do not merely compute the dynamics of a linearized version of our model around the balanced growth path, but the fully-fledged dynamics to an arbitrary degree of precision. This is of course crucial when doing policy analysis, for which comparing balanced growth paths or local dynamics is insufficient.

The rest of this paper is organized as follows. Section 2 discusses the related literature. Section 3 presents the model. Section 4 shows our analytical results (proofs are in Appendix A) and discusses three model extensions. Section 5 describes our numerical experiments. Section 6 concludes.

2 Related literature

Rivera-Batiz and Romer (1991) is an early attempt to analyze the implications of the Romer (1990) model in a multi-country world. The paper does not develop a multi-country model, but instead points out the presence of a scale effect on the growth rate, which implies that the world economy grows faster when ideas can flow freely across borders than if the world is divided into hermetically sealed jurisdictions.

Barro and Sala i Martín (1997) (see also Barro and Sala i Martín (2004), chapter 8) consider a two-country version of the Romer (1990) model. Country 1 is the leader, where all of the innovation takes place. Country 2 is the follower, where no innovation takes place; instead it imitates, at a cost and with a delay, the ideas invented in country 1. Our model shares with theirs the property that a lagging country can catch up quite quickly if it integrates into the world economy in the sense that it becomes open to adopting ideas from abroad. However, in Barro and Sala i Martín (1997) there can only be a single country at a time that innovates and hence only one country that contributes to expanding the world technological frontier. In our model, all countries innovate and all countries contribute to expanding the world frontier. For instance, while it is true that China is currently catching up to more technologically advanced countries like the United States, its researchers and entrepreneurs also contribute to the global frontier of knowledge.[5] Another important difference between our model and the one in Barro and Sala i Martín (1997) is that a doubling of the research effort in the leading country would more than double productivity growth in Barro and Sala i Martín (1997), just as in Romer (1990), while the effect is concave in our model and quite small in the baseline calibration.

Since Jones (1995b), it has been accepted wisdom that the scale effect property of the endogenous growth model of Romer (1990) is counterfactual and should be avoided in favour, perhaps, of a "semi-endogenous" growth model along the lines of Jones (1995a). A semi-endogenous growth model avoids the implication that the growth rate of productivity increases in the size of the population. However, in a semi-endogenous growth model, growth in productivity comes to a complete stop if the population stops growing. Indeed, the main feature of semi-endogenous growth

[5] See, for instance, Tollefson (2018).

models is that productivity growth is linearly increasing in the population growth rate. This is empirically problematic. Across countries, there is no evidence of a positive association between population growth and growth in output per capita. But perhaps the level at which this association operates is the world as a whole, not an individual country. That distinction does not exist in Jones (1995a), but it does in Jones (2002) which presents a multi-country version of the model in Jones (1995a).

The model in Jones (2002) features a positive relationship, but now at the global as opposed to the country level, between per-capita income growth and population growth, which is consistent with the time-series evidence. There is no such relationship across countries, because, in his model, all countries enjoy the same productivity growth rate. Meanwhile, in his model (as well as in Solow's and ours) "a higher population growth rate reduces the steady-state capital-output ratio because more investment must go simply to maintain the existing capital-output ratio in the growing population." Because productivity grows at the same rate in all countries, this means that output per capita is lower in countries experiencing a high population growth rate. However, the assumption in Jones (2002) that ensures that productivity growth is common across countries, is that, in the language of our model, all ideas are global. This in turn implies not only that all countries share the same productivity growth rate but also that that they have the same *level* of productivity, independent of country characteristics.

Klenow and Rodriguez-Clare (2005) follow Jones (1995a, 2002) in avoiding a scale effect on growth, but where the *level* of output per head may depend on policy as well as size. In their model, if all countries are fully open, then the *level* of productivity in each country depends on its characteristics, including policy. With incomplete openness, the level of productivity increases in country size. Our model has a similar feature, and we find it plausible. However, if population growth rates were allowed to vary across countries in Klenow and Rodriguez-Clare (2005), even if only along a transition, productivity growth rates would be positively associated with population growth rates, which is what Jones (2002) was aiming to avoid. By contrast, our model implies a negative relationship between population growth and productivity growth along a transition.

Peretto (2018) develops a closed-economy endogenous-growth model without a scale effect, where productivity growth is increasing in population growth, but does not necessarily come to a halt if the population stops growing.

Howitt (2000) shares some assumptions and implications with our paper. For a wide range of parameter values, countries in Howitt's model converge to the same growth rate, but to different levels, just as in our model. His model also exhibits catchup in the sense that otherwise similar countries with initially low productivity will catch up to countries whose initial productivity is higher. Meanwhile, in Howitt (2000), the size of a country's population has no direct effect on its level of productivity or on its growth rate, similar to our model. However, our models disagree in that our model, unlike Howitt's, *does* have a scale effect at the *global* level. The source of this difference is that Howitt achieves invariance to country size in a rather different way from how we do it. Our notion is that the level of productivity and its growth should not depend on how frontiers are drawn between jurisdictions, provided that there is a free flow of ideas (but not necessarily a free flow of capital, goods or labour). Splitting a jurisdiction in half should not affect the level of productivity nor its growth rate in either of the new jurisdictions, and certainly not the level or growth of productivity of the world as a whole.[6] In a setup with this property, such as ours, a larger country naturally has a larger impact on the common global growth rate than a smaller one does.

In Howitt's model the contribution of a country to the growth rate of the world's technological frontier *is exogenously fixed*, independently of a country's size. Howitt's model implies that a doubling of the population of each existing country would leave the long-run growth rate unchanged. Doubling the number of countries, however, would double the growth rate.

Another related paper is Eaton and Kortum (1999).[7] Their model shares with ours the property that all countries grow at the same rate in the long run and that the level of productivity depends on country-specific characteristics. However, in the baseline version of their model, each researcher/inventor is, other things being equal,

[6]Note that *merging* two jurisdictions only has unambiguous meaning if the jurisdictions only differ, if at all, with respect to population size. Splitting two jurisdictions is less problematic: it means splitting the population and letting the new jurisdictions inherit all other properties from the jurisdiction that they were carved out of.

[7]A very similar model is found in Eaton and Kortum (1996).

more productive in a big country than in a small one. That is not the case in the baseline version of our model; there, size does not matter for research productivity. If the baseline of Eaton and Kortum (1999) is changed to eliminate this size effect, by removing cross-country barriers to the flow of ideas, then all countries have the same level of productivity. That is not true in our model. In our model, a more innovative country has a higher level of productivity than a less innovative one, even if national borders are no obstacle to the free flow of ideas. The source of this result is our distinction between local and global ideas.

Sampson (2023) constructs a model of international trade, endogenous growth, and technology diffusion. Growth is generated by firms investing in superior technology, thereby increasing their productivity. There are technological spillovers across firms, both within and between countries. The focus of the paper is very different from ours, but the model in Sampson (2023) shares with our model the properties (ii), (iii), and (v) described in Section 1 of the present paper. With respect to property (i), it is noteworthy that Sampson does not even contemplate the possibility of full international integration when it comes to technology spillovers, but takes as given that spillovers are stronger within a country than between countries. Nevertheless, it seems very plausible to suppose that his model would share property (i) if such a possibility were contemplated. However, when it comes to property (iv), the global scale effect in Sampson (2023), i.e. the mapping from the level of population to the long-run growth rate of output per head, is, just as in Romer (1990), convex rather than concave. Other papers that investigate the dynamic gains from trade effects through knowledge spillovers include Buera and Oberfield (2020), Perla et al. (2021), Cai et al. (2022), and Hsieh et al. (2023).

Trouvain (2023) provides an explanation for the productivity slowdown in currently rich countries that has some intriguing similarities with our own account (see Section 5.3) of how measured research productivity may go through a period of decline as a large and initially backward country enters the world economy. In Trouvain (2023), the posited mechanism is that the frontier country puts less resources into adoption of new technologies, focussing instead on developing new ideas, as an initially backward region joins the world economy and expands the size of the market for new ideas. We view these two perspectives, our own and that of Trouvain (2023), as complementary.

3 The model environment

Our model is an open-economy version of the one presented in our previous work, Gross and Klein (2022), which in turn is a modified version of Romer (1990). A final good is assembled by competitive firms from a variety of intermediate goods, which are imperfect substitutes. Blueprints for new varieties of intermediate goods can be invented, and these new varieties are the source of endogenous growth. Varieties may be universally applicable ("global"), or be useful only in a specific location ("local"). There are several countries, each of which is constituted by a number of locations, which are in turn inhabited by many people. New varieties are produced by monopolistic firms, and already existing varieties (whose monopoly privileges have expired) are produced by perfectly competitive firms.

3.1 Locations and jurisdictions

There are J locations in the world. A location j is populated by a measure $\chi(j)$ of households. The population of a location defines the scope of a local idea, in the sense that the idea is useful to anyone living in that location. The scope of a global idea, on the other hand, is the population of the world as a whole.

A jurisdiction (or country) i is simply a set $\mathcal{A}_i$ of locations. Country i's population is then naturally defined as $\chi_i = \sum_{j \in \mathcal{A}_i} \chi(j)$. The distinction between location and jurisdiction is important, because in our model, the scope of a local idea is a location, not a country.

3.2 Households

A representative household living in location j consists of a worker and an innovator. Such a household chooses consumption $c_t(j)$, hours worked in production $l_t(j)$, innovation effort $h_t(j)$, and next-period assets $a_{t+1}(j)$ in order to maximize

$$U(j) = \sum_{t=0}^{\infty} [\beta(j)]^t \left[u(c_t(j), l_t(j), h_t(j))\right] \tag{1}$$

subject to

$$c_t(j) = w_t(j)l_t(j) + P_t(j)\hat{z}_t(j) + [1 + R_t(j)]\, a_t(j) - a_{t+1}(j) \tag{2}$$

for all $t = 0, 1, \ldots$, where $\hat{z}_t(j)$ is the measure of newly invented ideas; we specify the production function for ideas below. We think of innovators as inventors who auction off their newly invented blueprints and the accompanying monopoly rights at a price $P_t(j)$ per unit measure of ideas to the highest bidder. Alternatively, and equivalently, we can think of innovators as entrepreneurs who create new businesses and sell them to large investors/corporations. Workers receive a wage rate $w_t(j)$ and $R_t(j)$ is the rate of return on assets. Initial asset holdings $a_0(j)$ are exogenously given.

3.3 Ideas and research

An innovator in location j uses innovative effort $h_t(j)$ to produce $\hat{z}_t(j)$ new ideas according to the function

$$\hat{z}_t(j) = \eta(j) \cdot h_t(j) \left(\frac{(1-m)Z_t(j)}{[H_t(j)]^{\Lambda}} \right)^{\rho} \left(\frac{mZ_t^W}{\left[H_t^W\right]^{\Lambda}} \right)^{1-\rho}. \tag{3}$$

The parameter $\eta(j) > 0$ governs the productivity of the innovation process. There is a standing-on-shoulders and a stepping-on-toes component to innovation for both local and global ideas. An exogenous fraction m (for *mondial*) of all new ideas are global, and, correspondingly, a fraction $(1-m)$ are local. We denote the measure of all ideas previously (before period t) invented in location j by $Z_t(j)$. This suggests that $(1-m)Z_t(j)$, the measure of all previously invented local ideas in location j, is the relevant measure of local "shoulders" that inventors can use as a source of inspiration.[8] The parameter $\rho \in (0, 1)$ governs how important (old) local ideas are in the production of new ideas. Inventors also benefit from the measure of global ideas previously invented anywhere in the world, mZ_t^W, where

$$Z_t^W := \sum_{\ell=1}^{J} Z_t(\ell). \tag{4}$$

[8]The initial $Z_0(j)$ is exogenously given and its fraction of global ideas is also m.

Besides the positive externality of standing on shoulders, there is a negative congestion externality, which we refer to as stepping on toes; $\Lambda \in (0,1)$ governs the strength of this externality. The higher the aggregate effective research effort $H_t(j)$ in their location, the lower is the productivity of each individual inventor. Every inventor is small relative to her location, and so she regards $H_t(j)$ as exogenously given, though in equilibrium all inventors in a location exert the same amount of effort and we have $H_t(j) = \chi(j) \cdot \eta(j) \cdot h_t(j)$. A congestion effect also operates at the global level: the higher the total world-wide research effort H_t^W, the lower is the productivity of an individual researcher, where

$$H_t^W := \sum_{\ell=1}^{J} H_t(\ell). \tag{5}$$

The parameter ρ, which governs the relative importance of local and global "shoulders", also governs the relative importance of local and global congestion externalities.

We may think of the production function for ideas in the following way. Each innovator works on a number of projects and is unaware of whether a specific project may turn out to be local or global in applicability. A research project benefits from the existing stock of ideas, both local and global. At one extreme, when $\Lambda \to 0$, then all research projects benefit from all previously invented ideas. At the other extreme, when $\Lambda \to 1$, then old ideas may only be used for one research project at a time, and it is the measure of previously invented ideas per research project that determines research productivity. For intermediate values of Λ, previously invented ideas benefit several projects at a time, but there is some congestion, so that not all ideas can be used by all projects.

Notice that we assume that every jurisdiction is *fully integrated into the world economy with respect to the sharing of ideas*, in the following sense. As far as inventing new ideas is concerned, jurisdictional (country) boundaries do not matter. Only the local and total world stock of existing ideas matter for standing on shoulders, and only the local and total world research effort matters for stepping on toes.

Due to the assumed existence of a representative agent in every location j, the measure of new ideas invented in location j is $\widehat{Z}_t(j) = \chi(j)\widehat{z}_t(j)$ and the law of motion for the stock of ideas is

$$Z_{t+1}(j) = Z_t(j) + \widehat{Z}_t(j). \tag{6}$$

The distinction between local and global ideas gives rise to a distinction between "invented ideas" and "available ideas" in a given location. In the absence of frictions to the flow of ideas (see Section 4.2 for an analysis of such frictions), all global ideas ever invented are available everywhere, but the only local ideas available in a location are those invented in (and for) that location. The relevant concept for production in a given location is thus the measure of ideas specific to that location plus the measure of global ideas worldwide. For old ideas, the definition is

$$\mathbf{Z}_t(j) := (1-m)Z_t(j) + mZ_t^W \tag{7}$$

and, similarly, the definition for new ideas is

$$\widehat{\mathbf{Z}}_t(j) := (1-m)\widehat{Z}_t(j) + m\widehat{Z}_t^W, \tag{8}$$

where naturally $\widehat{Z}_t^W = \sum_{\ell=1}^{J} \widehat{Z}_t(\ell)$.

3.4 Production

For each idea, there is one type of intermediate good z. Just like the corresponding ideas, intermediate goods may be specific to a particular location; we then call them "local," or they may be of universal applicability; we then call them "global".

Firms rent capital $K_t(j,z)$ and hire labour $L_t(j,z)$ to produce intermediate goods according to the following Cobb-Douglas production function

$$Q_t(j,z) = K_t^{\alpha}(j,z)L_t^{1-\alpha}(j,z). \tag{9}$$

Capital depreciates at rate $\delta \in [0,1]$. Old (previously invented) ideas are in the public domain and the corresponding intermediate goods are provided competitively. New ideas are proprietary and the corresponding intermediate goods are provided monopolistically (possibly licensed by a patent-holder from another location). All producers of intermediate goods maximize profits.

The final good is produced by perfectly competitive, profit-maximizing firms that use the intermediate goods as inputs. Output in each location j is

$$Y_t(j) = \left[\int_{\mathbb{Z}_t(j)} Q_t^{\sigma}(j, z) d\nu(z) \right]^{1/\sigma} \tag{10}$$

where $\mathbb{Z}_t(j)$ is the set of ideas, old and new, available in location j; the corresponding measure of this set is $\mathbf{Z}_t(j) + \widehat{\mathbf{Z}}_t(j)$.

3.5 Market clearing and equilibrium

We assume that labour and capital are immobile across locations. However, it is conceptually straightforward to allow for country-wide or global capital markets.

We have labour, capital, and global goods market clearing:

$$L_t(j) := \int_{\mathbb{Z}_t(j)} L_t(j, z) = \chi(j) l_t(j) \; \forall \, j \tag{11}$$

$$K_t(j) := \int_{\mathbb{Z}_t(j)} K_t(j, z) = \chi(j) a_t(j) \; \forall \, j \tag{12}$$

$$\sum_{j=1}^{J} Y_t(j) = \sum_{j=1}^{J} \chi(j) c_t(j) + \sum_{j=1}^{J} [K_{t+1}(j) - (1 - \delta) K_t(j)] \tag{13}$$

Despite the coexistence of competitive and monopolistic firms, of global and local intermediate goods, the equilibrium output function is remarkably simple (proofs are in Appendix A).

Proposition 1 *Total output in location* j *is*

$$Y_t(j) = \text{TFP}_t(j) K_t^{\alpha}(j) L_t^{1-\alpha}(j), \tag{14}$$

where endogenous total factor productivity (TFP) *is given by*

$$\text{TFP}_t(j) = \frac{\left(\mathbf{Z}_t(j) + \widehat{\mathbf{Z}}_t(j) \sigma^{\sigma/(1-\sigma)} \right)^{1/\sigma}}{\mathbf{Z}_t(j) + \widehat{\mathbf{Z}}_t(j) \sigma^{1/(1-\sigma)}}. \tag{15}$$

TFP does not simply depend on the total measure of all available ideas, old plus new, $\mathbf{Z}_t(j) + \widehat{\mathbf{Z}}_t(j)$, but also on the relative measure of new vs. old ideas. This is because competitive and monopolistic firms will produce different quantities.

4 Main analytical results

We start this section by presenting a very convenient property of our model, namely that all locations will converge to the same rate of output growth on a balanced-growth path (BGP).[9] We then proceed to present the main analytical results at the *country* level for the baseline version of our model, and then move on to some extensions. All proofs are relegated to Appendix A.

Lemma 1 *The growth rate of output in location* j, $g_t(j)$, *on a BGP satisfies*

$$1 + g_t(j) = \left(1 + \widehat{\mathbf{Z}}_t(j)/\mathbf{Z}_t(j)\right)^{\frac{1-\sigma}{\sigma(1-\alpha)}}. \tag{16}$$

Our Lemma states that all growth eventually comes from the development of new ideas, endogenously driving TFP higher and higher. From the Lemma, our important result about growth convergence follows:

Proposition 2 *If each location converges to a BGP, then all locations converge to a common growth rate of output.*

While we have relegated all proofs to the Appendix A, we want to briefly explain the intuition for this result. The term $\widehat{\mathbf{Z}}_t(j)/\mathbf{Z}_t(j)$ in Equation (16) contains the growth rate of local and global ideas. By definition, global ideas grow at the same rate across all locations. The growth rate of local ideas depends on research effort and innovative productivity, but these are time invariant on a BGP; the stocks of ideas in different locations, however, vary over time. We can rewrite Equation (3) as

$$\frac{\widehat{Z}_t(j)}{Z_t(j)} = \left(\sum_{\ell=1}^{J} \frac{Z_t(\ell)}{Z_t(j)}\right)^{1-\rho} \eta(j)h_t(j)\left(\frac{(1-m)}{[H_t(j)]^{\Lambda}}\right)^{\rho}\left(\frac{m}{\left[H_t^W\right]^{\Lambda}}\right)^{1-\rho}. \tag{17}$$

[9]Because we assume a constant population, there is no useful distinction between growth in output and growth in output per capita. See Section 5.2 for an analysis of how the model works in the presence of population growth.

If we posit (in order to derive a contradiction) that location 1 keeps growing faster than others, then $Z_t(1)$ would keep growing relative to $Z_t(\ell)$ for all $\ell \neq 1$. Consequently, location 1's standing-on-shoulders terms $Z_t(\ell)/Z_t(1)$ would all converge to zero, except for its own shoulders term, $Z_t(1)/Z_t(1)$, which is obviously constant and equal to one. For all the other locations, we have a contradiction: the standing-on-shoulders terms $Z_t(1)/Z_t(\ell)$ would keep growing indefinitely, and hence the growth rate for any location $j \neq 1$ would keep increasing until it overtakes the growth rate of location 1.

4.1 Baseline version

For the rest of the paper, we make two important assumptions that enable us to compare country outcomes by *average* country characteristics, as opposed to needing, for example, the *whole distribution* of research productivity across locations within each country.

Assumption 1 The scope of local ideas is the same for all locations.

Assumption 1 can be implemented in two ways: (i) to assume an equal population in each location, and (ii) to adjust for location population size in the production of new ideas and intermediate goods. The first method is straightforward and we use it in most of the paper. In this case, since population units are arbitrary, we may as well normalize the measure per location to one, i.e. $\chi(j) = 1$ for all j.[10]

Assumption 2 All locations within the same country are symmetric.

Assumption 2 means that preferences of households, initial per-capita assets and ideas, and research productivities are identical for all locations within the same country. This assumption implies that we can use the notation η_i to denote research productivity in country i, i.e. in any location $j \in \mathcal{A}_i$. Similarly, we can use χ_i to denote the size (both in terms of population and number of locations) of a country. Incidentally, and this is relevant when we allow for population growth, we do not necessarily require each location to have the same population size; what matters is

[10]This method cannot be used, however, when allowing for population growth (one of our extensions; see Section 5.2). In that context we use the second approach; the details are explained in Appendix B.

that the scope of local ideas is the same across locations, and that is taken care of by Assumption 1.

In what follows, we will refer to country-level variables. These follow in a very straightforward manner from the corresponding location-level variables, thanks to Assumptions 1 and 2. Indeed, per-capita variables are exactly the same at the country level as at the location level; for example $h_{i,t} = h_t(j)$ for all $j \in \mathcal{A}_i$. For aggregates on the other hand, we have

$$Y_{i,t} = \sum_{j \in \mathcal{A}_i} Y_t(j) \tag{18}$$

$$Z_{i,t} = \sum_{j \in \mathcal{A}_i} Z_t(j) \tag{19}$$

$$\widehat{Z}_{i,t} = \sum_{j \in \mathcal{A}_i} \widehat{Z}_t(j) \tag{20}$$

$$K_{i,t} = \sum_{j \in \mathcal{A}_i} K_t(j) \tag{21}$$

$$L_{i,t} = \sum_{j \in \mathcal{A}_i} L_t(j). \tag{22}$$

It follows that

$$\mathbf{Z}_{i,t} = (1-m)Z_{i,t}/\chi_i + m \sum_{\ell=1}^{I} Z_{\ell,t} \tag{23}$$

$$\widehat{\mathbf{Z}}_{i,t} = (1-m)\widehat{Z}_{i,t}/\chi_i + m \sum_{\ell=1}^{I} \widehat{Z}_{\ell,t}. \tag{24}$$

We can then write total output in country i as

$$Y_{i,t} = \text{TFP}_{i,t} K_{i,t}^{\alpha} L_{i,t}^{1-\alpha}$$

where

$$\text{TFP}_{i,t} = \frac{\left(\mathbf{Z}_{i,t} + \widehat{\mathbf{Z}}_{i,t}\sigma^{\sigma/(1-\sigma)}\right)^{1/\sigma}}{\mathbf{Z}_{i,t} + \widehat{\mathbf{Z}}_{i,t}\sigma^{1/(1-\sigma)}}.$$

Note that location TFP is equal to country TFP under our assumptions. It is then clear that all *countries* converge to a common growth rate; a shared characteristic of all multi-country endogenous growth models that we are aware of. This is of course a highly desirable property of a multi-country growth model. If growth rates across countries do not converge to a common limit, then the country with the highest growth rate would eventually dominate the world economy, its share of world output converging to 100 percent; in this sense, the balanced growth path would be degenerate. The convergence of growth rates to a common limit is thus a necessary condition for computing both transitions and balanced growth paths: otherwise a non-degenerate BGP does not exist and we cannot compute a transition without a non-degenerate BGP.

Proposition 3 *If each country's economy converges to a balanced growth path, then the level of output per head on that balanced growth path depends on country characteristics. In particular, countries with a relatively high research productivity or research effort will have a relatively high level of output per head. If two countries differ only in population size, their level of output per capita will be identical.*

Proposition 3, that countries with different research productivities will have different levels of GDP per capita, even in the long run, is intuitive. However, it is worth noting that this feature is not shared by all endogenous growth models where countries are perfectly open to the exchange of ideas. A prominent example is Jones (2002), where all countries converge to the same level of productivity. This is because all ideas in Jones (2002) are global; it is the existence of local ideas that allows for cross-country differences in productivity in our model. The feature that population size does not affect productivity (at least for economies fully integrated into the world economy) has empirical support; see, for instance, Alesina et al. (2005).

While our model does not have a scale effect at the country level (country size does not affect growth), there is a scale effect at the global level (the world population does affect growth). To keep this scale effect in check, we make an additional assumption, described and discussed below.

Assumption 3 The parameter space is restricted so that $\dfrac{1-\sigma}{\sigma(1-\alpha)} \leq 1$.

Proposition 4 *The long-run global productivity growth rate is a strictly increasing function of the world population, holding the distribution of the population across countries and* *research effort per capita constant. Assumption 3 is a sufficient, but not necessary condition for it to also be strictly concave.*

Proposition 4 establishes that there is a positive global scale effect, i.e. that the global growth rate is greater the larger the total world population. However, unlike in Romer (1990) and Barro and Sala i Martín (1997), this scale effect is concave and not convex in population size. To understand how this result relates to our model features and Assumption 3, we rewrite the growth rate in Equation (16) on a BGP as a function of the global population x:

$$1 + g(x) = \left(1 + \mathring{\mathbf{Z}}(x)\right)^{\frac{1-\sigma}{\sigma(1-\alpha)}} . \tag{25}$$

We hold the population shares of each country χ_i and the scope of local ideas constant, i.e. the number of locations doubles when world population doubles. In this context, we define $\mathring{\mathbf{Z}} := \widehat{\mathbf{Z}}_{i,t}/\mathbf{Z}_{i,t}$. We can suppress any dependence on time and country, because, on a BGP, $\widehat{\mathbf{Z}}_{i,t}/\mathbf{Z}_{i,t} = \widehat{\mathbf{Z}}_{j,s}/\mathbf{Z}_{j,s}$ for all countries i, j and all time periods t, s. Holding research effort constant (which in any case does not respond very much to population size in our numerical exercises), we can write $\mathring{\mathbf{Z}}(x) = x^{(1-\Lambda)(1-\rho)}F$.[11]

The growth rate in available ideas, $\mathring{\mathbf{Z}}$, is thus a concave function of world population x. This concavity is guaranteed by two of the core features of our model: that local ideas contribute to the production of new ideas (and hence $\rho > 0$) and that there is some research congestion (and hence $\Lambda > 0$).

The growth rate of *output*, however, also involves the exponent $\dfrac{1-\sigma}{\sigma(1-\alpha)}$, and Assumption 3 restricts this to be no greater than one. Assumption 3 implies parameter restrictions on the capital exponent in the production function, α, joint with σ, which governs the elasticity of substitution between varieties. In this context, we may view the capital share parameter α as an amplifier of the growth effects of innovation. For a given σ, a higher α makes the exponent larger; in the limit, as α approaches 1, the exponent approaches infinity, so we want to restrict this amplification. In any case, empirically relevant parameter values easily satisfy the assump-

[11]We define $F := m^{1-\rho}(1-m)^{\rho}\left(\dfrac{\sum_{\ell=1}^{I}(\eta_\ell h_{\ell,t})^{\frac{1-\Lambda\rho}{1-\rho}}\chi_\ell}{\left(\sum_{\ell=1}^{I}\eta_\ell h_{\ell,t}\chi_\ell\right)^{\Lambda}}\right)^{1-\rho}$.

tion that $\dfrac{1-\sigma}{\sigma(1-\alpha)} \leq 1$. Our baseline choice for these parameters is $\alpha = 1/3$ and $\sigma = 6/7$, implying an exponent of 0.25. Even for values at the extreme of what we would consider reasonable, $\alpha = 2/5$ and $\sigma = 3/4$, the assumption still holds with an exponent of about 0.56.

Proposition 5 *Let $g(\chi)$ denote the global productivity growth rate, where χ is the vector of the population sizes of all countries, not necessarily holding world population constant. The distribution of population across countries, χ, generically has a non-zero effect on the long-run global productivity growth rate. In particular, if the population of country 1 contributes positively to $g(\chi)$, i.e. $\frac{\partial g(\chi)}{\partial \chi_1} > 0$, and the effective research effort per capita is higher in country 1 than in country 2, i.e. $\eta_1 h_{1,t} > \eta_2 h_{2,t}$, then $g(\chi)$ would be higher if the population of country 1 increased and the population of country 2 decreased by an equal amount, i.e. $\frac{\partial g(\chi)}{\partial \chi_1} > \frac{\partial g(\chi)}{\partial \chi_2}$, holding research effort per capita constant.*

Corollary 1 *If two countries have the same effective research effort per capita, i.e. $\eta_1 h_{1,t} = \eta_2 h_{2,t}$, and both contribute positively to the long-run global productivity growth rate, i.e. $\frac{\partial g(\chi)}{\partial \chi_1} > 0$ and $\frac{\partial g(\chi)}{\partial \chi_2} > 0$, then the impact of the larger country on productivity growth $g(\chi)$ will be larger, i.e. if $\chi_1 > \chi_2$ then $g(0, \chi_2, \chi_3, \ldots, \chi_I) < g(\chi_1, 0, \chi_3, \ldots, \chi_I)$.*

Proposition 5 states that the distribution of population across countries matters for the long-run growth rate. For example, a world where 99% of the global population live in a country with a high effective research effort and 1% in a country with a low effective research effort would enjoy a higher growth rate than a world with the reverse population weights. Meanwhile, Corollary 1 says that larger countries contribute more to global growth than smaller countries in the sense that if we eliminate the larger country, the growth rate drops by more than if we eliminate the smaller country. We find these two features intuitively very appealing. Indeed, Klenow and Rodriguez-Clare (2005, pp. 840-41) argue that the latter ("larger countries contribute more to world growth than smaller countries") is a desirable feature of an endogenous-growth model. Nevertheless, *neither* feature is present in their preferred model, nor in Jones (1995a, 2002) or Howitt (2000).

4.2 Extensions

We briefly discuss three extensions here; in Appendix B, available online, we spell out in more detail how we incorporate them into the model.

Incomplete Openness So far we have assumed that all countries are perfectly open to each other, i.e. that global ideas invented in any country are accessible in the same way in every country.[12] We now modify that assumption by making openness a matter of degree, in the spirit of McGrattan and Prescott (2009). An $I \times I$ matrix Θ_t represents the degree of openness, each element taking on a value in the interval $[0, 1]$, where the value 1 stands for complete openness and 0 for being completely closed. An arbitrary element (i, k) of Θ_t represents the openness of country i to ideas from country k. We will assume that each jurisdiction is fully open to itself so that $\Theta_t(i, i) = 1$ for each $i = 1, 2, \ldots, I$. Letting Θ_t depend on t allows us to model the gradual opening up of some or all jurisdictions.

It can be shown that all linked countries converge to the same growth rate. Two countries i, k are linked either directly in the sense that $\Theta_t(i, k) > 0$ or indirectly in the sense that there is a chain of non-zero values of Θ_t linking the two, e.g. via country ℓ so that $\Theta_t(i, \ell) > 0$ and $\Theta_t(\ell, k) > 0$.

It is easy to also show that a larger country will have a strictly higher TFP for the same number of ideas per capita, if the degree of openness is strictly less than one. More formally, if $\chi_1 > \chi_2$, $z_{1,t} = z_{2,t}$ and $\widehat{z}_{1,t} = \widehat{z}_{2,t}$, $\Theta_t(1, i) = \Theta_t(2, i)$ and $\Theta_t(i, 1) = \Theta_t(i, 2)$ for all $i > 2$, and $\Theta_t(1, 2) = \Theta_t(2, 1) < 1$, then $TFP_{1,t} > TFP_{2,t}$. This result is similar in flavour to Proposition 1 in Spolaore and Wacziarg (2005).

Population growth In our baseline model, the population of each country is constant; this extension deals with transitory population growth, where each country settles down to a finite population so that the balanced growth path features no population growth at all. We find this the most plausible, as unbounded population growth is eventually inconsistent with the finiteness of planet Earth. In any

[12]It is straightforward to have varying degrees of incomplete openness between locations, rather than countries, but we do not pursue this further here.

case, projections point towards negative global population growth in a few decades, while the population of many developed countries is already shrinking (or would be if not for immigration).

As the population grows, new locations may come into being; we assume that these new locations are "born" with the same average measure of local ideas per capita as old locations in the same jurisdiction. We find this assumption intuitively appealing. The obvious analogy is biological growth. As a body grows, it does so by cell division, resulting in two (nearly) identical copies of the original cell. These copies may then evolve in distinct directions after the division. An interesting and realistic implication of population growth is that the profitability of a global idea relative to a local idea increases.

Exogenous human capital Human capital accumulation may obviously play a role in accounting for cross-country income differences and thus to some extent economic growth. See, for instance, Hendricks and Schoellman (2017). Including a full-blown endogenous human-capital formation decision would go beyond the scope of this paper, but it is straightforward to incorporate exogenous human capital. We thus assume here that human capital per capita is given by the sequence $E(j)$ and that it converges to a finite constant for all j. Human capital is equally important for production and research and enters multiplicatively into the effective labour input.

5 Numerical results

In this section we explore the quantitative implications of our model. Here we are not interested in replicating a particular historic episode, but rather in investigating in more detail the properties of our model. We describe the parameterization as well as our numerical method in Appendix C, available online.

5.1 Scale Effects

Scale effects have been at the center of the debate on endogenous growth models for a long time. In Romer (1990), an economy that is twice as large as another one (but otherwise identical) would enjoy a growth rate that is (more than) twice as large. In Jones (1995a), and thus also Klenow and Rodriguez-Clare (2005), the size of an economy does not matter for economic growth on a BGP, but the population growth rate linearly determines economic growth. In Howitt (2000), it is a bit more difficult to assess the scale properties. Taken at face value, the model implies that a doubling of the population of a country (or the world) has no impact on the growth rate of the economy, while doubling the number of countries would double the growth rate.

In our model, the growth rate of per-capita income on a BGP depends positively, but non-linearly, on the size of the integrated world economy. We thus need to assess quantitatively how much size matters.

If we ignore the indirect effects, that the incentive to innovate increases due to a larger market size, then we can write the scale effects on the growth rate $g(\chi)$ as

$$1 + g(2 \cdot \chi) \approx \left(1 + (1 + g(\chi))^{\frac{\sigma(1-\alpha)}{1-\sigma}} \cdot 2^{(1-\rho)(1-\Lambda)} - 2^{(1-\rho)(1-\Lambda)}\right)^{\frac{1-\sigma}{\sigma(1-\alpha)}}. \qquad (26)$$

Using this formula, we can calculate that a doubling of the world population, keeping everything else equal, would lead to an increase in the annual growth rate from 2.0% to 2.1788%. It turns out that the indirect effects are indeed not that important here, and that the change in the actual growth rate due to a higher research effort per capita is small (2.1906% per annum, once this effect is taken into account). Notice that the approximation via logs, $g(2 \cdot \chi) \approx g(\chi) \cdot 2^{(1-\rho)(1-\Lambda)}$ is highly inaccurate.

If one thinks that research spillovers are larger and research congestion lower, so that $\rho = \Lambda = (1 - m) = 0.25$, then the annual growth rate would increase from 2.0% to 2.4106% after doubling the world population. If, on the other hand, research spillovers are smaller and research congestion higher, setting $\rho = \Lambda = (1 - m) = 0.75$, the increase is very small, from 2.0% to 2.0441%. The scale effects vanish almost completely when $\rho = \Lambda = (1 - m) = 0.9$, where a doubling of the world population increases the growth rate from 2.0% to 2.0070%.[13] In any case, for reasonable pa-

[13]Incorporating the effect of increased incentives to research changes the growth rate to 2.4501%

rameter values the scale effects are not so large that we would reject the model on these grounds. In particular, for any of these sets of parameter values, the implied increases in the growth rate from a doubling of the world population are less than the estimated regression coefficient from the data (see Appendix D, available online). This is actually reassuring, since historically many other factors were at play besides the population increases; for example, average education has increased and the world has become more integrated.

Openness and country-scale effects If countries are not perfectly open to each other, then there are scale effects at the country level. These scale effects of course do not concern growth rates on a BGP, since convergence to a common growth rate takes place for any strictly positive degree of openness, but affect the level of productivity. To assess how much country size matters in this regard, we consider the following exercise. We use the parameter values above, but investigate a world consisting of two countries, one being twice as large as the other, but otherwise identical. We set the degree of openness to each other permanently to 0.5. To ensure comparability with the previous exercise, we keep the world population size at 1.

When we set $\chi_1 = 2/3$ and $\chi_2 = 1/3$, the long-run annual growth rate is 1.9366%, slightly below the 2.0% in a fully integrated economy. The larger economy is naturally more productive due to the scale effects: country 1 is twice as large as country 2 and its output per capita is 6.11% higher. When we set $\chi_1 = 20/21$ and $\chi_2 = 1/21$, the long-run annual growth rate is 1.9908%, and country 1's output per capita is now 13.73% higher than in country 2.[14]

5.2 Effects of population growth

In the previous section, we explored the scale effects from population *size*. We now turn to the scale effects from population *growth*. As in Howitt (2000), the population

with $\rho = \Lambda = (1 - m) = 0.25$, to 2.0456% for $\rho = \Lambda = (1 - m) = 0.75$, and to 2.0071% with $\rho = \Lambda = (1 - m) = 0.9$.

[14]Income differences do not get much larger than this: when country 1 is 2 000 times as large as country 2, the long-run annual growth rate is 1.9999%, and country 1's output per capita is 14.66% higher than in country 2; when population differs by a factor of 20 000, the long-run annual growth rate is 2.0000% and the income per capita is 14.67% higher.

growth rate has a potentially ambiguous effect on economic growth, since it lowers the capital intensity through a capital dilution effect, but increases the incentive to innovate through a market size effect.

Since we assume that all countries converge to a finite population size in the long run, we have to consider transitions in order to study the effects of (differential) population growth. We study the following scenario: there are two countries that are identical except that, for the first two hundred years, their population growth rates differ. In particular, country 1's population remains constant, $\chi_{1,t} = 0.5$ for all t. Meanwhile, country 2's annual population growth rate is 1% at first and then linearly declines to zero, i.e. $n_{2,1} = 1.01^{20} - 1$ and $n_{2,10} = 1.001^{20} - 1$ and $n_{2,t} = 0$ for $t \geq 11$, with $\chi_{2,0} = 0.5$. All other parameters are as in the baseline scenario.

Panel (a) of Figure 1 shows the annualized growth rates of output for country 1 and country 2. Panel (b) of Figure 1 shows the relative output per capita in country 1 relative to country 2.

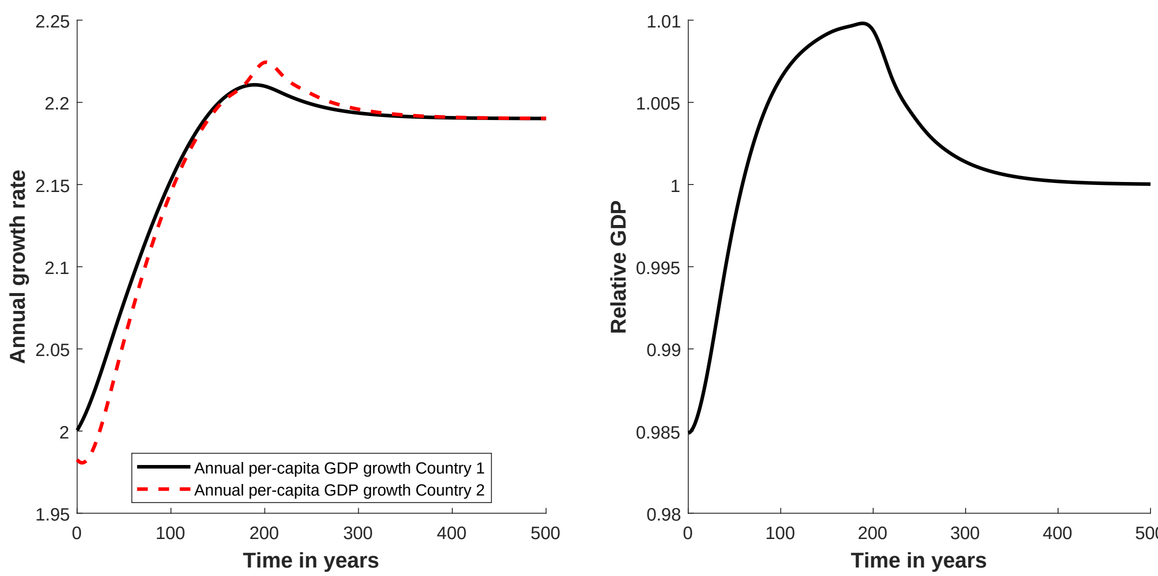

(a) Growth rates (b) Relative output per capita

Figure 1: Effects of population growth

As shown in Panel (a) of Figure 1, country 2 (whose population is growing) experiences slightly lower growth in output per capita in the short run (about 150 years in our model); subsequently, its per-capita output growth briefly exceeds that of country 1, and in the long run growth converges to a common rate. That common growth rate increases to 2.19%, due to scale effects: the population in country 2 roughly triples over time, implying that the world population approximately doubles.

The immediate effect of population growth is higher output in country 2; see Panel (b) of Figure 1, where even before the population has started to increase, people work harder in anticipation of needing more savings in the future. After about 60 years, output per capita in country 2 is lower than in country 1, however, and remains lower for the rest of the transition; when population growth stops after 200 years, it starts to converge to the same level as country 1. While the effects on output (and economic growth) are ambiguous, consumption per capita is always higher in country 1 than in country 2. In quantitative terms, population growth in our model has only small effects on relative productivity growth, output, or consumption.

5.3 Declining research productivity

Bloom et al. (2020) document that research productivity has exhibited a sustained gradual decline in the United States since the 1940s. Specifically, research activity has increased but the growth rate of output per head has remained roughly constant. This appears to be inconsistent with innovation-driven endogenous growth models such as Romer (1990). When allowing for a congestion effect in research, which is a simple modification of Romer (1990), research productivity will of course decline in response to an increase in research activity. However, as Bloom et al. (2020) convincingly argue, this cannot be the only factor at work in the U.S. experience. The point is that a roughly constant rate of innovation, measured as a roughly constant growth rate in total factor productivity (TFP), together with a sustained gradual decline in measured research productivity, is not implied by *any* degree of research congestion except at the limit where we are back to an exogenous growth model. Unless other factors are at work, a sustained ongoing increase in research activity, which is what we observe in the United States since the 1940s, leads to a sustained ongoing increase in the TFP growth rate which we do not observe. All a congestion effect does is to affect the magnitude of this effect. It does not remove it.

Moreover, if we focus only on a model's BGP, there can, by definition, be no other factors at work. Overall, this seems to be a fatal argument against innovation-driven endogenous growth models. Indeed, Bloom et al. (2020) conclude that sustained exponential growth in income per capita can only occur if there is a sustained exponential growth in research activity, suggesting a semi-endogenous growth model, perhaps along the lines of Jones (1995a), as the right one.

Nevertheless, what we argue here is that innovation-driven endogenous growth theory can be salvaged. To see how, note that the possibility remains that some *transitional* phenomenon may give rise to a sustained gradual decline in measured research productivity together with a roughly constant growth in output per head. As it turns out, this apparently remote possibility is not hypothetical at all, but a fairly straightforward implication of our model in two distinct scenarios that both have a pretty obvious empirical basis. The first one is a situation where an initially backward and isolated part of the world gradually joins the rest of the world economy in the sense that it starts taking part in the global sharing of ideas. There is no doubt that this sort of thing happened in the aftermath of the Second World War. After 1945, Japan, South Korea, Taiwan, Singapore and Hong Kong, and later China and Vietnam gradually joined the world economy. In a computational experiment inspired by this experience, we have no trouble generating a scenario where the leading country experiences a protracted decline in research productivity. This is exactly the sort of phenomenon that Bloom et al. (2020) appear to rule out for models in the style of Romer (1990).

The second scenario is one where a subset of the world suddenly improves its ability to produce new ideas, meaning that it is able to produce more research output for a given amount of inputs. This too qualitatively resembles a historical event, namely the scientific and industrial revolutions of the late 18th and early 19th centuries, when Europe and especially the United Kingdom started innovating at rates not seen before or elsewhere. Similarly, the United States became the world leader in innovation in the 20th century. In this scenario too, our model is capable of generating a sustained decline (after an initial surge) in observed research productivity in the rich part of the world. We now explore these two computational experiments in more detail.

Scenario 1 Let there be two (sets of) countries, which, as far as the flow of ideas is concerned, are initially completely closed vis-à-vis each other. Then they gradually open up. Specifically, suppose that, except for initial conditions, countries 1 and 2 are identical, and let the parameters values be the same as in our baseline calibration. Let the initial stock of ideas in country 1 be one, i.e. $Z_{1,0} = 1$, and let the initial capital per capita $a_{1,0}$ be equal to the long-run value from the baseline case. Now calibrate the initial stock of ideas in country 2, $Z_{2,0}$, so that country 2's productivity in time period $t = -1$ is one fourth of that in country 1, assuming that each is at that time on a BGP, not expecting the opening.

Meanwhile, the initial capital stock per capita in country 2, $a_{2,0}$, is set such that the capital/output ratio is the same in both countries in period $t = -1$ under the assumption that saving in period $t = -1$ is consistent with remaining on a BGP. In time period $t = 0$, the countries are still closed and the off-diagonal elements of $\Theta_{t=0}$ are equal to zero; the only change relative to period $t = -1$ is that agents now *know* that opening is going to happen. Then the countries gradually open up to each other: in period $t = 1$ (after 20 years) the degree of openness is 0.25, in $t = 2$ it is 0.5, in $t = 3$ it is 0.75, and from $t = 4$ on it is 1.

In Figure 2, we can see that observed research productivity in the initially rich country indeed declines, and remains depressed for a long time, before it finally reaches a higher level than before the new (initially poorer) country joined the world marketplace of ideas. Here we use $\frac{\hat{Z}_{i,t}/Z_{i,t}}{\chi_i h_{i,t}}$, the growth rate of ideas per unit of research effort, as a measure of observed research productivity; fundamental research productivity η remains unchanged by assumption in this scenario. At the same time, initial TFP growth remains roughly constant and then converges to a higher rate, due to the larger connected world population engaged in research.

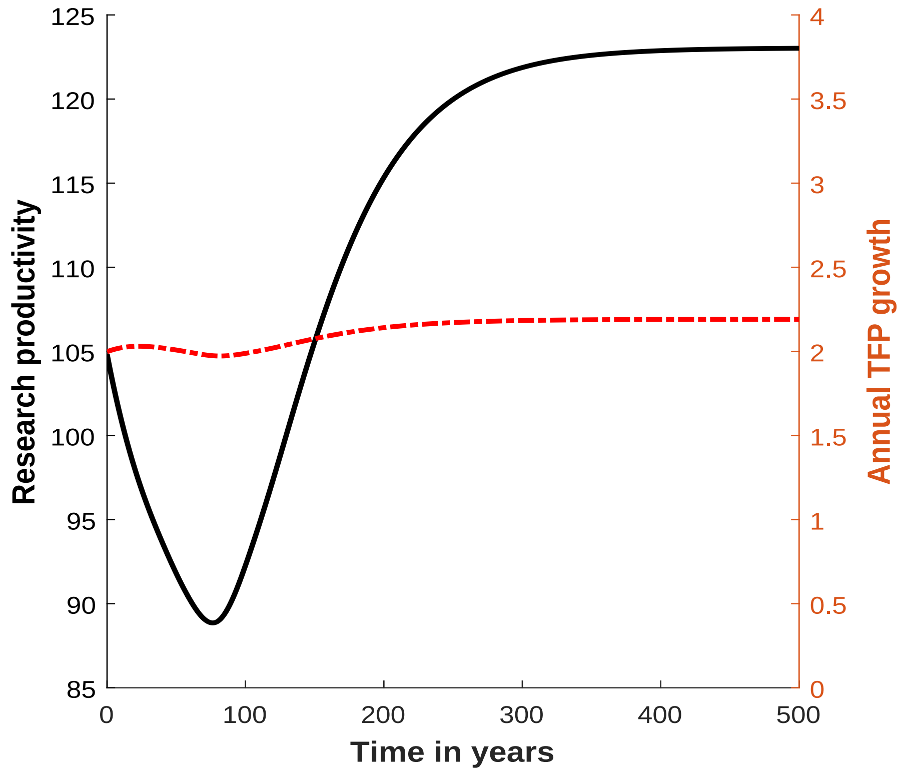

Figure 2: Research productivity over time in country 1.

What could explain this transitory, but very long-lived, decline in observed research productivity in our model? We rewrite Equation (17) at the country level:

$$\frac{\widehat{z}_{1,t}}{z_{1,t}} = (1-m)^{\rho}\left(\eta_1 h_{1,t}\right)^{1-\rho\Lambda}\left(\frac{mZ^{W}_{1,t}/z_{1,t}}{\left[H^{W}_{1,t}\right]^{\Lambda}}\right)^{1-\rho}. \tag{27}$$

Upon opening, the initially advanced country 1 has a large advantage in ideas; country 2 therefore initially adds very little to $Z^W_{1,t}$, the stock of ideas (shoulders to stand on) available to researchers in country 1. At the same time, country 1 researchers start competing with those from country 2 in the development of new ideas and country 2 contributes substantially to $H^W_{1,t}$, the measure of researchers competing to develop new ideas (stepping on toes). Due to the *global* congestion externality, researchers in country 1 become *less* productive as country 2 joins the world marketplace of ideas; the opposite happens in country 1.

The quantitative extent of this effect naturally depend on the parameter Λ; as it goes to zero, the congestion externality disappears and research productivity would initially increase in both countries due to the sharing of ideas between countries. Meanwhile, the protractedness of the transition depends on the parameter ρ which governs the strength of research spillovers. The smaller ρ, and hence the stronger the spillovers are, the more rapid the transition is.

Notice also that the parameter ρ governs not only positive spillovers across countries but also congestion spillovers. Indeed the congestion externality across countries is governed by the factor $\Lambda(1-\rho)$ rather than just Λ itself. A smaller ρ means stronger congestion effects across countries. With $\rho = \Lambda = 0.25$, for example, observed research productivity in country 1 initially declines by about the same amount as with $\rho = \Lambda = 0.75$ (because $\Lambda(1-\rho)$ remains constant) but recovers faster (because ρ is smaller), though it still remains below the value at $t = 0$ for 80–100 years.

Scenario 2 Country 1 experiences a sudden and permanent jump in fundamental research productivity η. Let there be two (sets of) countries that are fully integrated as far as the flow of ideas is concerned and initially ($t < 0$) identical. A $t = 0$, country 1 starts to enjoy a higher fundamental research productivity η. More precisely, we set all elements of Θ to one and follow the baseline calibration, except that, from $t = 0$ onwards, we set $\eta_1 = 2 \cdot \eta_2$.

Figure 3 shows that, in our second scenario, observed research productivity in country 1, after an initial surge, declines for an extended period of time. As above, we use $\frac{\hat{Z}_{i,t}/Z_{i,t}}{\chi_i h_{i,t}}$, the growth rate of ideas per unit of research effort, as our measure of observed research productivity; after the initial jump in country 1, fundamental research productivity remains constant.

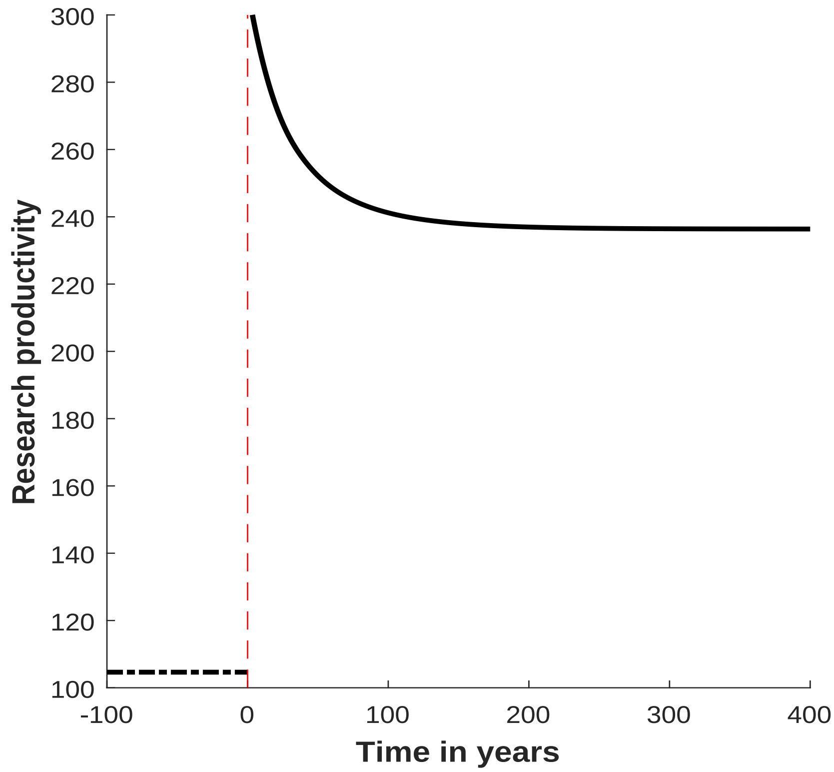

Figure 3: Research productivity over time in country 1.

To understand what happens in this scenario, we notice that when fundamental research productivity is higher in country 1, then researchers in country 1 can initially build on the ideas from both countries to an equal extent (since they are initially the same: $Z_{1,0} = Z_{2,0}$). As time goes on, however, country 1 becomes ever more technologically advanced relative to country 2, and hence it benefits to a lesser extent (relatively speaking) from the ideas in country 2 in its creation of new ideas. Formally, Equation (27) implies that $\widehat{Z}_{1,t}/Z_{1,t}$ increases in the ratio $Z_{2,t}/Z_{1,t}$. The observed research productivity therefore has to decline over time: the rate at which a single country can push ahead of other countries is diminishing in the technological distance between the advanced country and the rest of the world.

5.4 Catchup Growth

The previous section may give the impression that transitions are incredibly drawn out and that it would be impossible for a country to catch up quickly. While it is true that it takes a long time (at least with our baseline parameter values) for technology levels to roughly converge, output per capita can grow very quickly in a country that is technologically backward but then opens up to the rest of the world. We now illustrate that in a computational exercise. For this exercise, we use the same setup as in the first scenario of the previous section, except that country 2, the laggard country, now has a much smaller population than country 1; specifically $\chi_1 = 1$ and $\chi_2 = 0.05$. This allows us to focus on the unfolding events in country 2; country 1 is barely affected by the addition of the rather small country 2 to the world economy.

Panel (a) of Figure 4 shows the annualized growth rates (in percent) of output in this scenario.[15] Before opening, the advanced economy experienced an annual growth rate of 2%, while the developing economy grew at 1.26% a year. When opening, the developing economy initially grows at a pace of more than 7% a year, and then gradually slows down to the common rate of growth, shared with the advanced economy, which has slightly increased to 2.01%.

Panel (b) of Figure 4 shows the relative output per capita in the advanced compared to the emerging economy. Evidently catchup happens quite rapidly, though not as rapidly as in a Solow (1956) model. After one period (20 years), the ratio has been reduced by 63 percent. Consider, for comparison, a Solow model where two countries differ only in their initial capital stock. While the leading country is assumed to be on its BGP, the initial capital stock in the lagging country is set so that the ratio of leading/lagging country output is initially the same as in the computational experiment of the present section (6.81). Then if the capital share is 1/3, the ratio is reduced by 76 percent after 20 years. Thus our model has, broadly speaking, similar implications for rates of catchup as the Solow model, except that in our case the rate of catchup is governed not only by the capital share but by ρ, the parameter determining the strength of cross-country spillovers.

[15]Given that the model period length is 20 years, the formula is $\left[\left(Y_{t+1,2}/Y_{t,2} - 1\right)^{1/20} - 1\right] \cdot 100$.

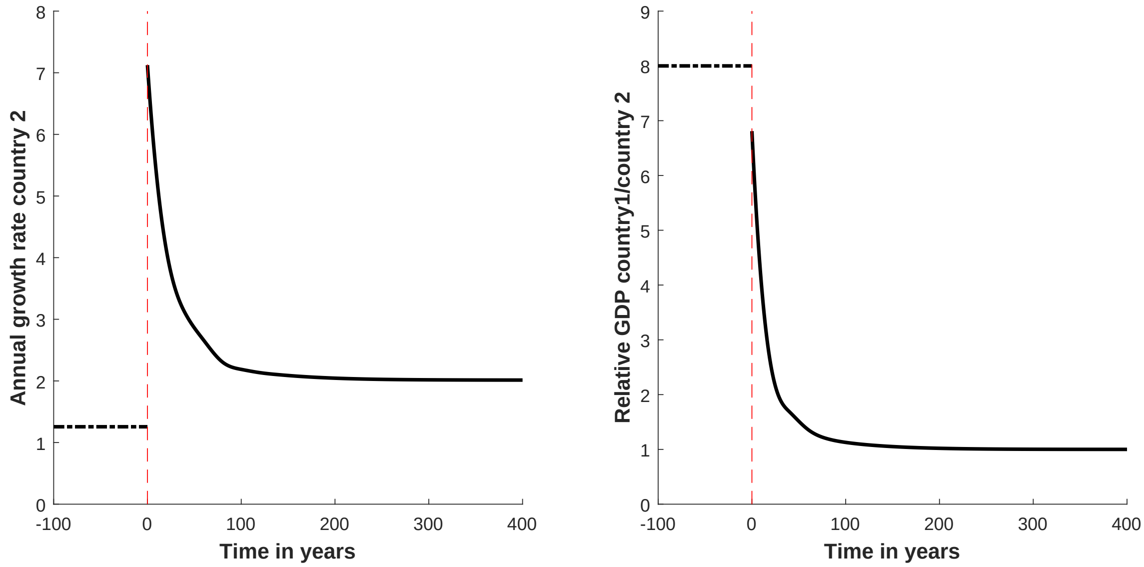

(a) Growth rates (b) Relative output per capita

Figure 4: Catch-up growth

To understand the determinants of this rapid convergence, we provide a decomposition of the growth rate of output in country 2:

$$\frac{Y_{2,t+1}}{Y_{2,t}} = \left(\frac{\mathbf{Z}_{2,t+1}}{\mathbf{Z}_{2,t}}\right)^{\frac{1-\sigma}{\sigma(1-\alpha)}} \cdot \left(\frac{K_{2,t+1}/Y_{2,t+1}}{K_{2,t}/Y_{2,t}}\right)^{\frac{\alpha}{1-\alpha}} \cdot \left(\frac{L_{2,t+1}}{L_{2,t}}\right) \cdot$$

$$\left[\frac{1+\left(\widehat{\mathbf{Z}}_{2,t+1}/\mathbf{Z}_{2,t+1}\right)\sigma^{\sigma/(1-\sigma)}}{1+\left(\widehat{\mathbf{Z}}_{2,t}/\mathbf{Z}_{2,t}\right)\sigma^{\sigma/(1-\sigma)}}\right]^{\frac{1}{\sigma(1-\alpha)}} \left[\frac{1+\left(\widehat{\mathbf{Z}}_{2,t}/\mathbf{Z}_{2,t}\right)\sigma^{1/(1-\sigma)}}{1+\left(\widehat{\mathbf{Z}}_{2,t+1}/\mathbf{Z}_{2,t+1}\right)\sigma^{1/(1-\sigma)}}\right]^{\frac{1}{1-\alpha}}. \quad (28)$$

The right-hand side of this equation consists of four factors separated by mutiplication symbols ($\cdot$). The first factor is the growth in the stock of available ideas. (Later, we will decompose this factor into the development of new ideas and the adoption of existing ideas from abroad.) The second factor is the growth in the capital-output ratio.[16] The third is the growth of production labour. The fourth and final factor is the growth rate in the degree of monopolization of the intermediate goods sector, which in turn of course depends on the change in the relative quantities of old

[16] In our model, the capital-labour ratio grows without bound, so this is not a very useful concept in this context. The capital-output ratio, on the other hand, remains constant on a BGP.

and new ideas. On a balanced-growth path, all growth ultimately stems from the first factor, the growth in the stock of available ideas, and in particular the development of new ideas. All other factors will be constant on a BGP, though they can be significant during the transition.

To get at the concept of the adoption of ideas in our framework, we look deeper into the factor $\mathbf{Z}_{2,t+1}/\mathbf{Z}_{2,t}$:

$$
\begin{gathered}
\mathbf{Z}_{2,t+1} - \mathbf{Z}_{2,t} = \\
(1-m)\left(Z_{2,t+1} - Z_{2,t}\right)/\chi_2 + m\left(Z_{2,t+1} - Z_{2,t}\right) + m\left(\Theta_{2,1,t+1} Z_{1,t+1} - \Theta_{2,1,t} Z_{1,t}\right) = \\
(1-m)\widehat{Z}_{2,t}/\chi_2 + m\widehat{Z}_{2,t} + m\Theta_{2,1,t+1}\left(Z_{1,t+1} - Z_{1,t}\right) + m\left(\Theta_{2,1,t+1} Z_{1,t} - \Theta_{2,1,t} Z_{1,t}\right) = \\
(1-m)\widehat{Z}_{2,t}/\chi_2 + m\widehat{Z}_{2,t} + m\Theta_{2,1,t+1}\widehat{Z}_{1,t} + m\left(\Theta_{2,1,t+1} - \Theta_{2,1,t}\right) Z_{1,t+1} = \\
\widehat{\mathbf{Z}}_{2,t} + m\left(\Theta_{2,1,t+1} - \Theta_{2,1,t}\right) Z_{1,t+1}.
\end{gathered}
$$

We can rearrange this so that

$$
\frac{\mathbf{Z}_{2,t+1}}{\mathbf{Z}_{2,t}} = 1 + \underbrace{\frac{\widehat{\mathbf{Z}}_{2,t}}{\mathbf{Z}_{2,t}}}_{\text{Development of ideas}} + \underbrace{m\left(\Theta_{2,1,t+1} - \Theta_{2,1,t}\right)\frac{Z_{1,t+1}}{\mathbf{Z}_{2,t}}}_{\text{Adoption of ideas}}. \tag{29}
$$

Growth in the stock of available ideas comes from new ideas, local and global ideas developed at home and global ideas developed abroad, and the adoption of existing ideas from abroad. The stock of global ideas abroad in time period $t+1$ is $mZ_{1,t+1}$ and a fraction $\Theta_{2,1,t+1} - \Theta_{2,1,t}$ are being newly adopted in the home country. For decomposition purposes, we calculate the fraction of the growth in available ideas that comes from the development of ideas vs. the adoption of ideas, and multiply it by the contribution of the growth in $\mathbf{Z}_{2,t+1}/\mathbf{Z}_{2,t}$ to output growth.

Table 1: Catchup growth: decomposition by source

Period	Total	Capital	Labour	Monopolization	Adoption	Innovation
0	7.124	−2.116	0.122	0.817	8.277	0.023
1	4.136	1.200	−0.025	0.036	1.597	1.328
2	3.105	0.528	−0.030	0.038	0.936	1.634
3	2.659	0.228	−0.034	0.024	0.641	1.800
4	2.305	0.187	−0.034	0.014	0.000	2.139
5	2.186	0.076	−0.025	−0.015	0.000	2.151
6	2.133	0.043	−0.018	−0.027	0.000	2.135
7	2.101	0.030	−0.012	−0.026	0.000	2.110
⋮	⋮	⋮	⋮	⋮	⋮	⋮
30	2.013	0.000	0.000	0.000	0.000	2.013

Note: Growth rates are reported in percent per year. Every period is 20 years.

As Table 1 shows, the annualized growth rate of output in the first twenty years after opening is a blistering 7.12%, then declines to a still sizable 4.14% in the following twenty years, and then gradually declines further. Capital accumulation in the first twenty years does not keep up with the vast increase in the mass of available intermediate goods, so the decline in the capital-output ratio is at first a net negative contributor to growth. If the capital-output ratio had remained constant in the first 20 years after opening, annual output growth would have been 2.12% higher. Subsequently, the capital-output ratio grows back to its long-run level and contributes significantly to output growth. Production labour expands slightly in the first twenty years after opening, contributing 0.12% to growth, but from there on has a small negative impact on growth as the labour supply returns over time to its long-run level. The adoption of existing ideas is key to the fast catchup growth; in the first twenty years after opening, output would have grown 8.28% per year from the adoption of existing ideas alone. In the second twenty years, the contribution is still 1.6%, then down to 0.94%, and finally 0.64%, after which the home country is perfectly open to the foreign country, so no further adoption of ideas can take place.

The development of new ideas barely contributes to growth at all at first, 0.02% in the first 20 years after opening, but then gradually contributes more, up to 2.14% 100 years after opening. After that (not shown in the table), new ideas contribute 2.15% to growth 120 years after opening, and then the contribution gradually declines to its new long-run level of 2.01%.[17] The change in the degree of monopolization of the economy contributes 0.82% to growth in the first twenty years after opening, but after that a rather negligible amount.

In summary, when the economy opens up, two things happen. First, researchers can use foreign ideas, and compete with foreign researchers (this is what we examined in the previous section). Second, global ideas from abroad become gradually available and the corresponding varieties can be produced; it is this second effect that leads to the large and sudden increase of output after opening.[18] The contribution to the growth rate from new ideas in the developing economy at first remains below the growth rate of the advanced economy and is only slightly more elevated after about one hundred years (before converging to the same value). The rapid convergence in terms of relative output per capita is thus mainly due to the adoption of existing ideas.

6 Conclusion

We have constructed a multi-country model of a world economy where the ultimate source of long-run growth in productivity is incentive-driven innovation. We have

[17]New ideas are of course being developed in the first twenty years after opening, but 99.7% of the change in available ideas between periods 0 and 1 come from the adoption of ideas. Before the adoption of ideas, new ideas contributed 1.26% to output growth (and thereby the entirety of growth). However, since intermediate goods are (imperfect) substitutes, the measure of newly developed ideas in period 0 does not contribute much to output growth because there are so many newly adopted ideas.

[18]This effect is, in principle, also taking place in the advanced economy. However, given the differences in size and in the initial stock of ideas, it is negligible in practice. The advanced economy is only very slightly affected by the reciprocal opening of the two economies. It is interesting to note, perhaps, that it grows slightly more slowly after the opening, because of the reduced research productivity highlighted in the previous section.

shown that it is broadly compatible with a large set of facts.

In our model, the long-run growth rate of productivity responds to policy. In consequence, a lot more is at stake in setting policy than what an exogenous or semi-endogenous growth model would imply. Arguably, it is this feature that made endogenous growth models so compelling when they first emerged. An obvious follow-up is to study the effects of policy on growth in our model, and to analyze optimal policy setting. For example, our model may be a good starting point for studying international patent policies.

The framework that we have established in this paper has the potential to be used to address many more questions. For instance, we may ask: under what circumstances might a country not contribute to the world frontier? Consequently, can we better understand how a world economy might evolve so that it consists of leaders and followers? This question can be addressed in a modified version of our model, where innovators choose how much effort to devote to the development of global vs. local ideas. We could also conduct a detailed growth-accounting study for specific historical episodes, such as the rise of the four "Asian Tigers," and more recently, of China. Another possibility is to tap more into the potential richness of location heterogeneity and study agglomeration and migration when locations differ in size (and other dimensions).

Throughout the paper, we have assumed some degree of mobility across borders of ideas, but no mobility of capital. This raises the question of what the implications are of capital mobility in our framework. Given capital mobility, we may also ask what the implications are of tax competition, i.e. the decentralized setting of policy in each jurisdiction rather than at the global level. How does policy differ depending on whether it is set by a benevolent ruler of the world as opposed to decision-makers in each jurisdiction? How does welfare compare in the two scenarios? We leave these important questions for future work.

References

Aghion, P. and P. Howitt (1992, March). A model of growth through creative destruction. *Econometrica 60*(2), 323–351.

Alesina, A., E. Spolaore, and R. Wacziarg (2005). Trade, growth and the size of countries. In P. Aghion and S. Durlauf (Eds.), *Handbook of Economic Growth* (1 ed.), Volume 1, Part B, Chapter 23, pp. 1499–1542. Elsevier.

Barro, R. J. and X. Sala i Martín (1997). Technological diffusion, convergence, and growth. *Journal of Economic Growth 2*(1), 1–26.

Barro, R. J. and X. Sala i Martín (2004). *Economic Growth* (2nd ed.). MIT Press.

Bloom, N., C. I. Jones, J. Van Reenen, and M. Webb (2020, April). Are ideas getting harder to find? *American Economic Review 110*(4), 1104–44.

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

Cai, J., N. Li, and A. M. Santacreu (2022, January). Knowledge diffusion, trade, and innovation across countries and sectors. *American Economic Journal: Macroeconomics 14*(1), 104–45.

Eaton, J. and S. Kortum (1996). Trade in ideas patenting and productivity in the OECD. *Journal of International Economics 40*(3-4), 251–278.

Eaton, J. and S. Kortum (1999). International technology diffusion: Theory and measurement. *International Economic Review 40*(3), 537–70.

Galor, O. (2010). Economic growth in the very long run. In B. L. Durlauf S.N. (Ed.), *Economic Growth. The New Palgrave Economics Collection.* Palgrave Macmillan, London.

Gross, T. and P. Klein (2022, May). Optimal tax policy and endogenous growth through innovation. *Journal of Public Economics 209*.

Hendricks, L. and T. Schoellman (2017, 12). Human capital and development accounting: New evidence from wage gains at migration. *The Quarterly Journal of Economics 133*(2), 665–700.

Howitt, P. (2000, September). Endogenous growth and cross-country income differences. *American Economic Review 90*(4), 829–846.

Hsieh, C.-T., P. J. Klenow, and I. Nath (2023). A global view of creative destruction. *Journal of Political Economy Macroeconomics 1*(2), 243–275.

Jones, C. I. (1995a). R&D-based models of economic growth. *Journal of Political Economy 103*(4), 759–784.

Jones, C. I. (1995b). Time series tests of endogenous growth models. *Quarterly Journal of Economics 110*(2), 495–525.

Jones, C. I. (2002). Sources of U.S. economic growth in a world of ideas. *The American Economic Review 92*(1), 220–239.

Klenow, P. and A. Rodriguez-Clare (2005). Externalities and growth. In P. Aghion and S. Durlauf (Eds.), *Handbook of Economic Growth* (1 ed.), Volume 1, Part A, Chapter 11, pp. 817–861. Elsevier.

McGrattan, E. and E. Prescott (2009). Openness, technology capital, and development. *Journal of Economic Theory 144*(6), 2454–2476.

McGrattan, E. R. and E. C. Prescott (2010). Unmeasured investment and the puzzling U.S. boom in the 1990s. *American Economic Journal: Macroeconomics 2*, 88–123.

Peretto, P. F. (2018). Robust endogenous growth. *European Economic Review 108*, 49–77.

Perla, J., C. Tonetti, and M. Waugh (2021). Equilibrium technology diffusion, trade, and growth. *American Economic Review 111*(1), 73–128.

Rebelo, S. (1991, June). Long-run policy analysis and long-run growth. *Journal of Political Economy 99*(3), 500.

Rivera-Batiz, L. A. and P. M. Romer (1991). Economic Integration and Endogenous Growth. *The Quarterly Journal of Economics 106*(2), 531–555.

Romer, P. M. (1990, October). Endogenous technological change. *Journal of Political Economy 98*(5), 71–102.

Sampson, T. (2023). Technology gaps, trade, and income. *American Economic Review 113*(2), 472–513.

Solow, R. M. (1956). A contribution to the theory of economic growth. *The Quarterly Journal of Economics 70*(1), 65–94.

Spolaore, E. and R. Wacziarg (2005, December). Borders and Growth. *Journal of Economic Growth 10*(4), 331–386.

Tollefson, J. (2018). China declared largest source of research articles. *Nature 553*.

Trouvain, F. (2023). Technology adoption, innovation, and inequality in a global world. Unpublished manuscript.