Uncertainty, Openness to Novelty, and Economic Growth

Maren Bartels* Johannes Binswanger† Manuel Oechslin‡

February 14, 2023

[Draft - Please do not cite or circulate]

Abstract

Successful innovations are a key driver of long-run economic growth. In practice, potential innovations come with a great deal of uncertainty, i.e., a dearth of objective information on the likelihood of eventual success or failure. We present a tractable growth model in which entrepreneurs are characterized by their openness to novelty. Greater openness makes it more likely that, against the backdrop of uncertainty, potential innovations are checked out; but greater openness can also lead to exuberance, to negative signs being ignored—and then to misallocation and crisis-induced paralysis. We analyze this trade-off and show that it implies a hump-shaped relationship between openness and long-run growth. The calibrated model predicts that over a significant part of the range the negative effect of openness dominates, a result we show to be consistent with the empirical pattern. On the other hand, the calibrated model suggests that heterogeneity in entrepreneurial openness to novelty helps growth. The magnitude of the effect is sizable.

JEL classification: O40, E71

Keywords: Innovation, uncertainty, openness to novelty, paralysis, economic growth

*University of Lucerne, Department of Economics, Switzerland: maren.bartels@unilu.ch.

†University of St. Gallen, Department of Economics, Switzerland: johannes.binswanger@unisg.ch.

‡*Corresponding author.* University of Lucerne, Department of Economics, Frohburgstrasse 3, 6002 Lucerne, Switzerland: manuel.oechslin@unilu.ch.

"*First, technology has to be invented or adopted. Human societies vary in lots of independent factors affecting their openness to innovation.*"
Jared Diamond, 2003.

1 Introduction

One of growth theory's key insight is that technological progress is the ultimate source of long-run economic development. Over the past four decades, the new (i.e., endogenous) growth theory has contributed to a much better understanding of innovation by private entrepreneurs, the main driver of technological progress (see, e.g., Jones 2019). Yet, until now, most works have abstracted from an important phenomenon that is inextricably connected to innovation and technological progress: uncertainty. When a potential innovation arrives, it is often impossible to come up with objective probabilities for different possible outcomes in terms of, e.g., eventual commercial success or failure, time to commercial success, or where exactly in the value chain the benefits will accrue. Goldfarb and Kirsch (2019), an extensive study of major technological innovations over the past 180 years, testifies to this.

Abstracting from uncertainty, while justified for the sake of parsimony, comes at a cost. Almost by construction, it eliminates what in practice are frequent companions of innovation—such as subjective prior beliefs that receive little discipline from facts but instead are colored by personality traits and past experiences; or progressions of exuberance, misalloction, disappointment, and possibly paralysis, sometimes summarized by the term "hype cycle".[1] Yet experience-colored subjective priors and hype cycles are not just innocuous by-products of uncertain innovations that can be safely ignored. Arguably, those phenomena will have feedback effects on technological progress and the pattern of economic development. Therefore, abstracting from uncertainty in the realm of technological innovation obscures potentially important determinants of long-run prosperity. This is the topic of the present paper.

More specifically, we are interested in the role of entrepreneurial openness to novelty in the context of innovation-driven economic growth, where innovation is understood in a broad sense that also includes innovations that lack a physical existence (such as novel financial technologies, e.g., mortgage securitization in the 1990s). If innovations tend to come with uncertainty, and thus with room for subjective priors, entrepreneurs' intrinsic attitudes towards innovation, anchored in the respective cultural background, are likely to matter (e.g., Guiso et al. 2006). But how exactly openness to novelty impacts technological progress is unclear. On the one

[1]See, e.g., Malmendier (2021) for a recent overview of the role of subjectivity and experience effects in (management) decision making. See, e.g., Dedehayir and Steinert (2016) on innovation hype cycles.

hand, openness makes it more likely that potential innovations are checked out, that discoveries are being made. On the other hand, lack of skepticism enables exuberance, leads to negative signs being ignored—and then to misallocation, disappointment, and possibly crisis-induced paralysis. Sketched in bold strokes, an economy with more open entrepreneurs is one in which economic dynamism alternates with crashes and stagnation; an economy with more skeptical entrepreneurs grows more evenly, produces fewer hype cycles, but more often misses out on productive innovations. Which one of the two economies grows faster? Economic history offers a first clue. In the 18th century, Europe started to pull ahead of China, the erstwhile leader, economically. One hypothesis says that this divergence reflects that, with useful knowledge more plentiful and sciences advancing, Europe's increasing openness to new ideas became an important advantage (see, e.g., Mokyr 1990, 2017, p. 188 and p. 300).

In the modern world, though, there seems to be no lack of openness to novelty. In fact, a first look at the data reveals a surprising pattern. Using cross-sectional data, we document a robust negative correlation between the growth rate of GDP p.c. and a new empirical proxy for openness to novelty that is based on World Value Survey data (see Section 2, Figure 1). The theoretical model set out below allows us to describe and analyze mechanisms—positive as well as negative—that link general entrepreneurial openness to novelty to long-run growth. The formal analysis identifies parameter constellations under which the relative strength of the mechanisms changes, giving rise to a hump-shaped relationship. Simulations of the model, which are based on parameters calibrated to match empirical moments in historical data, locate the "peak" close to the lower end of the plausible range of the openness parameter. The simulations thus predict that over a broad range the relationship between long-run growth and general openness to novelty is negative, consistent with the empirical pattern. In terms of magnitude, we find that going from the growth-maximizing level of openness to the upper end of the plausible range lowers annual GDP p.c. growth by a quarter of a percentage point. Interestingly, however, further simulations suggest a clear positive impact of heterogeneity in openness to novelty on growth. So, under uncertainty, belief heterogeneity—as is observed in many realistic situations (e.g., Gilboa et al. 2014)—is an advantage.

In the model, each period sees the exogenous arrival of what may prove to be a major innovation—or a failure.[2] The a priori probability that it is the former is unknown. However, entrepreneurs adopt an individual subjective prior belief about this probability. The individual prior is composed of a common belief anchor, reflecting the general attitude (the extent

[2] A major innovation is understood as a broadly applicable general-purpose technology (e.g., online sales and distribution) that presents an opportunity to a non-negligible share of all firms. According to the literature, a plausible period length is ten years (e.g., Goldfarb and Kirsch 2019; Dedehayir and Steinert 2016).

of the hype) towards the potential innovation, and an individual component that makes an entrepreneur more or less open relative to the general attitude. The belief anchor is a random variable that in practice is affected by random forces such as the (non-)appearance of a contagious narrative (e.g., Binswanger and Oechslin 2021; Shiller 2017, 2019). Its mean reflects general entrepreneurial openness to novelty. But the belief anchor is not immune to recent macroeconomic events. If in the previous period a hype, i.e., an omnipresence of "exuberant" priors among entrepreneurs, has ended in misallocation and disappointment, the current period sees a lower belief anchor, all else equal. This is to capture that in practice such a negative experience may lead to temporarily more conservative priors (e.g., Malmendier and Tate 2011; Dittmar and Duchin 2016; Guiso et al. 2018), not least by raising the salience of the downsides (e.g., Bordalo et al. 2012) of experimenting with uncertain innovations.

After having adopted individual subjective priors, all entrepreneurs provisionally decide on whether or not to allocate resources to the potential innovation, a new production method. Later on, those who have done so observe a noisy signal about the new method (success vs. failure) and then get a chance to reallocate the resources to the tested method, albeit at a loss. If there are entrepreneurs who continue with the new method, its quality is publicly revealed and, if a success, it will become next period's tested method. This set of assumptions assembles in one single model several aspects of the process of technological progress that so far have been considered separately. The model permits a systematic analysis of how the general openness, as well as heterogeneity in openness, relate to long-run economic growth. At a basic level, we find that openness to novelty plays an ambiguous role. On the positive side, greater openness lifts the chance that at least some entrepreneurs experiment with, and learn about, new methods—a necessary condition for technological progress. On the negative side, openness promotes the adoption of exuberant priors, i.e., priors that are sufficiently strong to make entrepreneurs continue with a new method even when facing a negative signal; as a result, there is a higher chance of misallocation and disappointment—and thus a heightened risk of short-run output losses and paralyzed experimentation in the future. Together with the theoretical analysis, the simulations clarify the relationship between the two mechanisms—and stress the importance of distinguishing between general openness and heterogeneity in openness.

This paper combines insights from two recent strands of literature, the literature on cultural influences on prior beliefs (e.g., Guiso et al. 2006; Guiso et al. 2008, Mokyr 2017) and the literature on how priors are impacted by experience effects (e.g., Malmendier and Tate 2011; Dittmar and Duchin 2016; Guiso et al. 2018; Malmendier 2021). Consistent with the former, we assume that actors' prior beliefs have a cultural component that gives them a measure of persistence. Consistent with the later, we assume that in the aftermath of negative experiences actors may temporarily pivot to more conservative priors. In terms of focus, our paper joins a series of contributions on the effect of culture—through beliefs, values, and preferences—on long-run economic development (e.g., Doepke and Zilibotti 2008; Galor and Özak 2016; Gorodnichenko and Roland 2017; Sunde et al. 2022).[3] Our attention does not lie on the patterns of transition from Malthusian period to the modern growth era, but rather on how (one aspect of) culture affects the growth rate of productivity in the modern era, an interest we share in particular with Gorodnichenko and Roland (2017) and Sunde et al. (2022). However, while Gorodnichenko and Roland (2017) explore the importance of individualism for innovation incentives (and Sunde et al. 2022 work with a human capital externality), we explore the impact of openness to novelty on the process by which exogenously arriving, uncertain innovations are checked out—and finally accepted or rejected in the practice of business.

By relaxing the common prior assumption, the paper at hand also connects to a recent literature that explores the implications of disagreeing, heterogeneous (prior) beliefs for, among other things, welfare analysis (e.g., Gilboa et al. 2014) or asset pricing and wealth dynamics (e.g., Cao 2018; Borovička 2020). To this list, we add the topic of long-run growth. In fact, the role of heterogeneity—or diversity—in long-run growth is another research topic that as attracted much attention lately. The literature considers diversity along various dimensions, among them ethnic (e.g., Alesina et al. 2003; Montalvo and Reynal-Querol 2021), genetic (e.g., Ashraf and Galor 2013), and birthplace (e.g., Alesina et al. 2016). In the literature, an important argument is that diversity may advance economic development by expanding the production possibility frontier. Provided that greater, e.g., genetic or birthplace-related diversity translates into greater belief heterogeneity, this paper describes a channel by which such diversity *accelerates* the expansion of the production possibility frontier.

The rest of the paper is organized as follows. The next section presents an empirical proxy for openness to novelty and explores its relationship with the growth rate of GDP p.c. Section 3 introduces the theoretical model, whose intra-temporal equilibrium is analyzed in Section 4. Abstracting from belief heterogeneity, Section 5 considers the openness-growth nexus from a theoretical as well as quantitative perspective. Section 6 then turns to the impact of belief heterogeneity on long-run economic growth. Section 7, finally, concludes.

[3] Falk et al. (2018), introducing the Global Preference Survey, provide an overview of how various preferences (among them for time and risk) correlate with economic outcomes such as GDP p.c. and total factor productivity.

2 Empirical Patterns

Before turning to a theoretical analysis of the relationship between openness to novelty and long-run growth, it is worthwhile to take a first look at the data. This section offers cross-sectional correlations, using a novel empirical measure of general openness to novelty extracted from the World Value Survey (WVS).[4] To construct our empirical measure, we make use of two particular WVS items that originate from the Portrait Values Questionnaire (Schwartz et al. 2001). These items present a short portrait of a person, describing that person's attitude. Respondents subsequently assess their own similarity with the described person on a five-point scale that ranges from 1 (not at all like me) to 6 (very much like me). The first item gives the following portrait: "It is important to this person to think up new ideas and be creative; to do things one's own way". Clearly, respondents that recognize a higher degree of similarity with this portrait tend to show a more open attitude towards novelty. The portrait described by the second item is as follows: "Adventure and taking risks are important to this person; to have an exciting life". The purpose of this item is to illicit a respondent's attitude towards risk and uncertainty, an important determinate of openness to novelty. Again, respondents that recognize a higher degree of similarity with this portrait tend to show a more open attitude towards novelty. To construct our measure of general openness to novelty, we first calculate at the country level the weighted average response per item (where the weights correct for imperfections in the sample representation of population shares); we then average across the two items in order to obtain a single country-level observation (the correlation between the two item averages is strong, with a Pearson correlation coefficient of 0.59).

The two WVS items we draw on are included in two (successive) survey waves only, 2005-2009 and 2010-2014. Since, arguably, openness to novelty matters for growth primarily in the long run, we average the two observations to obtain one observation per country for the entire ten-year period from 2005 to 2014. In cases where only one observation is available (56% of cases), we treat this observation as the ten-year average (justified by the fact that, as one would expect, our measure of openness to novelty is highly persistent). We capture long-run economic growth by the average annual growth rate of GDP per capita (p.c.) in the 2005-14 period and rely on GDP p.c. (PPP, constant 2011 international $) data from the World Bank's World Development Indicators database. Our sample counts 75 countries.

Figure 1 shows a partial residual plot of the relationship between general openness to novelty

[4] While the WVS works with representative population samples, our model economy (Section 3) focuses on openness to novelty among entrepreneurs. However, it is highly plausible that there is a strong positive correlation between the society-wide level of openness to novelty and such openness in the subgroup of entrepreneurs.

Figure 1: Openness to novelty and growth, all countries

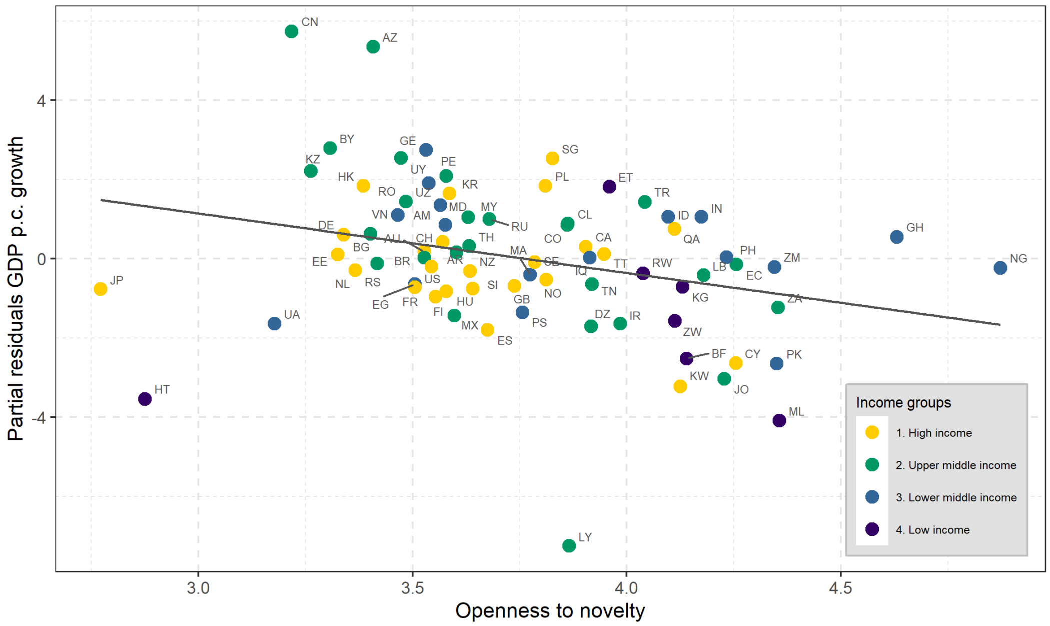

Notes. This is a partial residual plot. The underlying OLS model regresses average real GDP p.c. growth over the period from 2005 to 2014 on the mean value of openness to novelty, controlling for the log of real GDP p.c. in 2005. The coefficient of average openness to novelty is -1.67 and significant at the 1% level.

and the average annual growth rate of real GDP p.c. The underlying OLS regression controls for the level of GDP p.c. at the beginning of the period (2005) in order to account for convergence growth. The figure also shows two-letter country codes and indicates the income group a country (predominantly) belonged to according to the World Bank's classification for the corresponding years. Visual inspection suggests a negative correlation between openness to novelty and growth. This is borne out by the statistical analysis. The slope coefficient of the line of best fit is -1.67, with statistical significance at the 1% level. The relationship is also economically significant: a one-standard-deviation rise in openness to novelty is associated with a fall in the growth rate of 0.66 percentage points.

The negative relationship in Figure 1 is robust to a large number of modifications. In the rest of this section, we first discuss three obvious robustness checks. We then explore the relationship between openness to novelty and growth when low and lower-middle income countries are excluded from the sample. Finally, we add standard determinants of economic growth as additional controls to the regression equation.

As a first robustness check, we eliminate outliers (5% in each dimension). Doing so leaves the relationship intact. In fact, the slope coefficient turns even more negative (-2.19) and statistical significance rises to the 0.1% level. Second, when we rely on just one item to construct our openness measure, we find that the negative relationship is not exclusively driven by one of them: the pattern in Figure 1 holds in each case separately. Third, we ran a pooled regression that includes two observations per country (one for each wave).[5] Again, the findings confirm the baseline results. The coefficient is -1.79, significant at the 10% level.

In countries with lower incomes, economic growth (or the absence of it) is not primarily determined by the pace of innovation and adoption (but often by factors such as international aid, political instability, or internal conflict). When we exclude low income countries from the sample (7 observations), we find an economically stronger negative relationship, with a coefficient of -2.11, significant at the 0.1% level. When we additionally exclude lower-middle income countries (17 observations), the relationship becomes even stronger: the coefficient increases (in absolute terms) to -2.76, significant at the 1% level. Through the lens of our theoretical analysis, observing a stronger correlation when low-income countries, and then lower-middle-income countries, are excluded is exactly to be expected: as the practical relevance of the proposed openness-innovation nexus is limited to countries that are not too far away from the technology frontier, including countries with lower incomes in the sample is likely to weaken any correlation caused by the nexus among countries with higher incomes.

In the empirical growth literature, proxies for institutional quality and human capital are standard regressors. A priori, it is plausible to expect that our measure of general openness to novelty correlates with those regressors. For instance, one can imagine that eduction fosters curiosity and openness. For this reason, we ran another set of regressions that additionally and simultaneously include controls for institutional quality (mean of six governance indicators) and human capital (average years of schooling among adults aged 25+).[6] We find that our baseline results are robust to these modifications, too. When we use the full sample (countries with lower incomes included), the slope coefficient is still negative, but with a value of -1.07 it is smaller in absolute size and also less significant in statistical terms (10% level). However, when we exclude low income counties, or those with low and lower-middle incomes (see Figure 2), the value of the slope coefficient is again around -2 and statistical significance returns to the 5% level. So it does not appear that the baseline correlation between growth and openness to novelty is driven by the omission of other plausible growth determinants.

[5] Since only 44% of the 75 countries are part of both waves, the number of observations increases to just 108.

[6] The six governance indicators (voice and accountability, political stability and absence of violence/terrorism, government effectiveness, regulatory quality, rule of law, control of corruption) come from the Word Bank's World Governance Indicators database. Average years of schooling come from the Barro-Lee dataset (for seven countries, we had to turn to either the Wittgenstein Projection dataset or the UNESCO Institute for Statistics).

Figure 2: Openness to novelty and growth, countries with higher incomes, additional controls

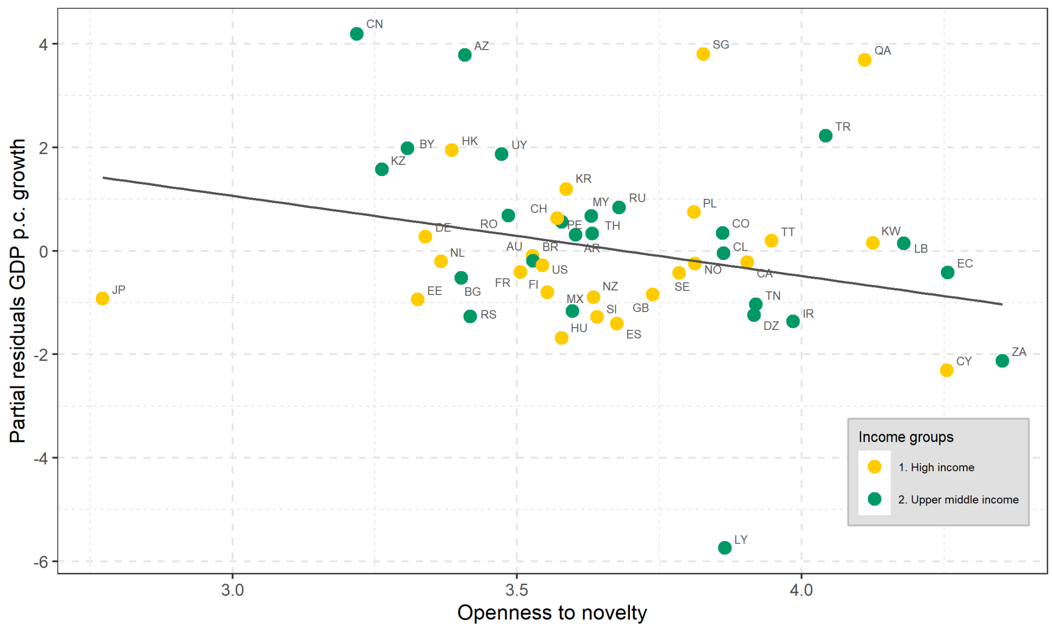

Notes. This is a partial residual plot for high and upper-middle income countries. The underlying OLS model regresses average real GDP p.c. growth over the period from 2005 to 2014 on the mean value of openness to novelty, controlling for the log of real GDP p.c. in 2005, institutional quality, and human capital. The coefficient of average openness to novelty is -2.14 and significant at the 5% level.

Overall, for the decade starting in 2005, our empirical analysis documents a robust negative correlation between economic growth and a new WVS-based measure of openness to novelty (controlling for initial GDP p.c. and other growth determinants). Obviously, the analysis does not establish causality. Still, the clear negative association may come as a surprise. The following section develops a theoretical model that allows us to explore the channels by which entrepreneurial openness to novelty may impact long-run economic growth. Simulations of the model clarify when we should expect the kind of negative relationship illustrated in Figure 1.

3 Model

3.1 Final Good Sector

The economy produces a single final good, which can either be consumed or converted into capital. The conversion is one-for-one. Capital depreciates at a rate of $\kappa \geq 0$. The final good is produced from capital and an intermediate good. We refer to the intermediate good as the "technology good". The rental rate of capital is denoted by r and the price of the technology good by m. All firms operating in the final good sector are price takers, both in the factor markets and in the output market. The final good is the numeraire.

We refer to Y_t as the aggregate supply of the final good in period t. The final good sector can be represented by a common CD aggregate production function

$$Y_t = F(K_t, X_t) = (K_t)^{\alpha}(X_t)^{1-\alpha}, \tag{1}$$

where $\alpha \in (0,1)$. K_t and X_t are the aggregate supplies of capital and the technology good, respectively. As will become clear, X_t can also be considered a measure of overall productivity.

The assumptions so far imply that the factor prices are given by

$$r_t = \alpha(K_t/X_t)^{\alpha-1} - \kappa \quad \text{and} \quad m_t = (1-\alpha)(K_t/X_t)^{\alpha}. \tag{2}$$

As the aggregate production function (1) has constant returns to scale, there are no profits in the final goods sector. The technology good is supplied by a continuum of mass one of independent entrepreneurs who also own and supply the capital stock.

3.2 Technology Good Sector

Two production methods. In each period t, the entrepreneurs that form the technology good sector have access to two different production methods, a tested (or "old") method and a new method. Each method produces an output that is proportional to the time spent on using that method. We use l_{it}^k for the time entrepreneur $i \in [0,1]$ in period t devotes to production method $k \in \{o, n\}$. As the total time endowment is normalized to 1, we have $l_{it}^o + l_{it}^n = 1$. The tested method produces an output of

$$x_{it}^o = A_t^o l_{it}^o, \tag{3}$$

where A_t^o measures the productivity of the tested method in period t. Similarly, the output of the new production method is given by

$$x_{it}^n = A_t^n l_{it}^n, \tag{4}$$

with the productivity variable depending on the "fundamental" of the method, denoted by F_t:

$$A_t^n = \begin{cases} \theta^H A_t^o & \text{if } F_t = H \\ \theta^L A_t^o & \text{if } F_t = L \end{cases}, \tag{5}$$

where $\theta^H > 1 > \theta^L$. For simplicity, we assume $\theta^L = 1/\theta^H$. So the new method is either more or less productive than the tested one, depending on the realization of the fundamental. Importantly, at the time entrepreneurs have to decide on the use of methods, they do not have any objective information on $\Pr[F_t = H]$. They are thus confronted with *fundamental uncertainty* (i.e. uncertainty in the sense of Knight 1921). We will specify below how entrepreneurs proceed in the face of such uncertainty.

If and only if a positive mass of entrepreneurs use the new method, F_t will be revealed by the end of period t. In case the new method is revealed to be a success ($F_t = H$), it will become the tested method in the next period. In case the new method is revealed to be a failure ($F_t = L$), today's tested method will be tomorrow's, too. The same holds if F_t is not revealed. In summary, we assume that

$$A_{t+1}^o = \begin{cases} A_t^n & \text{if } F_t \text{ revealed } \wedge \ F_t = H \\ A_t^o & \text{otherwise} \end{cases}. \tag{6}$$

But no matter which one of the cases in described in equation (6) prevails in period t, a "fresh" new method will arrive in period $t+1$.

Time allocation. In each period, entrepreneurs must decide on how to split their time endowment between the two methods. Following Binswanger et al. (2021), this happens in two steps. In a first step, entrepreneurs make a provisional choice. We use $\tilde{l}_{it}^k$ for the time entrepreneur $i \in [0,1]$ in period t provisionally allocates to production method $k \in \{o, n\}$. After some time, all entrepreneur who have allocated time to the new method receive an identical signal $S_t \in \{H, L\}$ about F_t. The signal's quality is given by $\sigma = \Pr[S_t = F_t] > 1/2$ for $F_t, S_t \in \{H, L\}$. A larger value of σ means a more precise signal. Although not explicitly modeled, we can imagine that the signal arrives at the end of a preparation stage, just before the start of the actual production stage. In a second step, considering the signal, entrepreneurs with $\tilde{l}_{\cdot t}^n > 0$ receive a chance to terminate the use of the new method and to reallocate the time originally devoted to the new method towards the tested one. The amount of time that can be reallocated is $\lambda \tilde{l}_{\cdot t}^n$, where $\lambda < 1$. The termination-induced loss of time may reflect that the time spent on preparation works is irrecoverable. In what follows, we make the following assumption:

$$\lambda > (1-\sigma)\theta^H + \sigma/\theta^H \tag{A1}$$

We are going to clarify the meaning of restriction (A1) in Section 4 below. Essentially, it permits the richest set of patterns when it comes to entrepreneurial time allocation.

For entrepreneurs who, in the first step, have decided not to allocate any time to the new production method, it follows that the provisional allocation and the final one are identical. For an entrepreneur i with $\tilde{l}_{it}^n > 0$, the final allocation is as follows:

$$(l_{it}^o, l_{it}^n) = \begin{cases} (\tilde{l}_{it}^o, \tilde{l}_{it}^n) & \text{if } i \text{ does not terminate use of new method} \\ (l_{it}^o + \lambda \tilde{l}_{it}^n, 0) & \text{otherwise} \end{cases}. \tag{7}$$

After the (potential) reallocation of time, no more decisions are to be taken and the technology good is being produced. Its aggregate supplies are given by

$$X_t = \int_0^1 (x_{it}^o + x_{it}^n)\, di = \int_0^1 (A_t^o l_{it}^o + A_t^n l_{it}^n)\, di. \tag{8}$$

As the number of entrepreneurs as well as their time endowments are kept constant, X_t is also a direct measure of the productivity of the technology good sector.

3.3 Technology Entrepreneurs

Uncertainty and beliefs. Owing to the lack of objective information on the probability that the new production method has an edge over the tested one, entrepreneurs adopt a subjective prior belief in this regard. We denote entrepreneur i's subjective prior about $\Pr[F_t = H]$ by p_{it}. Entrepreneurial priors are determined by two different factors:

$$p_{it} = p(\pi_t; \psi_i) = \pi_t^{(1/\psi_i)-1}. \tag{9}$$

The non-specific factor, $\pi_t \in (0,1)$, is called "belief anchor" and captures the general attitude towards the innovation. A high level means that the innovation comes with a considerable "hype". Binswanger and Oechslin (2021) offer a generic model in which π_t emerges endogenously as the outcome of a contest game between so-called "belief entrepreneurs"—two actors with (conflicting) interests in, respectively, favorable and unfavorable subjective priors. In Binswanger and Oechslin (2021), structural parameters as well as random forces—such as a

Figure 3: Distribution of subjective prior beliefs

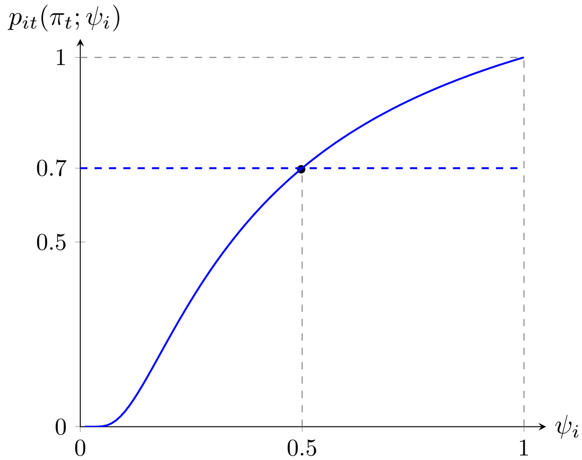

narrative suddenly "going viral" (e.g., Shiller 2017, 2019)—affect the equilibrium belief anchor.[7] Here we abstract from the structural part and model the belief anchor as influenced by random forces and recent economic events. The details are given below.

The second factor in equation (9) that affects p_{it} is an entrepreneur-specific characteristic, $\psi_i \in [1/2 - \delta, 1/2 + \delta]$, where $\delta \in (0, 1/2)$. Note that for $\psi_i = 1/2$, we obtain $p_{it} = \pi_t$; for all $\psi_i < 1/2$, $p_{it} < \pi_t$ and for all $\psi_i > 1/2$, $p_{it} > \pi_t$. As a result, entrepreneurs with $\psi_i > 1/2$ generally have more favorable beliefs towards innovative production methods than the median entrepreneur, while the reverse is true for entrepreneurs with $\psi_i < 1/2$. The ψ_is reflect individual personality traits. In particular, they capture *individual openness to novelty* (in terms of how businesses are run). We assume that individual openness follows a uniform distribution Ω on $[1/2 - \delta,\ 1/2 + \delta]$ with density $\omega(\psi)$. We thus deviate from the common prior assumption and assume that the entrepreneurs agree to disagree.

The parameter δ is our measure of *heterogeneity* in terms of openness among entrepreneurs. If $\delta = 0$, there is no heterogeneity at all: entrepreneurs' priors are identical to the belief anchor. For $\delta = 1/2$, there is maximum heterogeneity, with the resulting priors ranging from 0 to 1. Figure 3 illustrates the distribution of p_{it}, assuming maximum heterogeneity and $\pi_t = 0.7$. In what follows, we will rely on the notion of exuberance. An *exuberant prior belief* is defined as one being so favorable that the entrepreneur holding it continues with a new production method even in the event of a negative signal. Obviously, all else equal, greater openness to novelty implies a higher probability of holding an exuberant prior belief.

[7]That model's premise is that under fundamental uncertainty people are prone to heuristic belief formation based on analogies. This tendency is exploited by belief entrepreneurs who offer suitable analogies, wrapped in narratives, in order to steer subjective prior beliefs in a desired direction.

Having introduced openness to novelty, we can now complete the discussion of the belief anchor. The anchor consists of two components, the first exogenous and the second endogenous:

$$\pi_t = (\tilde{\pi}_t)^{1/\Psi} Q_t. \tag{10}$$

In equation (10), $\tilde{\pi}_t$ is an i.i.d. random variable with a uniform distribution on $[0, 1]$. The role of $\Psi \in (0, \infty)$ is to parameterize the level of *general openness to novelty*: the larger Ψ, the larger the mean of $(\tilde{\pi}_t)^{1/\Psi}$, which is given by $\Psi/(1 + \Psi) \in [0, 1)$.[8] While broadening the distribution of the ψ_is via δ will allows us to explore the impact of heterogeneity on growth, we can turn to Ψ to explore the effect of altering the level of general openness. Ψ can be called a "belief parameter" that may have a cultural component (e.g., Guiso et al. 2006), capturing a cultures' general attitude towards novelty. However, as noted by Mokyr (2017), not all people of the same culture share identical attitudes—hence the individual ψ_is.

The second component of the belief anchor, $Q_t \in \{0, 1\}$, allows for spillovers of disappointing past innovation experiences on current belief formation. In the fields of management and finance, a growing body of evidence suggests that the experience of negative shocks may alter decision makers' expectations and behavior (e.g., Malmendier and Tate 2011; Dittmar and Duchin 2016; Guiso et al. 2018). In particular, as Dittmar and Duchin (2016, p. 566) note, "experiencing troubles may alter risk preferences or expectations, and lead managers to implement more conservative policies", not least by making possible downsides more salient (e.g., Bordalo et al. 2012). Here, a potentially expectations-changing negative shock is understood as a situation in which, against the backdrop of exuberant entrepreneurial priors, a new production method is revealed to be a failure—with the result that high hopes are disappointed and the downside of working with a new production method gains in salience. The purpose of Q_t in belief anchor (10) is to capture the possible move towards temporarily "more conservative" entrepreneurial expectations in response to such a negative shock.[9]

For concreteness, assume the following. If in period $t - 1$ the new production method is revealed to be a failure, there is a chance that $Q_t = 0$ (as opposed to $Q_t = 1$). This chance is given by the share of entrepreneurs who prior to the disappointing revelation have held an exuberant prior. In the event of $Q_t = 0$, the period-t priors about the new production method are most conservative and, as a result, the process of technological improvements is completely paralyzed in that period; but the shift to maximally conservative priors is only temporary: in

[8] More formally, RV $(\tilde{\pi}_t)^{1/\Psi}$ has a first-order stochastic dominance over $\tilde{\pi}_t$ if $\Psi > 1$ (and vice versa if $\Psi < 1$).

[9] In practice, entrepreneurial expectations are not the only channel by which negative shocks may slow technological progress through innovation. If a shock triggers a banking crisis, firms' access to—in particular—external R&D finance worsens, with negative consequences for innovation activity (Hardy and Sever 2021).

the following period, Q returns to its positive value: $Q_{t+1} = 1$. Under paralysis, the downside of working with a new production method loses in terms of salience.

Consumption and saving. Entrepreneurs (or their dynasties) live forever. They have preferences over current consumption and future assets. Entrepreneur i's expected utility in t reads

$$U_{it} = \mathbb{E}_{p_{it}} \left\{ u(c_{it}, k_{it+1}) \right\} = \mathbb{E}_{p_{it}} \left\{ (c_{it})^{1-\beta} (k_{it+1})^{\beta} \right\}, \tag{11}$$

where $\mathbb{E}_{p_{it}} \{\cdot\}$ refers to expectations formed at the beginning of period t under the entrepreneur's subjective prior belief about $\Pr[F_t = H]$. Preferences are represented by a CD utility function $u(\cdot)$ and we use c_{it} and k_{it+1} to denote, respectively, entrepreneur i's consumption in period t and assets in period $t+1$. The parameter $\beta \in (0, 1)$ represents the weight with which future assets enter preferences. Given that entrepreneurs have incomplete knowledge about the economy, and thus standard intertemporal choice infeasible, including next-period assets is a tractable way of modeling the trade-off between present consumption and future opportunities (see Binswanger et al. 2022). Beyond that, future assets may enter preferences because of a "capitalist spirit" under which asset accumulation "is an end in itself" (e.g., Francis 2009, p. 396)—which is clearly plausible when entrepreneurs are concerned.[10]

In each period, entrepreneurs receive income from the sale of the intermediate good and from asset holdings. Entrepreneur i's flow budget constraint reads

$$k_{it+1} = (1 + r_t)k_{it} + m_t x_{it} - c_{it}, \tag{12}$$

where $x_{it} = x^o_{it} + x^n_{it}$ is total output across the two methods of production (equations 3 and 4, respectively). Below it will be helpful to work with the definition

$$z_{it} = (1 + r_t)k_{it} + m_t x_{it}. \tag{13}$$

The consumption-saving decision takes place towards the period end (see below for the exact timing of decisions). In particular, when deciding on c_{it}, entrepreneur i knows all the components of z_{it}. Given utility function $u(\cdot)$ and flow budget constraint (12), this means that the choice of consumption is a decision under certainty. As a result, standard calculations yield

$$c_{it} = (1 - \beta)z_{it} \quad \text{and} \quad k_{it+1} = \beta z_{it}. \tag{14}$$

[10] See Michaillat and Saez (2021) for additional justifications for incorporating wealth in the utility function.

Taking the expressions in (14) for current consumption and assets one period ahead into account, beginning-of-period expected utility (11) can be rewritten as

$$U_{it} = (1-\beta)^{1-\beta}\beta^{\beta}\ \mathbb{E}_{p_{it}}\left\{z_{it}\right\}. \tag{15}$$

It follows that when in period t entrepreneur i determines the allocation of time across the two production methods (old vs. new), the aim is to maximize $\mathbb{E}_{p_{it}}\{z_{it}\}$.[11]

3.4 Timing

Within each period t, there is a maximum of seven stages. The sequence of events is as follows:

1. From the previous period, all entrepreneurs $i \in [0,1]$ inherit their individual asset holdings, k_{it}, and the productivity level A^o_t (either A^o_{t-1} or A^n_{t-1}).
2. Nature draws the initially unobservable fundamental of the new production method, $F_t \in \{H, L\}$, as well as the belief anchor, $\pi_t \in [0,1]$.
3. Observing π_t, all entrepreneurs adopt a subjective prior belief about $\Pr[F_t = H]$, p_{it}, and then decide on the provisional time allocation, $(\tilde{l}^o_{it}, \tilde{l}^n_{it})$.

If non of the entrepreneurs provisionally allocates time to the new production method, stages 4 and 5 are skipped. Otherwise, the sequence continues with stage 4:

4. Nature draws the informative but noisy signal about the fundamental, $S_t \in \{H, L\}$.
5. Observing S_t, all entrepreneurs with $\tilde{l}^n_{it} > 0$ form their posterior belief, q_{it}, and then decide on the final time allocation, (l^o_{it}, l^n_{it}).
6. Production takes place, incomes are incurred (and the z_{it}s observed), and—provided that $l^n_{it} > 0$ for some entrepreneurs—F_t is inferred.
7. All entrepreneurs divide z_{it} between current consumption, c_{it}, and future assets, k_{it+1}.

As the above timeline makes clear, a period in the model starts with the arrival of an uncertain innovation and it ends with the (possible) resolution of uncertainty. Data from Goldfarb and Kirsch (2019), an extensive study that traces major technological innovations over the past 150 years, suggest that the median length of such uncertainty windows is 15 years. Dedehayir and Steinert (2016, p. 29), dealing with hype cycles surrounding innovations, argue that their

[11] The timing of decisions, together with the linearity of indirect utility in z_{it} (which is a property of homothetic preferences), implies that entrepreneurs behave in a risk-neutral way when choosing the time allocation.

length “may vary between two years and two decades, although so-called ‘normal technologies’ are anticipated to take five to eight years, [...].” Given this information, it seems reasonable to work with a period length of ten years when calibrating the model below.

4 Static Equilibrium

4.1 Allocation

Consumption and knowledge. To characterize the static equilibrium, we go backwards through the seven stages listed above. The decision problem at stage 7 is standard and already solved (equation 14). At stage 6, any entrepreneur i who has allocated time to the new production method can infer A_t^n from observing x_{it}^n (equation 4); information about A_t^n so derived instantaneously spreads and becomes common knowledge (F_t revealed).

Continuation/termination decision. Consider an entrepreneur i who at stage 3 has provisionally allocated a positive amount of time to the new production method ($\tilde{l}_{it}^n > 0$). Having observed Nature's draw of signal S_t at stage 4, the entrepreneur at stage 5 updates $p_{it} = \Pr_i[F_t = H]$, their subjective prior belief regarding the probability of $A_t^n > A_t^o$ (or $F_t = H$). Bayes' Rule implies that the posterior belief is given by

$$q_{it} = q(p_{it}, S_t; \sigma) = \begin{cases} \left[1 + \frac{1-p_{it}}{p_{it}} \frac{1-\sigma}{\sigma}\right]^{-1} & : \quad S_t = H \\ \left[1 + \frac{1-p_{it}}{p_{it}} \frac{\sigma}{1-\sigma}\right]^{-1} & : \quad S_t = L \end{cases}, \tag{16}$$

where the notation in equation (16) indicates that q_{it} is a function of the prior, the signal, and the signal's quality. Based on this posterior, entrepreneur i then decides whether to continue or to terminate the use of the new production method. Since the entrepreneur's objective is to maximize the expected quantity of output produced at stage 6 (see Proposition 1 below), the entrepreneur continues if the expected output following a continuation decision exceeds the (known) output following a termination decision. Using the definition of A_t^n in equation (5), and accounting for the fact that only a fraction λ of the time allocated to the new production method can be re-allocated to the old method, this is the case if

$$q_{it}\theta^H A_t^o \tilde{l}_{it}^n + (1 - q_{it})\theta^L A_t^o \tilde{l}_{it}^n \geq \lambda A_t^o \tilde{l}_{it}^n. \tag{17}$$

Figure 4: Posterior belief q_{it} as a function of p_{it} and S_t

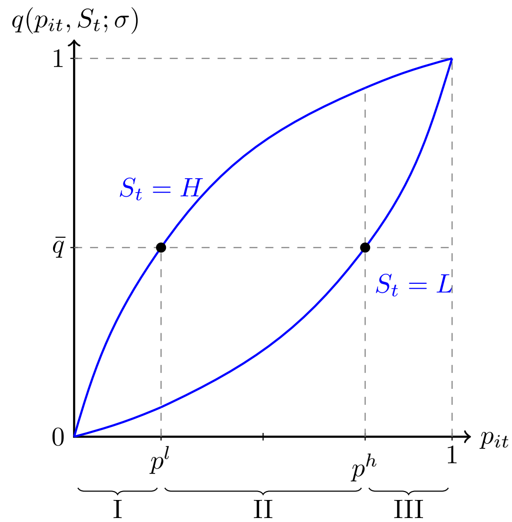

Rearranging terms in equation (17) gives the equivalent condition

$$q_{it} \geq \bar{q} \equiv \frac{\lambda - \theta^L}{\theta^H - \theta^L} \in (0, 1), \tag{18}$$

where $\theta^L < \lambda < \theta^H$ guarantees that $\bar{q}$ is strictly between zero and one. So, if q_{it} exceeds $\bar{q}$, the provisional time allocation is unaltered and entrepreneur i definitely employs the new method, while in the opposite event the entrepreneur turns away from it:

$$(l_{it}^o, l_{it}^n) = \begin{cases} (\tilde{l}_{it}^o, \tilde{l}_{it}^n) & \text{if} \quad q_{it} \geq \bar{q} \\ (l_{it}^o + \lambda \tilde{l}_{it}^n, 0) & \text{if} \quad q_{it} < \bar{q} \end{cases}. \tag{7'}$$

Figure 4 visualizes q_{it} and $\bar{q}$. No matter the realization of $S_t \in \{H, L\}$, q_{it} is a strictly increasing function of p_{it}, monotonically rising from zero to one; moreover, with the exception of the start and end points, $q(p_{it}, H; \sigma) > q(p_{it}, H; \sigma)$. A key implication of the figure is that the signal is not always decisive for the continuation/termination decision. In particular, if $p_{it} < p^l$, entrepreneur i terminates the use of the new method even in case of $S_t = H$; and if $p_{it} \geq p^h$, the entrepreneur continues even in case of $S_t = L$. Thus, in these two cases, the entrepreneur chooses to ignore the sole objective source of information. Only if $p_{it} \in \left[p^l, p^h\right)$, the signal determines the decision by the entrepreneur. The three different ranges are denoted by I, II, and III, respectively. For the two thresholds, we get

$$p^l = \frac{(1-\sigma)(\lambda - \theta^L)}{\lambda - \theta^L + \sigma(\theta^H + \theta^L - 2\lambda)}. \tag{19}$$

and

$$p^h = \frac{\sigma(\lambda - \theta^L)}{\theta^H - \lambda + \sigma(2\lambda - \theta^L - \theta^H)}. \tag{20}$$

Provisional time allocation. At stage 3, after having acquired subjective prior p_{it}, entrepreneur i must decide on the provisional division of the total time endowment across the two production methods. The entrepreneur does so with the aim of maximizing expected utility, U_{it}, which according to equation (15) is linear in $\mathbb{E}_{it}\{z_{it}\}$. In this regard, note the following:

PROPOSITION 1 *A provisional time allocation $(\tilde{l}^o_{it}, \tilde{l}^n_{it})$ that maximizes $\mathbb{E}_{p_{it}}[x_{it}(\tilde{l}^o_{it}, \tilde{l}^n_{it})]$, i.e., the quantity of output entrepreneur i expects to produce at stage 6, also maximizes entrepreneur i's expectation of end-of-period resources, $\mathbb{E}_{p_{it}}\{z_{it}\}$.*

Proof. See Appendix ■

So, when choosing the provisional time allocation, entrepreneur i can simply focus on the maximization of $\mathbb{E}_{p_{it}}[x_{it}(\tilde{l}^o_{it}, \tilde{l}^n_{it})]$. While any split of the total time endowment is possible, it is a priori clear that in order to maximize expected output the entrepreneur will either allocate the full endowment to the new method—i.e., $(\tilde{l}^o_{it}, \tilde{l}^n_{it}) = (0, 1)$—or to the old method—i.e., $(\tilde{l}^o_{it}, \tilde{l}^n_{it}) = (1, 0)$. This follows from the linearity of the two production functions in $l^{\cdot}_{it}$ (equations 3 and 4). To decide, it is thus sufficient that the entrepreneur answers a simple question: in case the entire time endowment were provisionally allocated to the new production method, would the output expected to be produced at stage 6 be greater or less than the output if the full endowment were allocated to the old method? Formally: $\mathbb{E}_{p_{it}}[x_{it}(0,1)] \gtrless \mathbb{E}_{p_{it}}[x_{it}(1,0)]$?

According to the analysis above, three cases must be distinguished. First, if p_{it} is in range I, entrepreneur i anticipates the unconditional termination of the new method at stage 5. Thus,

$$\mathbb{E}_{p_{it \in \mathrm{I}}}[x_{it}(0,1)] = \lambda A^o_t. \tag{21}$$

Equation (21) reflects that after the termination of the new method in stage 5 only a fraction λ of the total time endowment is still available for reallocation. Second, if p_{it} is in range II, entrepreneur i anticipates the termination (continuation) of the new method if $S_t = H$

Figure 5: Categorization of beliefs

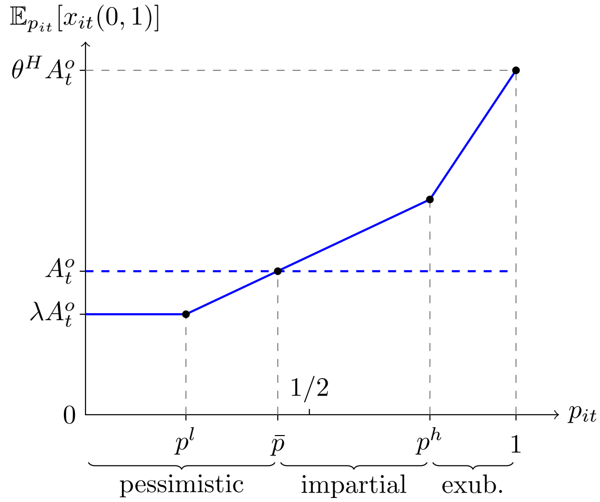

$(S_t = L)$. The analog to equation (21) is

$$\mathbb{E}_{p_{it\in\mathrm{II}}}[x_{it}(0,1)] = \left\{p_{it}\sigma\theta^H + (1-p_{it})(1-\sigma)\theta^L + [p_{it}(1-\sigma) + (1-p_{it})\sigma]\,\lambda\right\} A_t^o. \tag{22}$$

The first summand on the right-hand side of equation (22) captures the possibility $S_t = F_t = H$, which entrepreneur i thinks to occur with probability $p_{it}\sigma$; the second summand captures the possibility $S_t = H \neq F_t = L$, which entrepreneur i thinks to occur with probability $(1-p_{it})(1-\sigma)$; the expression in square brackets represents the (subjective) probability of $S_t = L$. Finally, if p_{it} is in range III, entrepreneur i anticipates the unconditional continuation of the new method in stage 5. As a result,

$$\mathbb{E}_{p_{it\in\mathrm{III}}}[x_{it}(0,1)] = \left[p_{it}\theta^H + (1-p_{it})\theta^L\right] A_t^o. \tag{23}$$

Figure 5 illustrates $\mathbb{E}_{p_{it}}[x_{it}(0,1)]$ as a function of p_{it}. As can be seen, the function is monotonically increasing and piecewise linear (with kinks at p^l and p^h).

Now consider the output produced at stage 6 if the full time endowment were allocated to the old production method. In that case, the output would simply be A_t^o, a level that is also shown in Figure 5. The properties of $\mathbb{E}_{p_{it}}[x_{it}(0,1)]$ as a function of p_{it} guarantee that $\mathbb{E}_{p_{it}}[x_{it}(0,1)]$ crosses the A_t^o-threshold exactly once from below. We refer to the level of p_{it} that equates $\mathbb{E}_{p_{it}}[x_{it}(0,1)]$ and A_t^o as $\bar{p}$. It is implicitly defined by

$$\mathbb{E}_{\bar{p}}[x_{it}(0,1)] = A_t^o \tag{24}$$

and must fall in the range $\left(p^l, 1\right)$, as is shown in Figure 5. The questions raised above can now be answered with the help of Figure 5: allocating the time endowment to the new production method is expected to lead to a larger stage-6 output if and only if $p_{it} \geq \bar{p}$. As a consequence,

$$(\tilde{l}_{it}^o, \tilde{l}_{it}^n) = \begin{cases} (1,0) & \text{if} \quad p_{it} < \bar{p} \\ (0,1) & \text{if} \quad p_{it} \geq \bar{p} \end{cases}, \tag{25}$$

where $\bar{p}$ can be calculated as

$$\bar{p} = \frac{1 - \theta^L + \sigma(\theta^L - \lambda)}{\lambda - \theta^L + \sigma(\theta^H + \theta^L - 2\lambda)}. \tag{26}$$

In Figure 5, we have $\bar{p} < 1/2 < p^h$, a constellation that is guaranteed by constraint (A1). Entrepreneurs with a prior belief $p_{it} < \bar{p}$ allocate their full time endowment to the old production method. Entrepreneurs with $p_{it} \in \left[\bar{p}, p^h\right)$ provisionally allocate their full time endowment to the new method and then let their continuation/termination decision be determined by the signal. Entrepreneurs with $p^h \leq p_{it}$ allocate their full time endowment to the new method and stick to that decision even if $S_t = L$. Following Binswanger et al. (2021), we refer to priors that fall into the first and the second category as *pessimistic* and *impartial*, respectively. As already specified in Subsection 3.3, prior beliefs that fall into the third category are referred to as exuberant. To summarize:

PROPOSITION 2 *Suppose assumption (A1) holds. Then, in stage 3, an entrepreneur i with*

- $p_{it} < \bar{p} < p^h$ *(pessimistic prior) allocates the full time endowment to the old production method (and has no opportunity to reconsider this decision);*
- $\bar{p} \leq p_{it} < p^h$ *(impartial prior) provisionally allocates the full time endowment to the new production method and in stage 5 terminates its use if and only if* $S_t = L$*;*
- $\bar{p} < p^h \leq p_{it}$ *(exuberant prior) allocates the full time endowment to the new production method and in stage 5 sticks to this decision even if* $S_t = L$*.*

Proof. In the text above. ■

As a benchmark, it might be interesting to consider how an entrepreneur committed to the principle of indifference (see, e.g., Gilboa 2009, pp. 17-19) would act. Here, the principle of indifference entails that, a priori, $\Pr[F_t = H]$ is treated as a realization of a random variable with a uniform distribution on $[0, 1]$. So the entrepreneur would not view any particular

Figure 6: Distribution of ψ and beliefs

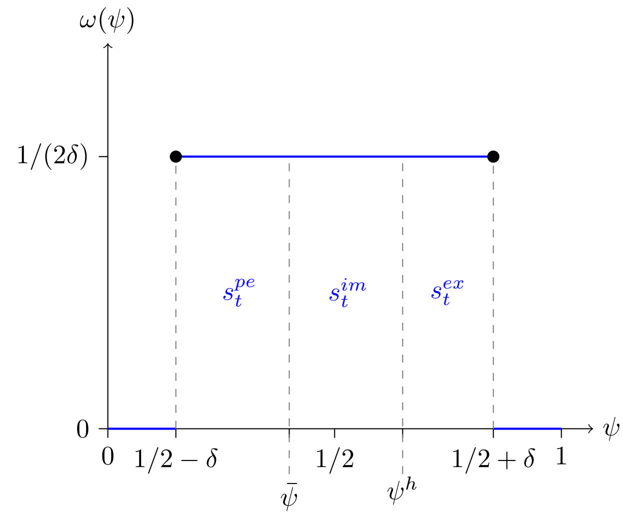

$\Pr[F_t = H]$ to be likelier than any other. One can show that such an "indifferent" entrepreneur would act in the same way as an entrepreneur with an impartial prior.[12]

4.2 Aggregation

Non-predetermined variables. The structure of the model is such that in general individual variables aggregate easily. Only the aggregation of the individual supplies of the technology good, specified in equation (8), requires attention. As entrepreneurs are heterogeneous in terms of openness to novelty, their behavior differs when it comes to the use of the two production methods. Yet as there are only three different types of behavior in this regards (reflecting pessimistic, impartial, and exuberant priors), aggregation is still straightforward.

As described above, it is $\bar{p}$ that separates pessimistic from impartial priors, while p^h marks the line between impartial and exuberant priors. We denote by s_t^{pe}, s_t^{im}, and s_t^{ex} the share of entrepreneurs who in period t hold pessimistic, impartial, and exuberant priors, respectively. Given belief anchor π_t, p_{it} is a strictly monotonous function of ψ_i (equation 9). We can thus find thresholds $\bar{\psi}$ and ψ^h such that all entrepreneurs with $\psi_i \in \left[0, \bar{\psi}\right)$ hold a pessimistic prior, all entrepreneurs with $\psi_i \in \left[\bar{\psi}, \psi^h\right)$ hold an impartial prior, and all entrepreneurs with $\psi_i \in \left[\psi^h, 1\right]$ hold an exuberant prior. Using equation (9), we obtain the following expressions:

$$\bar{\psi}(\pi_t) = [1 + \ln(\bar{p})/\ln(\pi_t)]^{-1} \quad \text{and} \quad \psi^h(\pi_t) = \left[1 + \ln(p^h)/\ln(\pi_t)\right]^{-1}. \tag{27}$$

[12] With the binary random variable F_t, an actor with a uniform prior acts in the same way as an actor with a prior of $1/2$. Since $\bar{p} < 1/2 < p^h$ (Figure 5), the indifferent entrepreneur thus behaves as the impartial.

The shares are now given as follows:

$$s_t^{pe} = \int_0^{\bar{\psi}(\pi_t)} \omega(\psi)\, d\psi \quad \text{and} \quad s_t^{im} = \int_{\bar{\psi}(\pi_t)}^{\psi^h(\pi_t)} \omega(\psi)\, d\psi \quad \text{and} \quad s_t^{ex} = \int_{\psi^h(\pi_t)}^{1} \omega(\psi)\, d\psi. \tag{28}$$

Figure 6 shows a possible equilibrium outcome in which all three shares are strictly positive. Such an outcome is likely if, as assumed in the figure, entrepreneurs are relatively heterogeneous in terms of openness to novelty (large δ). If δ were close to zero, $\omega(\psi)$ would be strictly positive only in a neighborhood around 1/2. With the values of $\bar{\psi}(\pi_t)$ and $\psi^h(\pi_t)$ shown in the figure, this would imply that the share of entrepreneurs with an impartial prior is 1.

Using the three shares, the aggregate supply of the technology good can now be written as

$$X_t = s_t^{pe} A_t^o + s_t^{im} \left[\mathbf{1}_{S_t=H} \cdot A_t^n + \mathbf{1}_{S_t=L} \cdot A_t^o \lambda \right] + s_t^{ex} A_t^n, \tag{29}$$

where $\mathbf{1}_{S_t=H}$ is an indicator variable that takes on the value 1 if and only if $S_t = H$ (an analogous definition applies to $\mathbf{1}_{S_t=L}$). Given X_t (and the predetermined aggregate capital stock), the aggregate output, Y_t, and the factor prices, r_t and m_t, follow immediately from equations (1) and (2), respectively. Moreover, from equations (13) and (14):

$$C_t = (1-\beta)[(1+r_t)K_t + m_t X_t]. \tag{30}$$

Predetermined variables. From equations (13) and (14), again, we get

$$K_{t+1} = \beta[(1+r_t)K_t + m_t X_t]. \tag{31}$$

Employing the notation introduced above, we can rewrite the law of motion for A_t^o (equation 6):

$$A_{t+1}^o = \begin{cases} A_t^n & \text{if } F_t = H \ \wedge \ (s_t^{im} \cdot \mathbf{1}_{S_t=H} > 0 \ \vee \ s_t^{ex} > 0) \\ A_t^o & \text{otherwise} \end{cases}. \tag{6'}$$

In equation (6'), the expression in parentheses replaces the term "F_t revealed" in equation (6). Finally, when it comes to the paralysis indicator, Q_{t+1}, we have to distinguish between two different cases. First, if $Q_t = 0$, it invariably follows that $Q_{t+1} = 1$. Second, if $Q_t = 1$, the paralysis indicator is a Bernoulli random variable whose success probability depends on the realization of the random variables F_t and s_t^{ex}. In particular,

$$\Pr[Q_{t+1} = 1 | Q_t = 1; F_t, s_t^{ex}] = \begin{cases} 1 & \text{if } F_t = H \\ 1 - s_t^{ex} & \text{otherwise} \end{cases}. \tag{32}$$

5 General Openness to Novelty and Growth

5.1 Analysis

From the perspective of progress through innovation, general openness to novelty is a double-edged sword. On the one hand, openness to novelty is a requirement for experimentation and learning. On the other hand, openness can lead to widespread exuberance—and then temporarily paralyze the process of technological improvements. In this section, we explore the relationship between the level of general openness, the pace of innovation, and economic growth. The study of heterogeneity in openness is left to the next section. We first provide analytical results relating openness and the pace of innovation and then turn to simulations to describe the behavior of the economy from a quantitative perspective. For both parts, we assume that the (unknown) probability of a new production method being a success is constant across time:

$$\Pr[F_t = H] = f > 0 \tag{A2}$$

In any given period, a necessary condition for experimentation with a new production method is that the previous period has not caused temporary paralysis through a combination of exuberance and disappointment. In particular, a positive chance of discovering a successful innovation in period t requires $Q_t = 1$. We now establish the probability of $Q_t = 1$, assuming that the entrepreneurs behave as described in Proposition 2.

PROPOSITION 3 *Suppose assumptions (A1) and (A2) holds and $\psi_i = 1/2$ for all i (no heterogeneity in openness to novelty). Then, for any arbitrary period t,*

$$\Pr[Q_t = 1] = \left\{1 + (1 - f)\left[1 - (p^h)^{\Psi}\right]\right\}^{-1}. \tag{33}$$

Proof. See Appendix ■

It follows from equation (33) that $\Pr[Q_t = 1]$ is a monotonically decreasing function of Ψ, reflecting that a higher level of general openness goes hand in hand with a generally higher chance of exuberant beliefs. Clearly, $\Pr[Q_t = 1]$ must influence the probability of discovering a successful innovation in period t. This probability is given as follows:

PROPOSITION 4 *Suppose assumptions (A1) and (A2) hold and $\psi_i = 1/2$ for all i (no heterogeneity in openness to novelty). Then, for any arbitrary period t,*

$$\Pr[\textit{new method in } t \textit{ revealed to be a success}] = f\frac{1-(1-\sigma)(p^h)^{\Psi}-\sigma(\bar{p})^{\Psi}}{1+(1-f)\left[1-(p^h)^{\Psi}\right]}. \tag{34}$$

Proof. See Appendix ■

In equation (34), the ambiguous role of the level of general openness to novelty, Ψ, is immediately apparent. An increase in Ψ raises the numerator, reflecting that greater openness lifts the probability that the belief anchor falls into a range that sparks the use of the new method (with or without listening to the signal). However, an increase in Ψ also raises the denominator: greater openness increases the chance of exuberance—which in period $t+1$ is followed by paralysis if $F_t = L$. The second effect strengthens as Ψ rises:

PROPOSITION 5 *Suppose that the signal's quality, $\sigma \in (1/2, 1]$, is sufficiently large such that*

$$f < \frac{2\sigma - 1}{\sigma}. \tag{35}$$

Then, $\Pr[\textit{new method in } t \textit{ revealed to be a success}]$ is a quasi-concave function of $\Psi \in [0, \infty)$. As Ψ rises from zero towards infinity, it monotonically increases from zero to some maximum level that is strictly greater than $f/(2-f)$ and then monotonically decreases towards $f/(2-f)$.

Proof. See Appendix ■

Figure 7 illustrates the success probability as a function of the level of general openness, assuming that condition (35) is satisfied. As Ψ rises from a low level, the probability of impartial beliefs increases at the expense of the probability of pessimistic beliefs. This must lift the success rate. However, at higher levels of Ψ, a further rise mainly increases the probability of exuberant beliefs at the expense of impartial beliefs. This reduces the success rate, provided that the signal (which matters under impartial beliefs only) is sufficiently informative and/or a low chance of $F_t = H$ makes temporary paralysis through a combination of exuberance and disappointment sufficiently likely. Viewed through the lens of the model, the negative empirical relationship between economic growth and general openness to novelty in Figure 1 suggests that in our sample most countries are beyond $\tilde{\Psi}$ in terms of openness.

As a final step before the turning to the simulations, we link the probability with which successful innovations are discovered to long-run productivity growth:

Figure 7: Openness to novelty and the pace of innovation—theory

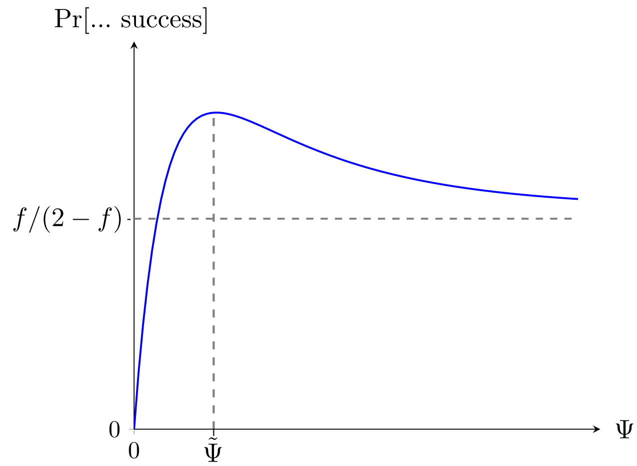

PROPOSITION 6 *Suppose assumptions (A1) and (A2) hold and $\psi_i = 1/2$ for all i (no heterogeneity in openness to novelty). Then, for any two periods t and $t+1$, the expected growth rate of the "technology frontier" is given by*

$$\mathbb{E}_f \left\{ \frac{A^o_{t+1} - A^o_t}{A^o_t} \right\} = (\theta^H - 1) f \frac{1 - (1-\sigma)(p^h)^\Psi - \sigma(\bar{p})^\Psi}{1 + (1-f)\left[1 - (p^h)^\Psi\right]}, \tag{36}$$

where the notation $\mathbb{E}_f\{\cdot\}$ indicates that the expectation is based on the true chance of $F_t = H$.

Proof. Follows from equations (5) and (6') and Proposition 4. ■

5.2 Simulation

Parametrization. While so far the analysis has provided a qualitative characterization of the relationship between general openness to novelty and the pace of innovation, it has not offered any quantitative insights. We now turn to a quantitative perspective, making use of the full model. In particular, we provide numbers on the effect of openness to novelty on the long-run growth rate of GDP p.c. in order to gauge the quantitative importance of two channels—experimentation-learning vs. exuberance-paralysis—identified above.

Moving to a quantitative perspective requires us to find a plausible range for general openness to novelty, Ψ. To this end, we turn to the motivating evidence in Section 2. Our empirical proxy for Ψ ranges from 1 to 6. The smallest and largest levels of openness to novelty in Figure 1 are approximately given by 2.5 and 5, respectively. So the smallest (largest) level exhausts 30% (80%) of the range. As $\Psi \in [0, \infty)$ is unbounded, this information does not immediately translate into a range for the model parameter. Yet the purpose of Ψ is to govern $\Psi/(1+\Psi) \in [0, 1)$, the mean of belief anchor π_t (conditional on $Q_t = 1$). Therefore, plausible theoretical counterparts for the two empirical values are implicitly defined by $\Psi/(1+\Psi) = 0.3$ (lower bound) and $\Psi/(1+\Psi) = 0.8$ (upper bound). From this, we obtain a Ψ-range of $[0.4, 4]$. As the empirical proxy stems from a single decade, we will work with a somewhat broader range: we set the lower and upper bound of Ψ such that $\Psi/(1+\Psi)$ lies in the range from 0.25 to 0.9. As a result, in the simulations below, Ψ runs from 0.33 to 9.

Abstracting from heterogeneity in openness to novelty, our model has eight structural parameters, all listed in Table 1, Panel A. We determine the parameters in three different ways. First, we assign standard values to α, κ, θ^H, and λ with the help of the existing macro and R&D literatures. Second, we turn to the motivating evidence presented in Section 2 to find a baseline value for Ψ. Third, we calibrate the parameters β, σ, and f by targeting growth and interest rate moments in a sample that spans about 200 years. Steps two and three are based on the Western Europe region, according to the classification by the Maddison Project (Bolt and Van Zanden 2020).[13] Western Europe is considered the cradle of the industrial revolution and has experienced innovation-driven growth in GDP per capita (p.c.) since more than 200 years (e.g., Mokyr 2017). The three ways are now explained in turn.

To assign values to α and κ, we consult Caselli and Feyrer (2007). Given production function (1), α corresponds to the reproducible capital's share of total output. The average among Western European countries is 20%. Regarding depreciation, Caselli and Feyrer (2007) work with a uniform rate of 6% at annual frequency. Hence we set κ to 79%, the corresponding rate at ten-year frequency. Akcigit and Kerr (2018) provide an indication for θ^H. From their study, one can infer that the innovation size of a major innovation, including all following-up improvements, approximately corresponds to a three-quarter leap in terms of productivity.[14] Accordingly, we choose 1.75 for θ^H. Finally, we work with $\lambda = 0.95$. So any termination-induced loss of time is just 5%, reflecting that the signal comes early.

The motivating evidence in Section 2 can not only be used to find an empirically plausible range for Ψ, but also to determine a baseline value that—jointly with the the values of α, κ, θ^H, and λ—can be used in the calibration below. The Western European average of our empirical measure for openness to novelty is 3.66, a value that exhausts 53% of the range. Following the logic used to find a plausible range for Ψ, 53% translate into $\Psi = 1.13$.

[13] We use the "region1950"-classification of the project's 2018 release. It includes 20 countries: AUT, BEL, CHE, CYP, DEU, DNK, ESP, FIN, FRA, GBR, GRC, IRL, ISL, ITA, LUX, MLT, NLD, NOR, PRT, SWE.

[14] We refer to the value of $\bar{s}/\eta$, where $\bar{s}$ is specified in equation (20) of Akcigit and Kerr (2018). The values of the parameters that enter equation (20) can be found in Tables 4 and 6 of that paper.

Table 1: Parameters and moments

Panel A: Parameters

Parameter	Description	Determination	Value
α	Repr capital's share	Literature	0.20
κ	Depreciation rate	Literature	0.79
θ^H	Innovation size	Literature	1.75
λ	Reallocation share	Literature	0.95
Ψ	Openness	Section 2	1.13
β	Weight of fut assets	Calibration	0.217
σ	Signal quality	Calibration	0.896
f	Obj $\Pr[F_t = H]$	Calibration	0.450

Panel B: Moments

	Av GDP p.c. growth	Sd GDP p.c. growth	Av Interest rate
Data	0.019	0.008	0.029
Model	0.019	0.007	0.029

Notes. This section abstracts from heterogeneity in individual openness to novelty ($\delta = 0$). The triple $(\theta^H, \lambda, \sigma) = (1.75, 0.95, 0.893)$ satisfies constraint (A1) that underlies Propositions 2 to 6.

The data for the empirical moments targeted in the calibration of the three remaining parameters come from two different historical datasets. The first dataset is the Maddison Project Database (Bolt and Van Zanden 2020), which offers consistent GDP p.c. figures (in 2011$) far back. From this data, we calculate for each Western European country the average annual GDP p.c. growth rate over the period from 1820 to 2018.[15] This set of 20 long-run growth rates forms the basis for two empirical moments, the average long-run growth rate of GDP p.c. among Western European countries (0.019) and the standard deviation of those long-run growth rates (0.008). The second historical dataset contains the underlying data to Schmelzing (2020), a paper that traces real interest rates over eight centuries. The dataset provides information for five Western European countries (GBR, NLD, DEU, FRA, ESP). For each of them, we calculate the average annual real interest rate over the period from 1820 to 2018. The average of the five rates (0.029) is the third empirical moment targeted in the calibration. The three moments are also listed in Table 1, Panel B.

Given the five parameters already fixed, the remaining parameters β, σ, and f are chosen so that the three empirical moments are matched by the corresponding model-generated moments. The latter are computed from 100,000 simulations of a 20-period (200 years) economy, after

[15] The start of the Industrial Revolution falls into the first two decades of the 19th century. If we go farther back than 1820, the data become increasingly scarce. For six countries, the period covered is less than 198 years.

Figure 8: Openness to novelty and GDP p.c. growth—simulation

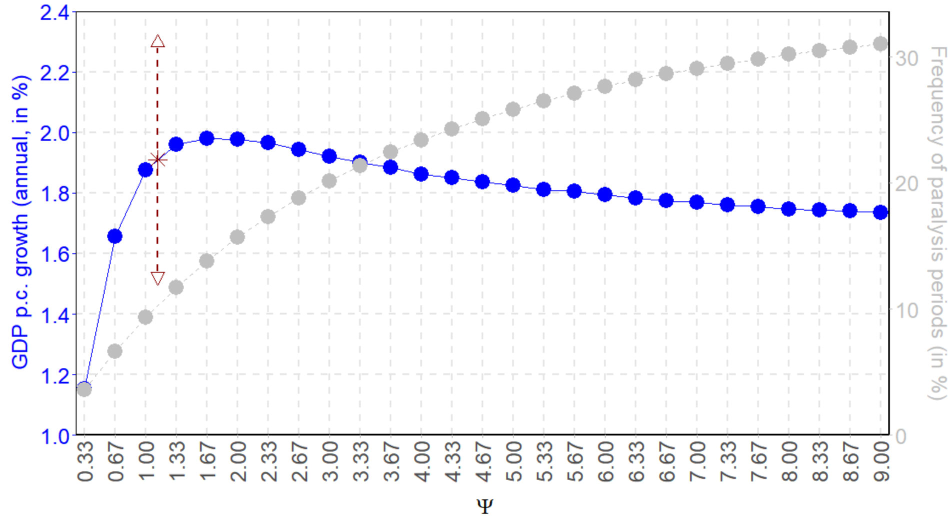

Notes. Each dot results from $100,000$ simulations of a 20-period (200 years) economy. See Table 1, Panel A, for the parameter values. The vertical red structure indicates the empirical value for average Western European GDP p.c. growth (star), as well as the standard deviation in that sample (arrow).

having transformed the simulated growth and interest data from ten-year frequency to annual frequency. Table 1, Panel B, shows that the combination of $\beta = 0.217$, $\sigma = 0.896$, and $f = 0.450$ leads to an almost perfect match between the empirical and the model-generated moments. It is easy to check that the triple $(\theta^H, \lambda, \sigma) = (1.75, 0.95, 0.896)$ satisfies constraint (A1) that underlines Propositions 2 to 6. Moreover, the pair $(\sigma, f) = (0.896, 0.450)$ satisfies condition (35). As a result, we know from Propositions 5 and 6 that there is a hump-shaped relationship between Ψ and the expected productivity growth rate. The simulation will show whether the peak in terms of GDP p.c. growth lies within the empirically plausible range of Ψ—and, if so, whether it is closer to the lower end or the upper end of the range.

Varying Ψ. Figure 8 shows average annual GDP p.c. growth and the frequency of paralysis periods as a function of openness to novelty, Ψ, where Ψ rises in increments of $1/3$ form the lower end of the plausible range $(1/3)$ to the upper end (9). Each dot results from $100,000$ simulations of a 20-period (200 years) economy that, with the exception of Ψ, is parameterized as shown in Table 1, Panel A. For comparison, the vertical red structure indicates the empirical value for average Western European GDP p.c. growth (cross), as well as the standard deviation in that sample (arrow). The horizontal position of the structure reflects that the Western European average of the empirical measure of openness to novelty translates into a Ψ of 1.13.

It turns out that the maximizer of GDP p.c. growth is close to the lower end of the considered Ψ-range. GDP p.c. growth first increases quickly in Ψ, reaches a peak, and then decreases monotonically and slowly as Ψ approaches the upper end of the range. This right-skewed shape is consistent with the empirical pattern in Figure 1 in the sense that over a large part of the range the relationship between Ψ and the growth rate of GDP p.c. is negative. So the simulation lends plausibility to the idea that, in parts, the negative correlation in Figure 1 actually picks up an overall adverse causal effect of openness to novelty on growth, i.e., that over a broad range the positive experimentation-learning channel is dominated by the negative exuberance-paralysis channel. Moving back in time, we can also link the asymmetric hump in Figure 8 to an important debate in economic history. In the 18th century, although similarly advanced in terms of property rights protection, economic liberties, and market integration, Europe started to significantly outperform China in terms of economic growth. Mokyr (1990, 2017) attributes this divergence, at least in parts, to an emerging difference in openness to new ideas. At the time, Europe became increasingly more open, while China remained conservative. In an era where useful knowledge became more plentiful, sciences were advancing, and civilizations established closer contact,[16] the emerging transformation in Europe proved highly advantageous in economic terms. The pattern in Figure 8 is consistent with this view in the sense that, when the starting point is low, increases in openness are predicted to cause steep gains in economic growth that allow the increasingly open entity to pull ahead.

Looking more closely at the simulated figures, we note that the effect of Ψ is sizable. As Ψ increases from its lower bound to the maximizer (1.67), annual GDP p.c. growth rises by 0.84 percentage points from 1.15% to 1.99%—only to fall overall by about a quarter of a percentage point to 1.73% (about one third of a standard deviation) as Ψ approaches its upper bound. At first sight, an annual difference of 0.26 percentage points may seem small. However, over time, such a difference carries weight. To see this, consider two economies, one with $\Psi = 9$ and the other one with $\Psi = 1.67$. Assume that initially the two economies have exactly the same level of GDP p.c. Then, after 100 years, the level of GDP p.c. in the former economy is only about 77% of the level in the latter. After 200 years, the ratio will stand at just 60%.

6 Heterogeneity in Openness to Novelty and Growth

We now change our focus, away from general openness to novelty towards heterogeneity in openness. In particular, this section sheds light on how economies that are identical in all

[16] When previously isolated civilizations connect, they may exchange, and put to work, useful practical and scientific knowledge. Mokyr (1990) calls the (possibly) resulting economic boom "exposure effect".

Figure 9: Heterogeneity in openness to novelty and GDP p.c. growth—simulation

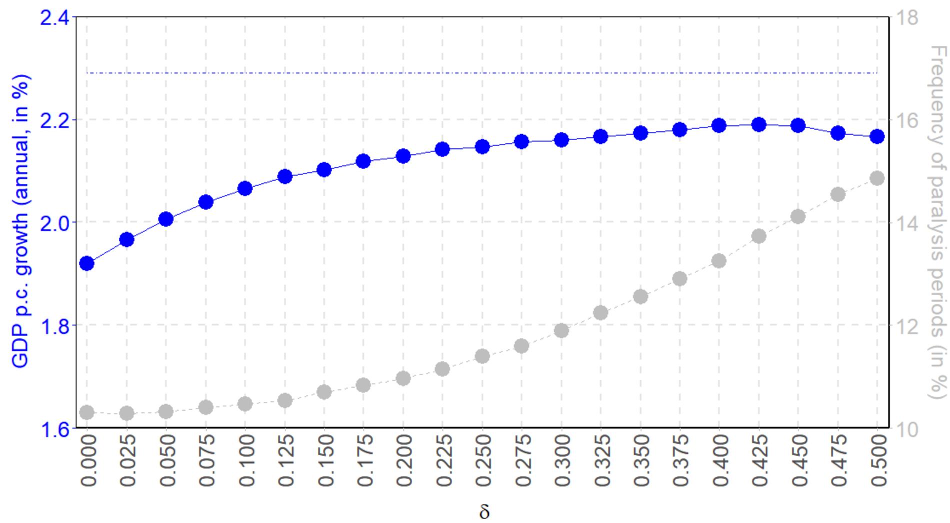

Notes. Each blue dot results from 100, 000 simulations of a 20-period (200 years) economy. See Table 1, Panel A, for the parameter values. The horizontal dashed line indicates GDP p.c. growth in a simulated economy that is populated by entrepreneurs who stick to the principle of indifference.

respects except for heterogeneity in openness perform in comparison. A priori, when it comes to heterogeneity in openness to novelty and GDP p.c. growth, there are stronger reasons to expect a clear positive relationship: a mean-preserving increase in heterogeneity makes some entrepreneurs more open for experimentation and learning—without giving a boost to the risk that a large number of entrepreneurs adopt an exuberant prior at the same time.

The simulation below again relies on the parameter values shown in Table 1, Panel A. However, deviating from the approach taken in Section 5, we no longer peg the level of heterogeneity in openness to novelty, δ, to a value of zero. Instead, we vary δ over the permissible range, while keeping the level of general openness to novelty, Ψ, fixed. Figure 9 shows the result: average annual GDP p.c. growth as a function δ, where δ rises in increments of 0.025 from zero to 1/2. Again, each dot results from 100, 000 simulations of a 20-period (200 years) economy that, with the exception of δ, is identically parameterized. For the most part of the range, the growth rate monotonically increases in δ, starting at a level of 1.91% p.a. and reaching a maximum value of 2.19% p.a. This is a difference of close to three tenth of a percentage points in terms of annual growth—a sizable effect. Only the final two increments, which also bring relatively large increases in the frequency of paralysis periods, cause a rather small fall in the growth rate. Thus, essentially, heterogeneity in openness to novelty helps growth.

It is worthwhile to discuss this finding in two different contexts. The first context is heterogeneity in prior beliefs. Although an obvious aspect of reality in many situations, this phenomenon used to attract little attention from economists.[17] But lately this has changed. Recent work focuses on, e.g., welfare analysis when agents have heterogeneous priors or on whether market selection would eliminate agents with distorted beliefs. In our setting, which emphasizes progress by means of uncertain innovations, belief heterogeneity is actually welcome in terms of long-run growth; from the perspective of the economy as a whole, it is a feature, not a bug. In fact, if belief heterogeneity is sufficiently pronounced, an economy populated by entrepreneurs that are subject to bouts of exuberance or pessimism behaves almost as if it were populated by "sober" entrepreneurs who stick to the principle of indifference (see Subsection 4.1). Specifically, simulations show that an economy populated by indifferent entrepreneurs (but otherwise equally parametrized) would grow at an annual rate of 2.29%, a level that is just marginally higher than the level achieved under pronounced heterogeneity.

The second context is diversity and long-run growth. The existing literature considers diversity along various dimensions. In Ashraf and Galor (2013), it is genetic diversity that impacts economic output. In their basic model, one effect of genetic diversity is to expand the production possibility frontier.[18] Similarly, Alesina et al. (2016) identify birthplace diversity as an important driver of an economy's production possibilities. To the extent that genetic/birthplace diversity and belief heterogeneity are positively correlated (which seems particularly plausible in the case of birthplace diversity), the current paper offers a mechanism by which diversity/heterogeneity lifts the pace by which the production possibility frontier expands. The economic logic is as follows. In the presence of uncertain innovations, heterogeneity in prior beliefs guarantees that some entrepreneurs are sufficiently open to experimentation—which is the only way the economy can learn which innovations "work" and which ones don't. At the same time, belief heterogeneity guarantees that some entrepreneurs are sufficiently skeptical to (temporarily) carry on with the status quo. In turn, this lowers the risk of resource misallocation at a massive scale—and thus that of a crisis that retards economic growth by depleting the capital stock and paralyzing the process of innovation.

[17] In economics, it is often assumed that individuals share a common prior. See, e.g., Morris (1995) for the main arguments in favor of the common prior assumption—and a critique of those arguments.

[18] Ashraf and Galor (2013) also point to negative effects of diversity like mistrust (also see Arbatlı et al. 2020). We could introduce a negative channel along those lines by assuming that skepticism slows down learning from more open entrepreneurs. Yet doing so would create a negative level effect rather than a negative growth effect.

7 Conclusion

This paper puts the spotlight on two staunch companions of innovation, uncertainty and subjective prior beliefs—companions that make what we call entrepreneurial "openness to novelty" an important determinant of the pace of innovation. The impact of such openness is complex: while enabling experimentation and learning, too much of it can slow the pace of innovation by producing progressions of hypes, crashes, and paralysis. We set up a tractable growth model that treats openness to novelty, in terms of the general level as well as heterogeneity, as a primitive of the economy. We then use a calibrated version of the model as a laboratory for the analysis of the effects of openness on long-run growth. Among other things, we find that economies that feature a more heterogeneous set of entrepreneurs grow faster. Heterogeneity allows for experimentation and learning without inviting hype cycles.

We see two broad avenues for future research. The first is an investigation into the origins of cross-sectional differences in openness to novelty, both in terms of general level as well as heterogeneity. Through what mechanisms can such differences emerge and persist? How do they relate to differences in risk preferences? A second avenue consists of exploring possible implications for policy. Economic policy, in many ways, affects real-world firms' incentives to experiment—or not to—with potential innovations. The present analysis suggest that the impact of policy may depend on openness to novelty: while policies that strengthen the incentives may pay off handsomely in some economies, their effects may be small or even negative in others. Working out the details will be the task of a different paper.

References

Akcigit, Ufuk and William Kerr (2018). "Growth through Heterogeneous Innovations". *Journal of Political Economy* 126.4, pp. 1374–1443.

Alesina, Alberto et al. (2003). "Fractionalization". *Journal of Economic Growth* 8, pp. 155–194.

Alesina, Alesina, Johann Harnoss, and Hillel Rapoport (2016). "Birthplace diversity and economic prosperity". *Journal of Economic Growth* 21, pp. 101–138.

Arbatlı, Cemal Eren et al. (2020). "Diversity and Conflict". *Econometrica* 88.2, pp. 727–797.

Ashraf, Quamrul and Oded Galor (2013). "The 'Out of Africa' Hypothesis, Human Genetic Diversity, and Comparative Economic Development". *American Economic Review* 103.1, pp. 1–46.

Binswanger, Johannes, Albert Flak, and Manuel Oechslin (2022). “Choice under Fundamental Uncertainty: The Case of Aggregate Consumption”. Unpublished manuscript.

Binswanger, Johannes, Anja Garbely, and Manuel Oechslin (2021). “Economic Forces Shape Beliefs: A Model of Investment under Fundamental Uncertainty”. (Link to paper).

Binswanger, Johannes and Manuel Oechslin (2021). “The Economics of Beliefs under Fundamental Uncertainty”. (Link to paper).

Bolt, Jutta and Jan Luiten Van Zanden (2020). “Maddison style estimates of the evolution of the world economy. A new 2020 update”. (Link to paper).

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2012). “Salience Theory of Choice Under Risk”. *Quarterly Journal of Economics* 127.3, pp. 1243–1285.

Borovička, Jaroslav (2020). “Survival and Long-Run Dynamics with Heterogeneous Beliefs under Recursive Preferences”. *Journal of Political Economy* 128.1, pp. 206–251.

Cao, Dan (2018). “Speculation and Financial Wealth Distribution Under Belief Heterogeneity”. *Economic Journal* 128 (614), pp. 2258–2281.

Caselli, Francesco and James Feyrer (2007). “The Marginal Product of Capital”. *Quarterly Journal of Economics* 122.2, pp. 535–568.

Dedehayir, Ozgur and Martin Steinert (2016). “The hype cycle model: A review and future directions”. *Technological Forecasting & Social Change* 108, pp. 28–41.

Dittmar, Amy and Ran Duchin (2016). “Looking in the Rearview Mirror: The Effect of Managers’ Professional Experience on Corporate Financial Policy”. *The Review of Financial Studies* 29.3, pp. 565–602.

Doepke, Matthias and Fabrizio Zilibotti (2008). “Occupational Choice and the Spirit of Capitalism”. *Quarterly Journal of Economics* 123.2, pp. 747–793.

Falk, Armin et al. (2018). “Global Evidence on Economic Preferences”. *Quarterly Journal of Economics* 133.4, pp. 1645–1692.

Francis, Johanna L. (2009). “Wealth and the capitalist spirit”. *Journal of Macroeconomics* 31, pp. 394–408.

Galor, Oded and Ömer Özak (2016). “The Agricultural Origins of Time Preference”. *American Economic Review* 106.10, pp. 3064–3103.

Gilboa, Itzhak (2009). *Theory of Decision under Uncertainty*. Cambridge University Press.

Gilboa, Itzhak, Larry Samuelson, and David Schmeidler (2014). “No-Betting-Pareto Dominance”. *Econometrica* 82.4, pp. 1405–1442.

Goldfarb, Brent and David A. Kirsch (2019). *Bubbles and Crashes. The Boom and Bust of Technological Innovations*. Stanford University Press.

Gorodnichenko, Yuriy and Gerard Roland (2017). “Culture, Institutions, and the Wealth of Nations”. *Review of Economics and Statistics* 99.3, pp. 402–416.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2006). “Does Culture Affect Economic Outcomes?” *Journal of Economic Perspectives* 20.2, pp. 23–48.

— (2008). “Social Capital as Good Culture”. *Journal of the European Economic Association* 6.2-3, pp. 295–320.

— (2018). “Time varying risk aversion”. *Journal of Financial Economics* 128, pp. 403–421.

Hardy, Bryan and Can Sever (2021). “Financial crises and innovation”. *European Economic Review* 138, Article 103856.

Jones, Charles I. (2019). “Paul Romer: Ideas, Nonrivalry, and Endogenous Growth”. *Scandinavian Journal of Economics* 121.3, pp. 859–883.

Knight, Frank (1921). *Risk, Uncenainty and Profit*. Houghton Mifflin Company.

Malmendier, Ulrike (2021). “Exposure, Experience, and Expertise: Why Personal Histories Matter in Economics”. *Journal of the European Economic Association* 19.6, pp. 2857–2894.

Malmendier, Ulrike and Jon Tate Geoffreyand Yan (2011). “Overconfidence and Early-Life Experiences: The Effect of Managerial Traits on Corporate Financial Policies”. *Journal of Finance* 66.5, pp. 1687–1733.

Michaillat, Pascal and Emmanuel Saez (2021). “Resolving New Keynesian Anomalies with Wealth in the Utility Function”. *Review of Economics and Statistics* CIII.3, pp. 197–215.

Mokyr, Joel (1990). *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press.

— (2017). *A Culture of Growth. The Origins of the Modern Economy*. Princeton University Press.

Montalvo, José G. and Marta Reynal-Querol (2021). “Ethnic Diversity and Growth: Revisiting the Evidence”. *Review of Economics and Statistics* 103.3, pp. 521–532.

Morris, Stephen (1995). “The Common Prior Assumption in Economic Theory”. *Economics and Philosophy* 11, pp. 227–253.

Schmelzing, Paul (2020). “Eight centuries of global real interest rates, R-G, and the ’suprasecular’ decline, 1311–2018”. (Link to paper).

Schwartz, Shalom H. et al. (2001). “Extending the Cross-Cultural Validity of the Theory of Basic Human Values with a Different Method of Measurement”. *Journal of Cross-Cultural Psychology* 32.5, pp. 519–542.

Shiller, Robert J. (2017). “Narrative Economics”. *American Economic Review* 107.4, pp. 967–1004.

— (2019). *Narrative Economics. How Stories Go Viral and Drive Major Economic Events.* Princeton University Press.

Sunde, Uwe et al. (2022). “Patience and Comparative Development”. *Review of Economic Studies* In press.

Appendix

Empirical Motivation.

Table 2: Baseline and Subsets

	Dependent variable: Average annual GDP p.c. growth			
	(1) Baseline	(2) W/o outliers	(3) W/o low income	(4) W/o low & lower middle income
Openness to Novelty	-1.669^{***}	-2.185^{***}	-2.110^{***}	-2.759^{***}
	(0.601)	(0.570)	(0.595)	(0.806)
log(GDP2005)	-1.379^{***}	-1.274^{***}	-1.802^{***}	-1.994^{***}
	(0.214)	(0.168)	(0.242)	(0.367)
Constant	21.815^{***}	22.829^{***}	27.721^{***}	32.133^{***}
	(3.487)	(3.150)	(3.674)	(4.787)
Observations	75	67	68	51
R^2	0.370	0.477	0.470	0.457
Adjusted R^2	0.352	0.460	0.454	0.435
Residual Std. Error	1.869	1.324	1.713	1.790
F Statistic	21.118^{***}	29.131^{***}	28.832^{***}	20.236^{***}

Note: *p<0.1; **p<0.05; ***p<0.01

Table 3: Additional Controls

	Dependent variable: Average annual GDP p.c. growth			
	(1) Baseline	(2) Add. Controls	(3) W/o low income	(4) W/o low & lower middle income
Openness to Novelty	-1.669^{***}	-1.068^{*}	-1.785^{**}	-2.142^{**}
	(0.601)	(0.632)	(0.684)	(0.940)
log(GDP2005)	-1.379^{***}	-1.874^{***}	-2.214^{***}	-2.767^{***}
	(0.214)	(0.327)	(0.358)	(0.600)
Quality of Institutions		0.143	0.388	0.600
		(0.344)	(0.374)	(0.500)
Human Capital		0.263^{**}	0.105	0.107
		(0.109)	(0.128)	(0.180)
Constant	21.815^{***}	21.897^{***}	29.419^{***}	36.296^{***}
	(3.487)	(3.998)	(4.878)	(6.587)
Observations	75	75	68	51
R^2	0.370	0.424	0.493	0.489
Adjusted R^2	0.352	0.391	0.461	0.444
Residual Std. Error	1.869	1.812	1.701	1.775
F Statistic	21.118^{***}	12.894^{***}	15.329^{***}	10.983^{***}

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Table 4: Each Item Separately

	Dependent variable:		
	Average annual GDP p.c. growth		
	(1) Baseline	(2) Item 1	(3) Item 2
Openness to Novelty	-1.669^{***}		
	(0.601)		
Item 1		-1.300^{**}	
		(0.552)	
Item 2			-1.347^{**}
			(0.531)
log(GDP2005)	-1.379^{***}	-1.260^{***}	-1.422^{***}
	(0.214)	(0.208)	(0.224)
Constant	21.815^{***}	19.969^{***}	20.322^{***}
	(3.487)	(3.287)	(3.239)
Observations	75	75	75
R^2	0.370	0.352	0.359
Adjusted R^2	0.352	0.334	0.342
Residual Std. Error (df = 72)	1.869	1.895	1.884
F Statistic (df = 2; 72)	21.118^{***}	19.565^{***}	20.203^{***}
Note:			*p<0.1; **p<0.05; ***p<0.01

Proof of Proposition 1. Consider entrepreneur i's expected end-of-period resources as a function of the provisional time allocation $(\tilde{l}_{it}^o, \tilde{l}_{it}^n)$, conditional on the aggregate supply of the technology good: $\mathbb{E}_{it}\{z_{it}|X_t\} = \mathbb{E}_{it}\left\{(1+r_t)\cdot k_{it} + m_t \cdot x_{it}\left(\tilde{l}_{it}^o, \tilde{l}_{it}^n\right)|X_t\right\}$. Given X_t, both factor prices, r_t and m_t, are fixed magnitudes (equation 2). Therefore,

$$\mathbb{E}_{it}\{z_{it}|X_t\} = (1+r_t)\cdot k_{it} + m_t \cdot \mathbb{E}_{it}\left\{x_{it}\left(\tilde{l}_{it}^o, \tilde{l}_{it}^n\right)|X_t\right\}. \tag{37}$$

Yet the conditional expectation on the right-hand side of equation (37) is independent of X_t:

$$\mathbb{E}_{it}\left\{x_{it}\left(\tilde{l}_{it}^o, \tilde{l}_{it}^n\right)|X_t\right\} = \mathbb{E}_{it}\left\{x_{it}\left(\tilde{l}_{it}^o, \tilde{l}_{it}^n\right)\right\}. \tag{38}$$

It follows that a provisional time allocation $(\tilde{l}_{it}^{o,*}, \tilde{l}_{it}^{n,*})$ that maximizes entrepreneur i's expected output also maximizes $\mathbb{E}_{it}\{z_{it}|X_t\}$ for any value of X_t. By the law of iterated expectations, the unconditionally expected value of z_{it} can be written as

$$\mathbb{E}_{it}\{z_{it}\} = \mathbb{E}_{it}\{\mathbb{E}_{it}\{z_{it}|X_t\}\}. \tag{39}$$

Given this, and since $(\tilde{l}_{it}^{o,*}, \tilde{l}_{it}^{n,*})$ maximizes $\mathbb{E}_{it}\{z_{it}|X_t\}$ for any aggregate supply of the technology good, $(\tilde{l}_{it}^{o,*}, \tilde{l}_{it}^{n,*})$ must be a maximizer of $\mathbb{E}_{it}\{z_{it}\}$, too.

Proof of Proposition 3. First consider the probability of $Q_t = 1$ conditional on Q_{t-1}. If $Q_{t-1} = 0$, we can immediately conclude that Q_t must be equal to one since in period $t-1$ exuberant beliefs cannot occur: $\Pr[Q_t = 1 \mid Q_{t-1} = 0] = 1$. However, provided that $Q_{t-1} = 1$, exuberant beliefs in period $t-1$ followed by paralysis in period t is a possibility. In fact, this situation occurs with probability $\Pr[s_{t-1}^{ex} = 1]\cdot\Pr[F_{t-1} = L]$—as can be inferred from equation (32), taking into account that without heterogeneity $s_{t-1}^{ex} \in \{0,1\}$. Now observe that $\Pr[F_{t-1} = L] = 1 - f$ and that the definition of the belief anchor in equation (10) implies $\Pr[s_{t-1}^{ex} = 1] = \left[1 - (p^h)^\Psi\right]$, where p^h is given by equation (20). As a result,

$$\Pr[Q_t = 1 \mid Q_{t-1} = 1] = 1 - (1-f)\left[1 - (p^h)^\Psi\right]. \tag{40}$$

By the law of iterated expectations, $\Pr[Q_t = 1]$ can be written as

$$\Pr[Q_t = 1] = \Pr[Q_{t-1} = 1]\cdot\Pr[Q_t = 1 \mid Q_{t-1} = 1] \\ + (1 - \Pr[Q_{t-1} = 1])\cdot\Pr[Q_t = 1 \mid Q_{t-1} = 0]. \tag{41}$$

Moreover, since the way entrepreneurs split their time endowments across production methods (Proposition 2) is time-invariant, any two periods $t-1$ and t are a priori perfectly symmetric—with the result that $\Pr[Q_t = 1] = \Pr[Q_{t-1} = 1]$. Taking account of this, and using the two conditional probabilities derived above, we obtain the equation stated in the proposition. ■

Proof of Proposition 4. Again, we start by considering conditional probabilities. If $Q_t = 0$, the chance that in period t the new method is revealed to be a success is zero. However, provided that $Q_t = 1$, this chance is $f\left\{\Pr[s_t^{ex} = 1] + \Pr[s_t^{im} = 1] \cdot \Pr[S_t = H \mid F_t = H]\right\}$, where the expression in braces gives the probability that the entrepreneurs adopt either an exuberant belief or an impartial belief and then receive a positive signal. Now observe that $\Pr[S_t = H \mid F_t = H] = \sigma$ and that the definition of the belief anchor in equation (10) implies

$$\Pr[s_t^{im} = 1] = \left[(p^h)^\Psi - (\bar{p})^\Psi\right] \quad \text{and} \quad \Pr[s_t^{ex} = 1] = \left[1 - (p^h)^\Psi\right], \tag{42}$$

where p^h and $\bar{p}$ are given by, respectively, equation (20) and (26). Combining all this results in

$$\Pr[\text{new method in } t \text{ revealed to be a success}] = \Pr[Q_t = 1] \cdot f\left[1 - (1-\sigma)(p^h)^\Psi - \sigma(\bar{p})^\Psi\right]. \tag{43}$$

Finally, by substituting the expression in equation (33) for $\Pr[Q_t = 1]$ we obtain the equation stated in the proposition. ■

Proof of Proposition 5. The limits follow immediately from $\lim_{\Psi\to 0}(\bar{p})^\Psi = \lim_{\Psi\to 0}(p^h)^\Psi = 1$ and $\lim_{\Psi\to\infty}(\bar{p})^\Psi = \lim_{\Psi\to\infty}(p^h)^\Psi = 0$. In order to prove quasi-concavity, we note that the first derivative being strictly positive/negative is equivalent to

$$\left\{-(1-\sigma)\frac{d[(p^h)^\Psi]}{d\Psi} - \sigma\frac{d[(\bar{p})^\Psi]}{d\Psi}\right\}\left\{1 + (1-f)\left[1 - (p^h)^\Psi\right]\right\} + \left[1 - (1-\sigma)(p^h)^\Psi - \sigma(\bar{p})^\Psi\right](1-f)\frac{d[(p^h)^\Psi]}{d\Psi} > / < 0. \tag{44}$$

By accounting for $d[(\bar{p})^\Psi]/d\Psi = \ln(\bar{p})(\bar{p})^\Psi$ and $d[(p^h)^\Psi]/d\Psi = \ln(p^h)(p^h)^\Psi$, and then rearranging terms, we obtain the equivalent condition

$$\frac{\sigma(1-f)\left[1 - (\bar{p})^\Psi\right] - (1-\sigma)}{\sigma\left\{1 + (1-f)\left[1 - (p^h)^\Psi\right]\right\}} < / > \frac{\ln(\bar{p})}{\ln(p^h)}\left(\frac{\bar{p}}{p^h}\right)^\Psi. \tag{45}$$

In what follows, we refer to the left-hand side and right-hand side of equation (45) as $LHS(\Psi)$ and $RHS(\Psi)$, respectively. As Ψ rises from zero towards infinity, $LHS(\Psi)$ increases from $-(1-\sigma)/\sigma < 0$ to $(-1+2\sigma-\sigma f)/(2\sigma-\sigma f) \in (0,1)$, where the strict positivity of the second limit follows from condition (35) that is stated in the proposition. $RHS(\Psi)$, on the other hand, is a monotonically decreasing function of Ψ, falling from a strictly positive level towards zero as Ψ rises from zero towards infinity.

Pr[new method in t revealed to be a success] is a quasi-concave function if there exists a $\tilde{\Psi}$ such that $LHS(\Psi) < RHS(\Psi)$ for all $\Psi < \tilde{\Psi}$ and $LHS(\Psi) > RHS(\Psi)$ for all $\Psi > \tilde{\Psi}$. A sufficient condition for the existence of such a $\tilde{\Psi}$ is that $LHS(\Psi)$, after turning positive, be a monotonically increasing function of Ψ. To establish that, in fact, $LHS(\Psi)$ is monotonically increasing, suppose it were not. Then, there must exist a threshold $x \in (0,\ (-1+2\sigma-\sigma f)/(2\sigma-\sigma f))$ such that $LHS(\Psi) = x$ has exactly three solutions. However, for any $\tilde{x}$ in the permissible range, $LHS(\Psi) = \tilde{x}$ has at most two solutions. To see this, observe that $LHS(\Psi) = \tilde{x}$ can be rearranged to obtain the equivalent equation

$$\frac{-1+[\sigma+\sigma(1-f)]\,(1-\tilde{x})}{\sigma(1-f)} = (\bar{p})^{\Psi} - \tilde{x}(p^h)^{\Psi}. \tag{46}$$

It is straightforward to show that, as Ψ increases from zero towards infinity, the right-hand side of equation (46) monotonically decreases from the strictly positive level $1-\tilde{x}$ to a negative minimum and then monotonically increases towards zero. Given this pattern, equation (46) can hold for at most two different Ψs. We conclude that the implied threshold x cannot exist—which is a contradiction that establishes the monotonicity of $LHS(\Psi)$. As a result, Pr[new method in t revealed to be a success] is quasi-concave function that monotonically increases on $[0,\tilde{\Psi})$ and monotonically decreases on $(\tilde{\Psi},\infty)$. ∎