When Does Impatience Pay? Social Mobility and Labour Market Returns*

Paul Baker†
Siyu Xu†

This paper studies how social mobility shapes parental expectations of the labour-market returns to patience, which in turn influence their incentive to cultivate patience in children. We develop a model in which the value of patience depends on the upward opportunities available for long-term investments to translate into educational and occupational advancement. When mobility is high, patient workers benefit more from investing and searching for better job opportunities, which raises parental incentives to instil patience. When mobility is low, by contrast, workers face fewer opportunities for upward advancement and impatience can be adaptive. In a rationed labour market, impatient workers require lower incentive-compatible wages and are therefore employed before patient workers. The resulting distribution of patience shifts the composition and productivity of future workers and entrepreneurs, which reshapes labour demand and mobility opportunities for the next generation, potentially generating a dynamic that can sustain multiple steady states with different patterns of social hierarchy and mobility. We further discuss how automation-induced productivity shocks can polarise the intergenerational formation of patience by hollowing out the middle-skill occupational ladder, reinforcing impatience among those whose expectations for upward mobility are diminished. This framework complements behavioural economics' traditional emphasis on psychological and cognitive mechanisms by showing how time preferences can also be socially structured by economic opportunities.

*We thank John Van Reenen for helpful comments and discussions, and conference participants at the Royal Economic Society Annual Conference for suggestions. All remaining errors are our own. The views expressed in this paper are those of the authors and do not necessarily reflect the views of any affiliated institution. No external funding was received for this research. The authors declare that they have no potential conflicts of interest.

†School of Management, University of Bath, Bath, United Kingdom.

Taste classifies, and it classifies the classifier. —*Pierre Bourdieu*

I. Introduction

Patience is generally viewed as a desirable economic preference for human-capital accumulation because patience increases an individual's willingness to invest in skills and persist in long-term plans, all of which are central to labour market success (e.g., Heckman, Stixrud, and Urzua, 2006; Golsteyn, Grönqvist and Lindahl, 2014; Hai and Heckman, 2022; Epper et al., 2026). The idea linking patience and labour-market advancement goes back at least to Max Weber's account of the "spirit of capitalism," in which disciplined work and restraint in consumption arose from religious doctrine and in turn shaped the conduct of work and the pursuit of gain (Weber, 1958).[1] This cultural perspective, which emphasises how values shape individuals' responses to economic opportunities, contrasts with Karl Marx's emphasis on material economic conditions as the underlying force shaping individuals' values (Marx, 1859). Existing work on the formation of patience largely follows cultural and economic explanations, emphasising cultural inheritance from the previous generation (Bisin and Verdier, 2001) or the social class position parents currently occupy (Doepke and Zilibotti, 2008). A more recent literature points to genetic endowments that parents can partially observe in their children, leading them to adjust their investments in children's noncognitive traits accordingly (Belsky et al., 2016; Papageorge and Thom, 2020).

Yet these perspectives are status-quo-based[2], linking patience formation to families' current social class, material circumstances or endowments. This implies similar incentives to cultivate patience among families in similar positions. However, previous explanations of patience formation overlook an important economic incentive for parental investment—expectations of

[1] Weber's (1958) account of the Protestant ethic offers this mechanism. Under Calvinist predestination, salvation was fixed in advance and unaffected by conduct; the resulting anxiety was resolved by reading worldly success as a sign of election rather than a means to it. Disciplined, future-oriented behaviour thus arose from doctrine rather than from material incentives.

[2] The definition of status-quo-based follows the behavioural literature on status quo bias, where the current state serves as a reference point and departures from it are evaluated as gains or losses (Samuelson and Zeckhauser, 1988; Kahneman, Knetsch, and Thaler, 1991).

upward mobility (Dizon-Ross, 2019). Empirical evidence shows that similar initial conditions may not produce similar outcomes when opportunity structures differ. For instance, Tucker-Drob and Bates (2016) show that the interaction between genetic endowments and social class in shaping intelligence is stronger in the United States, where developmental resources are more unequally distributed, than in the more equal settings of Western Europe and Australia. Given that upward-mobility expectations are higher in more unequal settings than in more equal countries (Borgonovi and Pokropek, 2025), differences in opportunity structures may generate different parental investment incentives. More generally, because social mobility varies systematically across institutional and labour-market environments (e.g., Bénabou and Ok, 2001; Chetty et al., 2014; Doepke and Zilibotti, 2017), the return parents expect from cultivating patience may depend not only on their current class position, material circumstances, or endowments, but also on the future opportunities they expect their children to face. We therefore depart from a status-quo-based view of patience formation and focus instead on expectations of future mobility. This leads to our research question: how does social mobility shape parents' expectations of the labour-market returns to patience and, in turn, their incentives to cultivate patience in their children? Given the documented positive relationship between patience and aggregate productivity (Sunde et al., 2022), a further question is whether differences in patience formation generate systematic differences in aggregate productivity, feeding back into labour-market conditions and future mobility.

To answer these questions, we develop a model that links a macro labour market to a micro household decision. We first characterize the labour market, where labour demand, wage-setting, productivity, and social mobility operates through the labour-market conditions that govern the return to patience. Workers can be either patient or impatient. Patient workers are more productive in production tasks, but they hold higher reservation wages as they gain higher human capital and weigh future outcomes more heavily. We then embed these labour market returns in a micro-level model of patience formation. Following Becker and Mulligan (1997) and Doepke and Zilibotti (2008), we treat patience as an endogenously formed preference rather than a fixed endowment, with parental investment increasing the probability that a child develops into a patient type. Parents judge whether that investment is

worthwhile from the degree of social mobility, which signal how far patience is likely to be rewarded in the child's future labour market. Households differ in how much these signals bear on their decision. For elite entrepreneurial families, whose children can enter the family enterprise, the return does not depend on being hired, so these families always choose to invest regardless of conditions in the wage labour market. Wage-worker families, by contrast, face a trade-off. Conditional on employment, patient children earn more because they require higher incentive-compatible wages. When labour demand is limited and employment is rationed, however, impatient workers can be employed at lower incentive-compatible wages and are therefore selected into employment before patient workers.

The key question in the model turns on how wage-worker parents respond to these labour market signals. Motivated by List et al., (2021), who show that only a sufficiently large shift in expectations translates into changes in parental behaviour, we model the micro-level household response to the macro labour-market conditions operate through two stages: expectation and action. In the first stage a high social mobility environment raises parents' expectation-based reference point[3] for their children's upward mobility. Crucially, this reference point is forward-looking and anchored on what parents expect their children to be able to achieve, rather than on the family's current status quo. In this way, children's labour market outcomes falling below this benchmark are therefore experienced as losses, which strengthens the parents' motivation to act. In the second stage, conditional on the first stage's expectation, parents decide whether to instil patience through costly cognitive effort. For wage-worker families, social mobility shapes the extent to which children can access high-wage labour markets. When social mobility falls below the threshold, impatience can function as an adaptive preference: impatient labourers more readily accept existing wage offers and exert effort under limited incentives, whereas patient labourers demand stronger incentives and are therefore disadvantaged. When high-wage opportunities are accessible, by contrast, these incentives are provided automatically, so the relative disadvantage of patience disappears. Hence parents who anticipate that their children will remain in low-wage labour markets may rationally refrain from costly investment in patience, allowing impatience to persist as an adaptive response

[3] See Freeman (2019).

to constrained labour market opportunities. Meanwhile, patience also carries a present cost for parents: patient children are less likely to contribute to within-family public goods today (Baland and Robinson, 2000), instead deferring gratification toward future-oriented activities such as higher education. The action trade-off is therefore that the expected economic gain from cultivating patience must outweigh both the cognitive cost of parental investment and the loss of current contributions to within-family public goods.

These investments in patience may in turn influence the demand side of the labour market through the composition and productivity of entrepreneurs in the economy. When wage-worker families invest less in cultivating patience, fewer children develop the time preferences associated with high-productivity entrepreneurship, reducing the future share of productive entrepreneurs and potentially keeping labour demand low.

In the baseline model, productivity is given conditional on entrepreneurial type, so patience formation affects aggregate productivity through the composition of entrepreneurs rather than by changing type-specific productivity itself. We then extend the model to consider technological change. Motivated by Doepke and Zilibotti (2008), who study how the Industrial Revolution changed the economic returns to class-specific preferences, we model automation as a productivity shock that simultaneously raises firm productivity and displaces labour from automatable tasks. Its effect on labour demand is therefore ambiguous: productivity gains may expand firm scale and create new jobs, while task displacement may eliminate existing middle-skill progression routes (Acemoglu and Restrepo, 2019; Acemoglu and Loebbing, 2026). When displacement dominates mobility opportunity creation, the middle of the occupational ladder contracts and the returns to patience diverge across families. Those that retain access to high-return opportunities have stronger incentives to cultivate patience, whereas those whose children are increasingly confined to low-wage employment have weaker incentives to invest in patience formation, reinforcing the persistence of impatience.

Importantly, one caveat of our model is that the intergenerational transmission of patience may partly reflect the inheritance of cognitive ability rather than the parents' active transmission of preferences. Given cognitive ability is correlated with patience (e.g., Benjamin et al., 2013), and cognitive ability itself is transmitted across generations (e.g., Anger and Heineck, 2010),

then the observed parent-child correlation in patience could partly capture the transmission of cognitive endowments. However, recent experimental evidence argues that patient choices do not require cognitive capacity alone; they also require individuals to exert cognitive effort by representing future utility, keeping delayed rewards salient, and inhibiting the appeal of immediate gratification (Bortolotti et al., 2021). Because cognitive effort is costly, individuals may exert it only when the perceived stakes are sufficiently high. This distinction is crucial to our argument. Parents may possess the cognitive ability to understand the value of patience, but whether they actively cultivate patience in their children depends on whether they expect patience to generate large future returns. This is consistent with experimental evidence from parental-belief interventions that providing information and guidance can change parents' beliefs about the effects of parental investment. However, simple educational interventions often produce limited behavioural change, whereas only more intensive belief shifts in parental beliefs can alter parental investment decisions (List et al., 2021). In our framework, parents respond to economic incentives only when their beliefs about the expected payoff to cultivating patience exceed a sufficiently high threshold. Therefore, rather than treating patience transmission as a consequence of inherited cognitive ability, our model allows cognitive ability to condition parents' endogenous investment in patience formation. Parents are more willing to cultivate patience when children's cognitive ability raises the expected return to long-horizon human-capital investment, consistent with the idea that parental investment responds to signals of children's ability (Cunha et al., 2026). Patience, in turn, supports human-capital accumulation by increasing children's ability to persist in schooling and long-term investments (Sutter et al., 2013; Golsteyn et al., 2014).

Our theoretical contribution is to develop a socially structured framework for time-preference formation that complements behavioural economics' traditional emphasis on psychological and cognitive mechanisms (Mullainathan and Thaler, 2001; DellaVigna, 2009). In our model, we show that labour market conditions and social mobility shape time preferences by determining the expected returns to patience and, in turn, parents' incentives to cultivate patience in their children. This formalizes Bourdieu's theory of field and habitus that individuals' orientation toward the future reflects the interaction

between their dispositions and "the chances objectively offered to them by the social world" (Bourdieu, 1990).[4] From this perspective, variation in patience need not be viewed simply as a departure from rational behaviour, but may instead reflect adaptation to differences in expected economic opportunities. Our framework complements the economics of culture, particularly Acemoglu and Robinson's (2021) and Acemoglu's (2022) view of cultural configurations as fluid and responsive to economic incentives, by providing a formal bridge between macroeconomic opportunity structures and the behavioural economics of preference formation.

This paper also contributes to the growing literature on the origins of time preferences (e.g., Doepke and Zilibotti, 2008; Falk et al., 2021; Chowdhury et al., 2022). Much of this literature emphasizes the intergenerational transmission of preferences within families. Doepke and Zilibotti (2008), for example, develop a theory of preference formation under financial market imperfections, arguing that middle-class families engaged in occupations requiring sustained effort, skill, and experience endogenously cultivate patience and a strong work ethic, whereas upper-class families relying on rental income foster a culture oriented toward leisure. More recently, Agostinelli, Doepke, Sorrenti, and Zilibotti (2025) extend this framework by emphasizing how parents optimally choose which traits to instil in their children in response to economic incentives. We build on this insight but shift attention from parents' current occupational position to the forward-looking opportunities they expect their children to face. In our framework, social mobility changes the expected labour market return to patience and therefore changes the incentive to invest in its transmission.

A central implication is that the intergenerational transmission of impatience can be broken when the expected return to cultivating patience is sufficiently high. This possibility is consistent with evidence that authoritative parenting (high control and high responsiveness) can counteract the intergenerational transmission of impatience (Brenøe and Epper, 2022). Further evidence that patience is malleable comes from Alan and Ertac (2018), who

[4] Bourdieu (1990) describes this interaction explicitly: "In fact, a given agent's practical relation to the future, which governs his present practice, is defined in the relationship between, on the one hand, his habitus with its temporal structures and dispositions towards the future . . . and on the other hand a certain state of the chances objectively offered to him by the social world."

implement a randomized intervention promoting forward-looking behavior among children and show that it causally increases patience, with long-lasting effects. Therefore, we assume that parenting effort shifts the probability that a child develops patience, while the parent's own type increases, but does not determine, the likelihood of same-type transmission. When expected upward mobility is low, an impatient parent has little incentive to incur the cost of cultivating patience, so impatience is more likely to persist across generations. When expected mobility is sufficiently high, however, the anticipated return to patience can induce even an impatient parent to invest actively in patience formation, increasing the probability that the child becomes patient and thereby weakening—and, when the response is sufficiently strong, potentially breaking—the intergenerational persistence of impatience. Our contribution is therefore to endogenize the strength of preference transmission itself by embedding parenting choices in a labour-market equilibrium in which social mobility, labour demand, and employment opportunities determine the expected return to cultivating patience.

Understanding the role of parental expectations is also important from a policy perspective. In our framework, parents' investment in cultivating patience is driven not simply by the belief that patience is desirable, but by the expectation that it will improve their children's future outcomes. This implies that policies aimed at strengthening patience transmission should not only encourage parents to value delayed gratification, but also alter the perceived payoff to doing so. Information, guidance, and school-to-work support may matter because they make future educational and occupational pathways more visible and credible, thereby raising parents' expectations about the returns to cultivating patience. However, our model also implies belief-based interventions are unlikely to be effective when the underlying opportunity structure remains weak. If local labour markets offer low mobility, low returns to long-horizon investment, or few credible advancement opportunities, parents may rationally remain reluctant to bear the short-run costs of instilling patience. This matters for inequality because weak mobility environments may discourage disadvantaged families from investing in the traits – patience – that support human-capital accumulation, thereby reinforcing intergenerational persistence. Effective policy therefore requires both shifting parental beliefs

and improving the real opportunity structure through which patience can be converted into human-capital investment.

The rest of the paper is organized as follows. Section II presents motivating evidence on the relationships among social mobility, parents' efforts to instil patience, and labour market structure. Section III develops the baseline model, linking labour market equilibrium and social mobility to parents' incentives to cultivate patience and characterizing the resulting intergenerational dynamics. Section IV extends the framework to technological shock, examining how automation affects the returns to patience through labour demand and the accessibility of upward mobility opportunities. Section V concludes.

II. Motivating Evidence

The motivating evidence is organized from macroeconomic conditions to micro-level preference formation. We begin with the aggregate relationship the model seeks to explain: whether societies with greater social mobility exhibit greater patience. We then ask whether this relationship varies with occupational structure, as social mobility should have a stronger effect on the returns to patience among wage-employed households, where returns depend more on formal career progression. Moving to the household level, we examine whether more patient workers have higher reservation wages, providing a microeconomic foundation for the model's employment and rationing mechanism. Finally, we turn to the household and explore whether parents who expect greater upward mobility for their children make stronger efforts to instill patience. Note that this evidence is descriptive rather than causal and is intended to motivate the key relationships formalized in the model.

II.A. Patience and Social Mobility

Figure I presents cross-country and within-country evidence on the relationship between social mobility and patience. Panel A plots residuals from separate regressions of country-level measures of patience from the Global Preferences Survey[5] (Falk et al., 2016, Falk et al., 2018) and the

[5]From the Global Preferences Survey (2012). https://gps.econ.uni-bonn.de/home

Global Social Mobility Index[6] on GDP per capita and risk-taking. We control for GDP per capita to separate mobility from resource endowments and liquidity constraints (Sunde et al., 2022), for risk-taking because the extent to which individuals pursue mobility opportunities may itself depend on their willingness to take risks (Bonin et al., 2007; Jaeger et al., 2010). Panel A shows countries with higher-than-predicted mobility also exhibit higher-than-predicted patience (n=49, r=0.42, p=0.012). As cross-country comparisons may be confounded by substantial differences in institutions, culture, and other country-specific factors, Panel B turns to within-country variation across U.S. states. In Panel B, state-level investing horizons and absolute income mobility are separately residualized on the gender, age, education, and income composition of the state sample, as well as state-level risk-taking. States with higher-than-predicted absolute mobility tend to exhibit longer-than-predicted investing horizons (n=46, r=0.30, p=0.045). Taken together, the two panels show that greater social mobility is associated with both greater patience and more future-oriented economic planning, conditional on relevant economic and demographic characteristics.

II.B. The Mobility–Patience Link Depends on Labour-Market Structure

Following Doepke and Zilibotti (2008), we allow the returns to patience to vary across occupational settings and distinguish between wage-employment and self-employment households. For wage-employment households, the return to patience depends on whether formal labour markets provide opportunities for advancement through education, training, and career progression. DellaVigna and Paserman (2005), for example, show more impatient workers set lower reservation wages and are therefore more likely to accept a job conditional on receiving an offer. For self-employment households, patience may instead be rewarded through delayed returns and breakthrough innovations (Hirshleifer et al., 2012). As household-level self-employment is not consistently available in the cross-country data, Figure II uses country-level self-employment rates[7] to capture the broader occupational structure shaping

[6] From the World Economic Forum (2020). https://www.weforum.org/publications/global-social-mobility-index-2020-why-economies-benefit-from-fixing-inequality/

[7] From the World Bank indicator "Self-employed, total (% of total employment) (modeled ILO estimate, which reports the share of self-employed workers in total employment. https://data.worldbank.org/indicator/SL.EMP.SELF.ZS

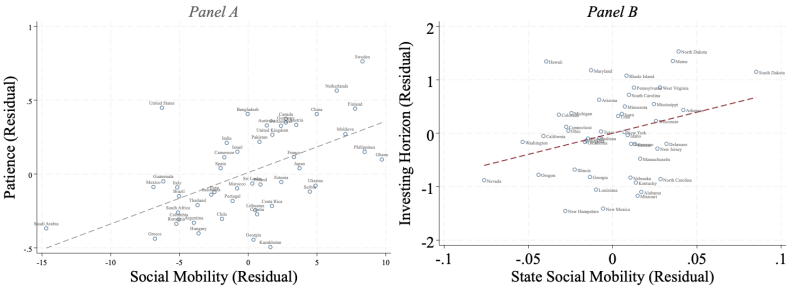


Figure I
Country-level Patience and Social Mobility Score and International Income Mobility

Notes. Panel A plots country-level measures of patience against an objective index of social mobility. The patience index is constructed using data from the Global Preferences Survey and aggregated at the country level. Objective social mobility is measured by the Global Social Mobility Index (GSMI) developed by the World Economic Forum (2020). The GSMI is a composite index designed to capture countries' capacity to generate intergenerational mobility by benchmarking the structural conditions that enable equality of opportunity. The index aggregates 51 indicators grouped into ten pillars: health, education access, education quality and equity, lifelong learning, technology access, work opportunities, fair wage distribution, working conditions, social protection, and inclusive institutions. Each pillar receives equal weight, and the overall index ranges from 0 to 100, with higher values indicating greater scope for upward social mobility. GDP per capita is measured in PPP-adjusted terms for 2019, expressed in constant 2021 international dollars, using data from the World Bank's World Development Indicators based on the International Comparison Program (worldbank.org/en/programs/icp/data). Country-level risk-taking is constructed from the Global Preferences Survey and aggregated in the same way as the patience measure.Each observation represents a country. In Panel B, state-level absolute mobility is calculated as the average of available birth-cohort-specific estimates for the 1940, 1950, 1960, 1970, and 1980 cohorts from Chetty et al. (2017). Absolute mobility is defined as the proportion of children whose household income at age 30 exceeds their parents' household income at age 30, after adjusting for inflation. Investing horizons are obtained from the Federal Reserve Bank of New York Survey of Consumer Expectations Planning Horizons Study, which measures respondents' reported time horizon when thinking about or planning investment decisions. The planning-horizon variables are coded on a 1-9 scale, where larger values indicate longer horizons: next day, next week, next two weeks, next month, next few months, next year, next few years, next 5-10 years, and longer than 10 years. Each observation in Panel B represents a U.S. state. The dashed lines report fitted linear relationships between the residualized variables.

these incentives. Panel A of Figure II shows that the association between social mobility and patience depends on the prevalence of self-employment. In countries where self-employment is less common, patience rises sharply with social mobility, consistent with the idea that its return depends more strongly on access to education, training, and formal career progression. By contrast, in countries with a high share of self-employment, average patience varies little with social mobility. In these settings, patience may already be rewarded by the occupational structure itself, because self-employment involves delayed returns, long-horizon investment, and the need to sustain effort before payoffs are realised. Social mobility therefore adds less to the return to patience, producing a weaker mobility–patience gradient. Panel B of Figure II uses parents' reported importance of instilling tolerance and respect for others, drawn from the World Values Survey (wave 7, 2017-2022, Haerpfer et al., 2022), as a proxy for the relational component of long-horizon child traits. Although the data do not directly measure parents' desire to instil patience, tolerance and respect capture traits that support sustained cooperation and repeated interaction, which are increasingly rewarded in modern labour markets (Deming, 2017). Doepke and Zilibotti (2008) similarly argue that parents who expect their children to rely on labour income have stronger incentives to instil tolerance for sustained effort. This interpretation is particularly relevant for self-employment, where returns often depend on long-term relationships with customers, suppliers, and collaborators (Kwon et al., 2013). Consistent with this idea, Panel B of Figure II shows that in low-self-employment economies, parents place greater importance on instilling tolerance and respect when social mobility is high. By contrast, in high-self-employment economies, this parental value varies much less with social mobility, as self-employment itself requires patience through delayed returns regardless of mobility conditions. Together with Panel A of Figure II, this pattern supports the broader argument that the cultivation of patience and related long-horizon traits depends not only on social mobility conditions, but also on the occupational structure through which these traits generate returns.

II.C. Patience and Wage

Figure III examines the relationship between forward-looking financial planning and reservation wages. The horizontal axis measures respondents'

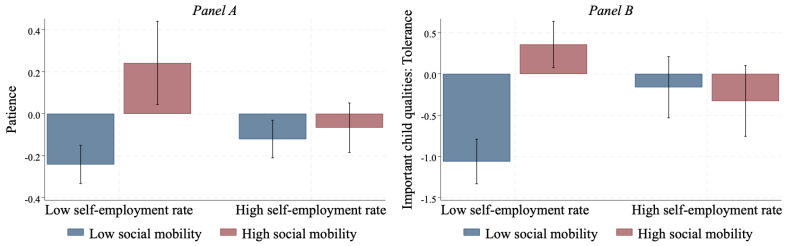


Figure II

Social Mobility and Patience under Different Occupational Structures

Notes. In Figure II, countries are divided into high- and low-self-employment groups according to whether their self-employment rate is above or below the sample mean. Self-employment is measured using the World Bank indicator "Self-employed, total (% of total employment)," based on modeled ILO estimates. Countries are similarly classified into high- and low-social-mobility groups using the sample mean of the Global Social Mobility Index. Panel A reports mean country-level patience from the Global Preferences Survey across the four resulting groups. Panel B reports the corresponding mean importance parents attach to instilling "tolerance and respect for other people" as a child quality, using data from World Values Survey Wave 7. Bars show group means and vertical lines denote 95% confidence intervals.

average investing and retirement planning horizon, while the vertical axis reports the mean log annual reservation wage at each horizon. A clear positive relationship emerges: individuals with longer planning horizons tend to require higher wages to accept a new job (n=869, r=0.226, p=0.001). This evidence provides empirical motivation for a key assumption of our model. If longer planning horizons proxy greater patience or future orientation, the positive association with reservation wages suggests that more patient workers require higher compensation before accepting employment. This higher acceptance threshold is central to the model's labour market mechanism: when labour demand is weak, patient workers are more likely to be rationed out of employment, whereas impatient workers, with lower reservation wages, are more readily employed.

II.D. Upward Expectations and Patience Instilling

Finally, we test whether parents' expectations about their children's upward mobility predict the patience they cultivate, using the Programme for International Student Assessment (PISA 2022). PISA, a triennial OECD assessment of representative samples of 15-year-old students, conducted in 57 countries since 2000. It surveys students and, in participating systems, their

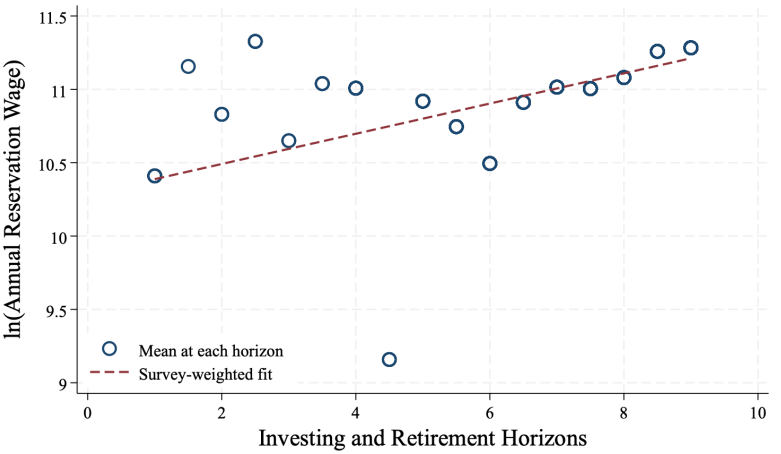


Figure III
Planning Horizon and Reservation Wage

Notes. Data are drawn from the Federal Reserve Bank of New York Survey of Consumer Expectations (SCE) Planning Horizons Study and Labor Market Survey. Respondents in the September 2023 Planning Horizons Study are matched to their nearest available Labor Market Survey observation. The planning-horizon measure is the respondent-level mean of investing and retirement horizons. Each item ranges from 1 to 9, where 1 denotes the next day and 9 denotes more than 10 years. "Does not apply to me" responses are treated as missing. Reservation wage is measured as the lowest pre-tax wage or salary a respondent would accept for a job in a line of work they would consider. Annual reservation wages below or equal to $1,000 or above $1,000,000 are excluded, and the remaining values are log transformed. Each circle reports the mean log reservation wage at the corresponding planning-horizon value. The dashed line is fitted using an individual-level linear regression with FRBNY respondent weights.

parents, eliciting children's behavioural dispositions, parents' engagement in their children's development, and detailed socioeconomic information.

We derive a household's upward-mobility expectation from the PISA 2022 parent questionnaires. We measure parents' current occupational standing by converting each parent's occupation into an International Socio-Economic Index of Occupational Status (ISEI) score and, where both parents are employed, assigning the household the higher of the two (HISEI). We measure parents' aspiration for their child from the question asking what occupation they expect their child to hold in the future, which we likewise convert into an expected ISEI score (PAREXPT). We define a household's *Occupation upward expectation* as the gap between the two, PAREXPT- HISEI: a higher value means parents expect their child to attain an occupation ranked further above the family's current standing.

We construct a *Children's Patience Index* by averaging each student's standardized answers to - on a five-point scale from "strongly disagree" to "strongly agree," - items such as "I stop to think before acting," "I think carefully before doing something," and, reverse-coded, "I rush into activities without thinking" and "I am more impulsive than most people I know" ($\alpha = 0.753$) . We use these items as a proxy for patience, motivated by the close conceptual link between patience, self-control, and impulsivity in the literature on intertemporal choice (e.g., Stevens et al., 2008; Alan, 2015). A higher value indicates a more patient, less impulsive child.

We construct a *Patience Instilling Index* by averaging standardized responses to the future-oriented cultivation items, which ask parents how often, from "never or hardly ever" to "every day or almost every day," they engage in activities such as talking to their child about the importance of completing upper-secondary education, talking to their child about his/her future education, and encouraging their child to get good marks ($\alpha = 0.812$). We interpret these items as a proxy for parental efforts to cultivate future-oriented behaviour. This interpretation is motivated by Alan and Ertac (2018), who show that interventions encouraging children to imagine future selves and adopt forward-looking behaviour increase patience. A higher value indicates more intensive cultivation of future-oriented behaviour.

We control for family income to separate upward expectations from current household resources, and for GDP per capita to separate mobility

expectations from broader differences in national prosperity. Student age and gender account for demographic variation in children's behaviour and parental involvement.

Table I
Upward Expectations and Patience

	(1) Parent instill patience	(2) Children's patience
Occupation upward expectation	0.0812***	0.1193***
	(0.0175)	(0.0158)
Family income	-0.0254***	0.0067***
	(0.0019)	(0.0017)
Age	-0.0668***	0.1225*
	(0.0070)	(0.0111)
Female	-0.0066	0.1225***
	(0.0070)	(0.0064)
GDP per capita	-0.0005***	-0.0007***
	(0.0000)	(0.0003)
N	40,838	36,885
R^2	0.0333	0.0133

Notes. We report OLS regressions of standardised patience measures on parents' occupation-based upward expectation and controls. Standard errors in parentheses. Significance levels: *** $p < 0.01$.

Table I reports non-causal regressions of the two patience measures on parents' upward expectation. Parents who expect greater upward mobility for their children cultivate patience more intensively (column 1) and have children who are themselves more patient (column 2), both significant at the 1% level. The association is present after conditioning on family income and on country-level GDP per capita, so it does not simply reflect that richer families or richer countries invest more. Notably, family income enters column (1) with a negative coefficient. Conditional on parents' upward expectations, higher-income families do not invest more heavily in future-oriented parenting. This pattern is difficult to reconcile with a simple resource-based account in which wealthier families cultivate patience because they have greater capacity to invest in their children. Instead, the positive coefficient on upward expectations suggests that what matters is the perceived scope for the child to improve on the family's current position. Parents appear to invest more when they expect patience to help their child move further upward, rather than simply when the household is richer. These correlations therefore motivate our modelling choice that parental investment responds to the expected return

to upward mobility, rather than to the family's current socioeconomic position alone.

III. Model

III.A. Labour Market Environment

Following Acemoglu (2022) and the occupational-choice approach of Doepke and Zilibotti (2008), we distinguish between entrepreneurship and production work. This structure allows patience to generate different returns across occupations: workers depend on wages and employment opportunities, while entrepreneurs determine labour demand and offer a route for upward mobility. It therefore provides a simple structure through which household preference formation can feed back into labour-market conditions and social mobility. Individuals who are not in entrepreneurship or production work are unemployed and receive zero income (as an innocuous normalization). Each individual (h) in generation (t) is characterized by one of two time-preference types, $T_t^h \in \{I\,, \quad P\}$, where I indicates impatience and P indicates patience. We assume an individual's time preference is public knowledge. This assumption allows the type-specific incentive constraints to translate directly into employment decisions.

Motivated by evidence that patience is positively associated with productivity (Sunde et al., 2022), we assume that as a production worker, each impatient individual can supply one efficiency unit of labour, while each patient worker can supply $\gamma > 1$ efficiency units of labour (i.e. such a worker is γ times more productive per hour as an impatient worker). Both types of workers are also subject to a moral hazard problem. Specifically, an individual can "exit" and "search another job", in which case she will not be productive. Motivated by Topel and Ward (1992), who show that job mobility is an important source of wage growth in workers' careers, we assume that a worker search for another job will receive an additional benefit $\upsilon > 0$ and will fail to secure a new job with probability $\delta \in (0,\ 1)$. In addition to the difference in productivity, the two types also differ in their labour market behaviour. These differences follow from a single mechanism. Search effort is costly today and pays off only in the future, so an impatient worker who discounts that payoff

more heavily searches less intensively. Waiting involves the same trade-off: rejecting an offer in order to hold out for a better one sacrifices current earnings for a delayed gain, which makes it less attractive to the impatient and gives them lower reservation wages. DellaVigned and Paserman (2005) document both patterns among unemployed job-seekers.The same trade-off governs employed workers, for whom separating means giving up current earnings for uncertain future prospects. Impatient workers are therefore less likely to exit. We model this by assuming that impatient workers incur an additional cost $\kappa \in (0,\ \upsilon)$ when they exit, while patient workers do not. When either type of worker exerts effort (rather than shirking), every efficiency unit of labour has the same productivity. Finally, we assume that there is limited liability, thus workers cannot be punished with negative wages.

We first derive the incentive compatibility constraints for patient and impatient workers. Since the marginal product of each efficiency unit of labour (conditional on non-shirking) is the same, it is convenient to formulate these constraints in terms of wage per efficiency unit of labour, w. A remark on the wage structure is in order, because it underlies the entire aggregation that follows. We have written a single wage per efficiency unit, w_t, that applies to the labour hours of patient and impatient workers alike. This is not an assumption imposed for convenience but an equilibrium property of a competitive market. Conditional on incentive compatibility — that is, once each type is paid at least its reservation wage so that no one shirks — one efficiency unit of labour supplied by a patient worker and one efficiency unit supplied by an impatient worker are perfect substitutes in production: each contributes identically to output. If patient efficiency units commanded a strictly higher price per unit, every entrepreneur would substitute entirely toward impatient efficiency units, and conversely. Hence both kinds of efficiency units must trade at the same price w_t whenever both are employed. The productivity advantage of patient workers therefore operates through the *quantity* of efficiency units they supply per hour — a patient worker supplies $\gamma > 1$ efficiency units while an impatient worker supplies one — and not through a separate price for their labour.

Then the incentive compatibility constraint for patient workers is,

$\gamma w \geq (1 - \delta)\gamma w + \upsilon.$

The left-hand side is the wage that the worker will get if he does not exit. The right-hand side is his expected utility if he exits: with probability $1-\delta$ he does not have search friction and receives the total wage γw (γ units of labour times wage per efficiency unit of labour, w), as well as the utility from exiting υ; and with probability δ he has search friction and receives just the utility from exiting. Rearranging, we obtain the incentive compatibility constraint for patient types:

$w \geq \underline{w}_P \equiv \frac{\upsilon}{\delta\gamma}$, (IC patient)

which specifies the minimum wage ($\underline{w}_P$) that a patient worker needs to be paid so that he does not exit and defines as the minimum wage at which a patient individual could be employed as a production worker.

The incentive compatibility constraint for impatient workers is similar except that such a worker has one efficiency unit of labour and incurs the additional cost κ if she exits. This implies a similar incentive compatibility constraint for patient workers:

$$\begin{aligned} w &\geq (1-\delta)w + \upsilon - \kappa \\ &\geq \underline{w}_I \equiv \tfrac{\upsilon-\kappa}{\delta}. \end{aligned} \quad \text{(IC impatient)}$$

This equation defines $\underline{w}_I$, analogously as the wage at which an impatient individual can be employed as a production worker. The assumption that $\kappa < \upsilon$ implies that an impatient worker, working at the minimum wage consistent with non-exiting, $\underline{w}_I$, will still receive a positive wage.

It is clear that in equilibrium, both incentive compatibility constraints, (IC patient) and (IC impatient), will have to hold, such that there will be no additional benefits from searching for another job. Expected stage utilities, the u_t^h terms, are therefore simply equal to labour market earnings of workers or the total incomes of entrepreneurs.

We next impose a crucial assumption for the rest of our analysis.

Assumption 1.

$$\kappa > \frac{\gamma - 1}{\gamma}\upsilon.$$

Assumption 1 requires the additional search cost faced by impatient workers to be sufficiently large relative to the gains from searching, after accounting for the productivity advantage of patient workers. This assumption implies that

$$\underline{w}_P > \underline{w}_I,$$

Thus, even taking into account the advantages that patient workers have - such as higher reservation wages, and more intensive job search - the moral hazard problem associated with their search effort is worse than that of impatient workers. Specifically, the wage required to ensure that patient workers exert sufficient search effort, $\underline{w}_P$, is higher than the incentive-compatible wage for impatient workers, $\underline{w}_I$. This generates a key trade-off in our framework: from an efficiency perspective, patient workers may be preferable due to their greater attachment and productivity in job search; however, employing patient workers necessitates paying higher wages to overcome search frictions and ensure incentive compatibility. Therefore, depending on overall labor market conditions, equilibrium may still favour impatient workers.

Having established the differences between patient and impatient individuals in the labour market as workers, we now turn to analyze their differences in entrepreneurship. Motivated by evidence that entrepreneurship requires long-horizon investment and delayed consumption (Doepke and Zilibotti, 2013; Dawson and Failla, 2022), patient individuals are assumed to be more productive as entrepreneurs ($i.e.,\ A_P > A_I > 0$). In particular, we assume that an individual of type T_t^hhas access to the production technology

$A_T l_t^{\alpha}$,

where l_t denotes the number of production workers hired at time t, and $\alpha \in (0, 1)$ which indicates diminishing returns.

III.B. Labour Demand and the Productivity Ranking

Having specified the incentive-compatibility constraints that determine the minimum wage at which each type of worker can be employed, we now turn to the demand side of the labour market. The purpose of this subsection is threefold. First, we derive the labour demand of an individual entrepreneur from profit maximization and show why, in a competitive market, a single wage per efficiency unit must prevail. Second, we establish the productivity ranking between patient and impatient entrepreneurs, which will later determine who is hired first under rationing and who finds it profitable to remain in entrepreneurship. Third, we aggregate individual demands to obtain the

economy-wide labour-demand schedule and isolate the first macroeconomic feedback of the model: the dependence of aggregate labour demand on the composition of the entrepreneurial pool.

Because labour markets are competitive, the entrepreneur takes the wage per efficiency unit w_t as given and chooses l_t to maximize profit. The first-order condition equates the marginal revenue product of an efficiency unit of labour to its price. After collecting the constant terms generated by the Cobb–Douglas structure, this condition can be written as

$$\alpha^{1+\alpha}(1-\alpha)^{1-\alpha}A_T l_T\,(w_t)^{\alpha-1} = w_t$$

Solving for the optimal labour demand of a type-T entrepreneur gives

$$l_T^*\,(w_t) = \left(\frac{\alpha^{1+\alpha}(1-\alpha)^{1-\alpha}A_T}{w_t}\right)^{\frac{1}{1-\alpha}}$$

Two features of this expression deserve comment. First, labour demand is decreasing in the wage, $\frac{\partial l_T^*}{\partial w_t} < 0$, as required for a well-behaved demand schedule. Second, and central to what follows, the only way in which the entrepreneur's type enters demand is through the productivity term A_T. A more productive (patient) entrepreneur demands strictly more labour at every wage.

Lemma 1 (Productivity ranking of entrepreneurs).

For any wage $w > 0$, a patient entrepreneur demands strictly more labour and earns strictly higher profit than an impatient entrepreneur[8]:

$$l_P(w) > l_I(w) \quad \textit{and} \quad \Pi_P(w) > \Pi_I(w)$$

We now aggregate across the entrepreneurial pool. Let the measure of active entrepreneurs in period t be E_t, and let μ_t^e denote the share of those entrepreneurs who are patient. Since each entrepreneur of a given type

[8] *Proof.* Substituting the optimal demand back into the technology and netting out the wage bill, the maximized profit of a type-T entrepreneur is

$$\Pi_T(w) = A_T{}^{\frac{1}{1-\alpha}}\,w^{-\frac{\alpha}{1-\alpha}}$$

Both $l_T(w) = \left(\frac{\alpha^{1+\alpha}(1-\alpha)^{1-\alpha}A_T}{w}\right)^{\frac{1}{1-\alpha}}$ and $\Pi_T(w)$ are strictly increasing in A_T because the exponent $\frac{1}{1-\alpha}$ is strictly positive for $\alpha \in (0,1)$. Since $A_P > A_I$, both inequalities follow immediately.

demands the type-specific quantity derived above, aggregate labour demand is the population-weighted sum of patient and impatient demands:

$$L_t^D\left(w \mid E_t, \mu_t^e\right) = E_t\left[\mu_t^e l_P(w) + \left(1-\mu_t^e\right) l_I(w)\right]$$

Substituting the type-specific demand functions and collecting the common wage term yields the closed-form aggregate schedule

$$L_t^D\left(w \mid E_t, \mu_t^e\right) = E_t w^{-\frac{1}{1-\alpha}}\left[\mu_t^e A_P{}^{\frac{1}{1-\alpha}} + \left(1-\mu_t^e\right) A_I{}^{\frac{1}{1-\alpha}}\right]$$

Lemma 2 (First macroeconomic feedback).

Holding the wage and the measure of entrepreneurs fixed, aggregate labour demand is strictly increasing in the patient share of entrepreneurs[9]:

$$\frac{\partial L_t^D}{\partial \mu_t^e} = E_t w^{-\frac{1}{1-\alpha}}\left(A_P{}^{\frac{1}{1-\alpha}} - A_I{}^{\frac{1}{1-\alpha}}\right) > 0$$

Lemma 2 records the first of two macroeconomic feedback results in the model and established an economic interpretation. A larger share of patient entrepreneurs raises aggregate labour demand, because patient entrepreneurs run more productive firms and therefore hire more efficiency units of labour at any given wage. This matters for the household problem analysed below: a higher patient share tightens the labour market and thereby raises the probability that a patient child from a worker household is actually employed in production work rather than rationed out of it. Composition on the entrepreneurial side spills over into the employment prospects of patient workers, linking the two sides of the social ladder.

III.C. Labour-Market Equilibrium with Rationing

As the previous subsection characterized labour demand, we now establish labour supply and describe how the market clears. The key conceptual point is that, because patient and impatient workers require different minimum wages to satisfy incentive compatibility, the market need not clear through the wage alone. When labour demand is insufficient, the higher reservation wage of

[9] *Proof.* Differentiating the closed-form schedule with respect to μ_t^e gives the displayed derivative. The prefactor $E_t w^{-\frac{1}{1-\alpha}}$ is positive, and the bracketed term is positive because $A_P > A_I$ and the exponent $\frac{1}{1-\alpha}$ is positive. Hence the derivative is strictly positive.

patient workers can leave them rationed even though they are individually more productive. The equilibrium is therefore characterized by a wage together with a rationing probability, and three regimes arise depending on the strength of demand.

Let μ_t^h denote the share of patient children among worker households. A patient worker supplies γ efficiency units and an impatient worker supplies one, so the effective labour supply of worker-household children, measured in efficiency units, is

$$L_t^S = (1 - E_t)\left[\left(1 - \mu_t^h\right) + \gamma\mu_t^h\right]$$

The term $1 - E_t$ is the measure of worker households; within it, a fraction $1 - \mu_t^h$ are impatient (each contributing one efficiency unit) and a fraction μ_t^h are patient (each contributing γ efficiency units). We now state the equilibrium regimes formally.

Proposition 1: Labour-market equilibrium and rationing of patient workers

Fix E_t, μ_t^e, μ_t^h. The labour market admits exactly one of three regimes, indexed by the strength of demand evaluated at the patient incentive-compatible wage w_P.

Case 1: Full employment

If demand is high enough that the market-clearing wage satisfies $w_t \geq w_P$, then all workers are employed and the clearing condition is

$$L_t^D\left(w_t \mid E_t, \mu_t^e\right) = (1 - E_t)\left[\left(1 - \mu_t^h\right) + \gamma\mu_t^h\right]$$

The patient employment probability is $q_t = 1$; patient workers earn γw_t and impatient workers earn w_t. Patience carries a strictly positive labour-market return.

Case 2: Full rationing of patient workers

If demand is so low that, even at w_P, firms do not demand enough labour to employ all impatient workers,

$$L_t^D\left(w_P \mid E_t, \mu_t^e\right) < \left(1 - \mu_t^h\right)(1 - E_t)$$

then the equilibrium wage falls in $w_t \in [w_I, w_P)$, the patient incentive constraint fails, and $q_t = 0$. Only impatient workers are employed and patience earns no return in production work.

Case 3: Intermediate demand and partial rationing

If demand at w_P is high enough to employ all impatient workers but not all patient workers,

$$\left(1 - \mu_t^h\right)(1 - E_t) \leq L_t^D\left(w_P \mid E_t, \mu_t^e\right) < (1 - E_t)\left[\left(1 - \mu_t^h\right) + \gamma\mu_t^h\right]$$

then the equilibrium wage equals $w_t = w_P$, impatient workers are hired first, and patient workers receive residual jobs with probability

$$q_t = \frac{L_t^D\left(w_P \mid E_t, \mu_t^e\right) - \left(1 - \mu_t^h\right)(1 - E_t)}{\gamma\mu_t^h\,(1 - E_t)}$$

with $q_t \in [0, 1)$.

The three regimes partition the value of demand at w_P relative to the two supply benchmarks $\left(1 - \mu_t^h\right)(1 - E_t)$ (impatient supply) and L_t^S (total supply), and these benchmarks are ordered because $\gamma > 1$. In Case 1, demand weakly exceeds total supply, so the wage rises until the market clears at $w_t \geq w_P$ with every worker employed, giving $q_t = 1$. In Case 2 demand at w_P is below impatient supply; the wage cannot stay at or above w_P because patient workers, who require w_P, cannot be profitably employed, so the wage drops into $[w_I, w_P)$, the patient incentive constraint is violated, and $q_t = 0$. In the intermediate case the wage is pinned at w_P: it cannot rise (excess supply remains because not all patient workers can be placed) yet it cannot fall (impatient demand alone would be unsatisfied, so firms bid the wage up to the patient constraint). The rationing probability then solves the residual market-clearing condition: total demand minus the impatient workers hired first, divided by the efficiency units of patient labour seeking work, $\gamma\mu_t^h\,(1 - E_t)$. The bounds on q_t follow directly from the two inequalities defining the regime.

The intermediate regime is the economically central case and organizes the household analysis that follows. The wage is fixed at w_P, so the expected wage income of a patient worker is $q_t\gamma w_P$ while an impatient worker earns

w_P with certainty, being hired first. Absent any route into entrepreneurship, patience is rewarded in production work if and only if expected patient income exceeds certain impatient income.

Lemma 3

In the intermediate regime, production work rewards patience if and only if the patient employment probability exceeds the inverse productivity advantage:

$$q_t \gamma w_P > w_P \quad \Longleftrightarrow \quad q_t > \frac{1}{\gamma}$$

The interpretation is a trade-off between a productivity premium and rationing risk. A patient worker who is employed produces γ times as much, but is employed only with probability q_t. The productivity advantage dominates the rationing risk precisely when $q_t > \frac{1}{\gamma}$. When patient workers are unlikely to be placed, worker-household parents have little reason, through the production-work channel alone, to incur the cost of cultivating patience.

III.D. Social Mobility and Entry into Entrepreneurship

Social mobility enters the model as the second channel through which patience can be rewarded. So far patience has paid off only through production work, and Lemma 3 shows that this channel can be weak when patient workers are rationed. The mobility channel opens an alternative: a patient child from a worker household may bypass production work entirely and become an entrepreneur. Whether this route is viable depends on how costly it is to enter, which is precisely where social mobility operates.

Children of entrepreneur households inherit their parents' firm and enter entrepreneurship at no cost. Children of worker households, by contrast, must pay an entry cost to acquire the same opportunity. We model this barrier as proportional to the wage and decreasing in the mobility environment M_t:

$$B_t = b(M_t) w_t, \quad with \quad b^{'}(M_t) < 0$$

Writing the barrier as a fraction $b(M_t)$ of the wage keeps the entry decision scale-free in the wage, so that comparisons between the production-work payoff and the entrepreneurial payoff do not depend mechanically on the

wage level. A higher M_t — greater social mobility — lowers the effective barrier $b(M_t)$, making the entrepreneurial route more accessible to children of worker households.

A patient child from a worker household compares two payoffs. Remaining in production work yields expected income $q_t \gamma w_t$; entering entrepreneurship yields net income $\Pi_P w_t - b(M_t) w_t$. Entry is the preferred option when

$$\Pi_P w_t - b(M_t) w_t \geq q_t \gamma w_t$$

Consider the benchmark in which the patient worker is fully employed ($q_t = 1$) and the wage sits at the patient incentive-compatible level ($w_t = w_P$). The entry condition becomes

$$\Pi_P w_P - b(M_t) w_P \geq \gamma w_P$$

Dividing by the wage isolates the mobility barrier and defines the first mobility threshold.

Lemma 4

Define M^* by $b(M)^* = \frac{\Pi_P w_P}{w_P} - \gamma$. Then, at the full-employment benchmark, a patient child of a worker household prefers entrepreneurship to production work if and only if $M_t > M^*$.

Proof.

Divide the entry condition $\Pi_P w_P - b(M_{t)} w_P \geq \gamma w_P$ by $w_P > 0$ to obtain $b(M_t) \leq \frac{\Pi_P w_P}{w_P} - \gamma$, i.e. $b(M_t) \leq b(M^*)$. Because $b^{'}(M_t) < 0$, the map bM_t is strictly decreasing, so $b(M_t) \leq b(M^*)$ holds exactly when $M_t \geq M^*$, with strict preference for $M_t > M^*$.

When patient children of worker households are exactly indifferent between production work and entry, the wage adjusts to maintain the indifference condition

$$\gamma w^* M_t = \Pi_P w^* M_t - b(M_t) w^* M_t$$

Using the closed-form profit function and solving for the wage gives

$$w^* M_t = \left[\frac{A_P^{\frac{1}{1-\alpha}}}{bM_t + \gamma} \right]^{1-\alpha}$$

This equilibrium wage is decreasing in the barrier $b(M_t)$ and hence increasing in mobility M_t: easier entry raises the outside option of patient workers and bids up the wage at which they are indifferent between working and entering.

A second threshold governs whether inherited impatient entrepreneurs choose to remain in entrepreneurship. An impatient entrepreneur stays only if the firm's return covers the opportunity cost of the labour market; the marginal case is reached when the indifference wage $w^* M_t$ rises to the impatient productivity term A_I.

Lemma 5

Define M^{**} by $w^* M^{**} = A_I$, equivalently $b(M^{**}) = \left(\frac{A_P}{A_I}\right)^{\frac{1}{1-\alpha}} - \gamma$. If $M_t > M^{**}$, inherited impatient entrepreneurs exit and the entrepreneurial sector becomes increasingly dominated by patient entrepreneurs.

Proof.

Set $w^* M_t = A_I$ in the closed-form wage and solve for the barrier. Raising both sides to the power $\frac{1}{1-\alpha}$ on the right and inverting gives $b(M^{**}) = \left(\frac{A_P}{A_I}\right)^{\frac{1}{1-\alpha}} - \gamma$. Because $b(M_t)$ is strictly decreasing in M_t, the condition $b(M_t) < b(M^{**})$ — under which the indifference wage exceeds A_I, so an impatient entrepreneur does better in the labour market than by operating the firm — is equivalent to $M_t > M^{**}$. Exit of impatient entrepreneurs raises the patient share μ_t^e, which by Lemma 2 further raises labour demand.

The two thresholds M^* and M^{**} were derived at the full-employment benchmark $q_t = 1$. For intermediate mobility, patient workers are only partially employed, so the relevant comparison must account for the employment probability q_t. We therefore define a q-dependent barrier threshold that nests the previous case.

Lemma 6

Define

$$\tilde{b_{q_t}} = \frac{\Pi_P w_P}{w_P} - q_t \gamma \quad = \quad A_P{}^{\frac{1}{1-\alpha}} w_P{}^{-\frac{1}{1-\alpha}} - q_t \gamma$$

A patient child from a worker household has a credible entrepreneurial route if $b(M_t) < \tilde{b}(q_t)$; if instead $b(M_t) > \tilde{b}(q_t)$, patience is rewarded

neither through production work nor through entry. The threshold nests Lemma 4: at $q_t = 1$, $\tilde{b(1)} = b(M^*)$.[10]

We can now collect these results into a single object: the total labour-market return to patience for a worker household, which is the payoff under the best available option (production work or entry) net of the impatient benchmark income.

Let $Y_I^W w_t$ denote the income an impatient worker-household child obtains. The total market return to patience is

$$R_t^W(M_t, q_t, w_t) = \max\{q_t \gamma w_t, \Pi_P w_t - b(M_t) w_t\} - Y_I^W w_t$$

In the intermediate regime the wage is $w_t = w_P$ and the impatient benchmark is $Y_I^W w_P = w_P$, so the return specializes to

$$R_t^W(M_t, q_t, w_P) = \max\left\{q_t \gamma w_P, \Pi_P w_P - b(M_t) w_P\right\} - w_P$$

This expression is the bridge between the labour-market equilibrium and the household parenting mechanism of the next section. The return to patience is governed by two margins. The production-work margin, $q_t \gamma w_P$, is high when patient workers are likely to be employed (Lemma 3); the entrepreneurial margin, $\Pi_P w_P - b(M_t) w_P$, is high when mobility is strong enough that the barrier $b(M_t)$ falls below $\tilde{bq_t}$ (Lemma 6). Both margins are increasing in social mobility — directly through the lower entry barrier, and indirectly through the composition feedback of Lemma 2, whereby a higher patient share of entrepreneurs raises labour demand and hence q_t. The household problem in the next section takes $R_t^W(M_t, q_t, w_P)$ as the economic payoff that parents weigh against the cognitive-effort cost of cultivating patience.

Definition (Lower mobility threshold). Define the lower mobility threshold $\bar{M}$ as the mobility level at which the total market return to patience just vanishes,

[10] *Proof.* With partial employment the entry condition $\Pi_P w_P - b(M_t) w_P \geq q_t \gamma w_P$ replaces γ by $q_t \gamma$ on the right-hand side. Dividing by w_P yields $b(M_t) \leq \frac{\Pi_P w_P}{w_P} - q_t \gamma = \tilde{bq_t}$. Using $\Pi_P w_P = A_P^{\frac{1}{1-\alpha}} w_P^{-\frac{\alpha}{1-\alpha}}$ gives the second expression. Setting $q_t = 1$ recovers $b(M^*)$ from Lemma 4.

$$b(\bar{M}) = \widetilde{b}\,(q_t)\,, \quad equivalently \quad R_t^W\left(\bar{M}, q_t, w_P\right) = 0$$

For $M_t < \bar{M}$ the barrier is so high that neither the production-work margin nor the entrepreneurial margin makes patience worthwhile, so worker-household parents instill patience through neither channel: this is the informational/cognitive-effort trap, region (i) of the figure. For $M_t > \bar{M}$ a credible entrepreneurial route exists and the return to patience becomes positive. Because $b^{'}(M_t) < 0$ and $\widetilde{b}\,(q_t) > b\,(M^*)$ whenever patient workers are rationed ($q_t < 1$), the three thresholds are ordered $\bar{M} < M^* < M^{**}$, matching the regions of the figure: the return to patience first turns positive at $\bar{M}$, entry overtakes fully-employed production work at M^*, and inherited impatient entrepreneurs exit at M^{**}.

III.E. Heterogeneity by Household Type

We now turn to our household settings. We assume the economy is in discrete time and is populated by a continuum 1 of non-overlapping generations. Each individual lives for one period and begets a single offspring. In line with Cunha and Heckman (2007) model that parents directly shape child preferences, the main interaction within the family is a parent's choice of time preferences for their offspring. Instantaneous (stage) utilities are defined over expected net income. In particular, an individual has overall utility

$$U_t^h = \lambda\left(u_t^h - B_t^h\right) + (1-\lambda)\left(u_{t+1}^h - B_{t+1}^h\right)$$

where u_t^h denotes this individual's expected stage utility, B_t^h is a non-pecuniary cost arising from barriers to upward mobility faced by individuals from low-SES backgrounds entering high-SES domains and u_{t+1}^h *and* B_{t+1}^hdenote her offspring's expected income and entry cost, and $\lambda \in (0, 1)$. This specification captures the extent to which parents care about their offspring's net income and social mobility, which is linked to parental investment and style, but not about their continuation utility. In addition, we focus on the case where $\lambda \to 1$, so that the individual cares much more about herself net income than her offspring. This will ensure that current occupation decisions are not made in order to reduce the next generation's entry costs.

For now, we specify that each parent h (of generation t) decides their "parenting strategy" -patience or impatience - corresponding to their decision about whether to impart patience or impatience values to their offspring. These parenting strategy decisions are denoted by $T'^{h}_{t} \in \{P, \quad I\}$. If a parent of type T^{h}_{t} chooses parenting strategy T′, then the type of the offspring will be $T^{h}_{t+1} = P$ with probability $\pi_{TT'}$, and $T^{h}_{t+1} = I$ with probability1-$\pi_{TT'}$.We assume that the transmission principle (i.e. $\pi_{TT'}$) satisfies the following two natural conditions:

$1 > \pi_{P\cdot} > \pi_{I\cdot} > 0$ and $1 > \pi_{\cdot P} > \pi_{\cdot I} > 0$,

implying that all else equal, parents who are themselves patient are more likely to have patient offspring (Brown and van der Pol, 2015), parents who are choosing to impart patience values are more likely to have patient offspring, and all these probabilities are between 0 and 1.

A natural concern is whether impatient parents can plausibly transmit patience to their children. Our argument does not imply that impatient parents become patient themselves. Instead, it implies that the parenting choice may change when the expected return to patience changes. This logic is consistent with preference reversals, whereby observed choices need not reflect a fixed preference ordering but may shift when the decision environment changes how alternatives are evaluated (Tversky and Thaler, 1990). Both theoretical and empirical work support this possibility. Theoretical, expectations-based reference dependence provides a foundation for preference reversals because reference points are formed endogenously from expectations. When expectations change, the same option can generate different gain-loss utility relative to different reference points, so it may be evaluated differently across environments. This can lead to violations of standard choice axioms such as Independence and, in some cases, the Weak Axiom of Revealed Preference (Freeman, 2019). In our setting, upward mobility changes the reference point against which parents evaluate the value of cultivating patience. When upward expectations are weak, impatient parents may rationally allow impatience to persist. When upward expectations are strong, however, not cultivating patience may be perceived as a loss relative to the child's attainable future position. This can generate a preference reversal in parenting choices, as even impatient parents may exert costly cognitive effort cultivating patience. Empirically, Brenøe and Epper (2022) show that

the authoritative parenting style can attenuate the transmission of impatience rather than merely reproducing parents' own time preferences. Therefore, in our model, parental type T_t^h and parenting strategy $T_t'^h$ need not coincide. An impatient parent with $T_t^h = I$ may optimally choose $T_t'^h = P$ when expected upward mobility is high. We discuss in detail the conditions under which this type of preference reversal can occur in next section.

III.F. Stage 1: Expectations-Based Reference Points and the Parenting Choice Set

Parents do not observe the true return to patience. The labour-market equilibrium derived above delivers that return as an object, $R_t^W(M_t, q_t, w_P)$, yet each parent must act on a forecast of it rather than on its realized value. Let parent h in a household of background $j \in \{W, E\}$ form an expectation of the return to cultivating patience, conditional on an information set I_{ht}:

$$m_{ht}^j = E(R_t^j \mid I_{ht})$$

For a worker household this expectation is taken over the worker-market return,

$$m_{ht}^W = E(R_t^W(M_t, q_t, w_P) \mid I_{ht})$$

The information set I_{ht} collects the signals from which parents infer mobility prospects: class background, neighbourhood experience, local mobility, school information, observed role models, and realized labour-market outcomes. Two comparative statics inherited from the labour-market block discipline this expectation. Because higher mobility lowers the entry barrier and raises the probability that patience is rewarded, the expected return rises with mobility,

$$\frac{\partial m_{ht}^W}{\partial M_t} > 0$$

and because a larger patient share among entrepreneurs raises labour demand and hence the employment probability of patient workers (Lemma 2), it rises with the composition of the entrepreneurial pool,

$$\frac{\partial m_{ht}^{W}}{\partial \mu_{t}^{e}} > 0$$

Expectations thus transmit the macroeconomic state of the economy into the household. The parent then chooses between two parenting plans, $D_h = \{d_h^0, d_h^P\}$, where d_h^0 denotes default parenting, which preserves the household's baseline routine, and d_h^P denotes active patience-instilling parenting. Active parenting requires the parent to incur a cognitive and psychological effort cost: monitoring the child's educational trajectory, encouraging delayed gratification, enforcing future-oriented discipline, interpreting opportunities, and sustaining a long-horizon strategy. Let the probability that the child becomes patient be π_h^0 under default and π_h^P under active parenting, with $\pi_h^P > \pi_h^0$, and write the effectiveness of active parenting as $\Delta\pi_h = \pi_h^P - \pi_h^0$. The effort cost is

$$K_h = \frac{\kappa_0}{a_h\,(1+\rho_h)}$$

where a_h is parental cognitive bandwidth and ρ_h is parental expertise in interpreting education-, labour-market-, and mobility-relevant opportunities; greater bandwidth or expertise lowers the cost of cultivating patience. Valuing each plan by the expected child return net of effort gives $V_h(d_h^0 \mid m_{ht}^j) = \pi_h^0 m_{ht}^j$ and $V_h(d_h^P \mid m_{ht}^j) = \pi_h^P m_{ht}^j - K_h$, so the material gain from active parenting is

$$\Delta V_h^j = \Delta\pi_h m_{ht}^j - K_h$$

Under standard preferences a parent would cultivate patience whenever $\Delta\pi_h m_{ht}^j \geq K_h$. The model departs from this benchmark by making the parent's expected plan a reference point against which the realized plan is evaluated, which is the subject of the next subsection.

Personal Equilibrium and the Preferred Personal Equilibrium

Following the logic of expectations-based reference dependence, a parenting plan is sustainable only if it is optimal given that the parent expects herself to carry it out. A plan $d_h \in D_h$ is a personal equilibrium if it is a best response to itself,

$$U_h(d_h \mid d_h) \geq U_h(d_h^{'} \mid d_h) \;\; for\ all \;\; d_h^{'} \in D_h$$

and the preferred personal equilibrium (PPE) is the personal equilibrium with the highest ex ante value,

$$PPE(D_h) = \underset{d_h \in PED_h}{arg\ max}\, U_h(d_h \mid d_h)$$

Equivalently, the parent first identifies which plans are self-confirming, then selects the self-confirming plan she would most prefer to commit to ex ante. Loss aversion disciplines the comparison: if she expects default parenting, switching to active parenting is felt as a loss of parental effort; if she expects active parenting, lapsing into the default is felt as a loss of the child's mobility opportunity. Let $\lambda > 0$ measure the strength of loss aversion. With the default as reference point, the gain from switching to active parenting is

$$U_h(d_h^P \mid d_h^0) - U_h(d_h^0 \mid d_h^0) = \Delta\pi_h m_{ht}^j - (1+\lambda)K_h$$

so, default parenting is a personal equilibrium when $\Delta\pi_h m_{ht}^j \leq (1+\lambda)K_h$. With active parenting as reference point, the gain in its favour is

$$U_h(d_h^P \mid d_h^P) - U_h(d_h^0 \mid d_h^P) = (1+\lambda)\Delta\pi_h m_{ht}^j - K_h$$

so active parenting is a personal equilibrium when $(1+\lambda)\Delta\pi_h m_{ht}^j \geq K_h$. Ranking each plan against itself, active parenting is the preferred personal equilibrium whenever $\Delta\pi_h m_{ht}^j \geq K_h$, that is, once the expected return clears

$$m_{ht}^j \geq \bar{m}_h^{PPE}, \;\; \bar{m}_h^{PPE} = \frac{K_h}{\Delta\pi_h}$$

The threshold rises with the cost of active parenting and falls with its effectiveness, $\frac{\partial \bar{m}_h^{PPE}}{\partial K_h} > 0$ and $\frac{\partial \bar{m}_h^{PPE}}{\partial \Delta\pi_h} < 0$; and since $K_h = \frac{\kappa_0}{a_h(1+\rho_h)}$, it falls in both bandwidth and expertise. Parents with more bandwidth or better knowledge of mobility-relevant opportunities require a lower expected return before active parenting becomes self-confirming.

III.G. Stage 2: Four Parenting Regimes

The PPE threshold partitions expected returns into four regimes that mirror the four regions of Figure I. Loss aversion drives a wedge λ between the expectation at which active parenting first becomes self-confirming and the higher level at which it becomes preferred, so the interior of the figure contains a band in which patience is feasible but not yet chosen.

Region I (Low-expectation trap).

If $m_{ht}^{j} < \frac{K_h}{(1+\lambda)\Delta\pi_h}$, active parenting is not even a personal equilibrium; the only self-confirming plan is the default, $PE(D_h) = \left\{d_h^0\right\}$, so $PPE(D_h) = d_h^0$. Expecting to remain in the default routine, the parent finds active cultivation too costly to begin. This is region (i) of Figure I, where expectations are so low that patience is never contemplated.

Proposition 2 Low-mobility PPE trap

If social mobility is low, so that labour demand is weak and patient worker-household children face high rationing risk, then default parenting is the preferred personal equilibrium for worker households: when $m_{ht}^{W} < \bar{m}_h^{PPE}$, $PPE(D_W) = d_W^0$. In the intermediate labour-demand regime this is equivalent to $q_t < \frac{1}{\gamma} + \frac{K_h}{\gamma w_P \Delta\pi_h}$. The patient share in worker households converges to $\bar{\mu}_I^h$ and the economy settles at the low-patience steady state $\left(\mu^h, \mu^e\right) = \left(\bar{\mu}_I^h, \bar{\mu}^e\right)$.

The intuition is that low mobility validates low expectations. When patience is unlikely to be rewarded through either production work or entry, active parenting is too costly to be self-confirming, parents remain in the default, and impatience persists across generations. This is the formal counterpart of regions (i)-(ii) in Figure V.

Region II (Cognitive-effort trap).

If $\frac{K_h}{(1+\lambda)\Delta\pi_h} \leq m_{ht}^{j} < \frac{K_h}{\Delta\pi_h}$, active parenting is a personal equilibrium but not the preferred one. The parent recognises that cultivating patience could be self-confirming, yet its ex-ante value still falls short of the default, so she does not switch. The obstacle is not missing information but an expected return too low to justify the cognitive and psychological cost of changing behaviour. This is region (ii) of Figure V, and it lies below the cognitive-effort threshold $\bar{R}$ drawn there.

Region III (Expectation-driven patience).

If $\frac{K_h}{\Delta\pi_h} \leq m_{ht}^j < \frac{(1+\lambda)K_h}{\Delta\pi_h}$, active parenting is the preferred personal equilibrium, $PPED_h = d_h^P$. The parent now expects to cultivate patience, and that expectation is self-confirming: failing to act would beequivalent to forgoing a valuable mobility opportunity. This is region (iii), where instilling patience begins.

Proposition 3: High-mobility PPE activation

If social mobility is high enough to raise the expected return above the PPE threshold, worker households cultivate patience: when $m_{ht}^W \geq \bar{m}_h^{PPE}$, $PPE(D_W) = d_W^P$. In the intermediate regime this is equivalent to $q_t \geq \frac{1}{\gamma} + \frac{K_h}{\gamma w_P \Delta\pi_h}$, and the patient share in worker households moves toward $\bar{\mu}_P^h$. High mobility changes the parent's self-confirming expectation: she no longer expects to remain in the default, but to act, and the expectation sustains itself because not acting would mean losing the opportunity.

Region IV (Preference-reversal region).

Impatient parents have a weaker default tendency to cultivate patience and may face a higher effort cost, so $K_I > K_P$ or $\Delta\pi_I < \Delta\pi_P$, and their threshold is higher, $\bar{m}_I^{PPE} > \bar{m}_P^{PPE}$. The steeper slope in the high-mobility region captures an over-optimistic expectation-amplification effect. Motivated by Alesina, Stantcheva, and Teso (2018), once upward mobility becomes sufficiently credible, parents may treat it not merely as a probability but as an attainable reference point for the child's future. At this point, expectations can become aspirational and may rise faster than calibrated mobility signals.

Proposition 4: Preference reversal among impatient parents

Let $\bar{m}_I^{PPE} = \frac{K_I}{\Delta\pi_I}$. If mobility is high enough that $m_{It}^W \geq \bar{m}_I^{PPE}$, then $PPE_I(D_I) = d_I^P$. Even impatient parents actively instill patience. Their own type remains impatient, but their parenting plan changes because the expected return is high enough to make active cultivation the preferred personal equilibrium; impatience is no longer mechanically reproduced in the next generation. This is region (iv) of Figure V, where even impatient parents instil patience.

III.H. Labour-Market Equilibrium and the Worker-Parent Switching Condition

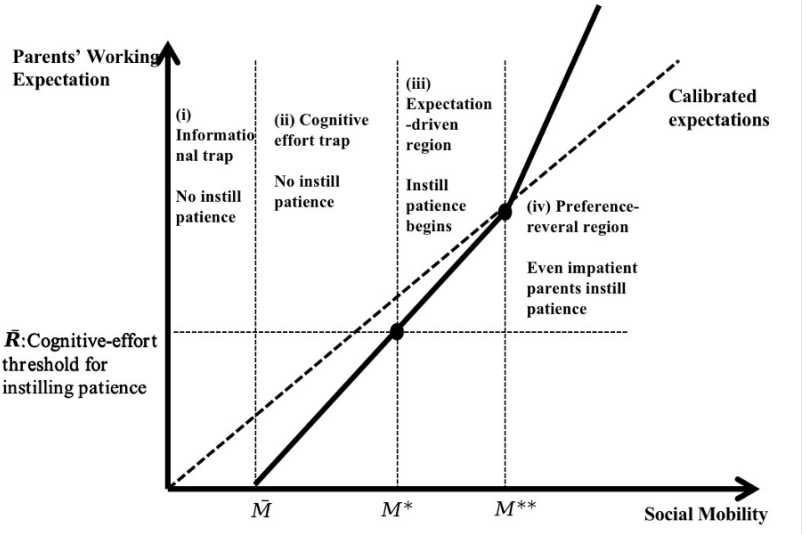


Figure V
Four Parenting Regimes

The PPE threshold can now be embedded in the labour-market equilibrium. Worker parents cultivate patience when $m_{ht}^W \geq \bar{m}_h^{PPE}$, which using the market return reads $E[R_t^W(M_t, q_t, w_P) \mid I_{ht}] \geq \bar{m}_h^{PPE}$. In the intermediate labour-demand regime, where $w_t = w_P$ and patience is rewarded only through production work, the market return is

$$R_t^{W,L} = q_t \gamma w_P - w_P$$

so that, under unbiased expectations, the worker-parent switching condition becomes $q_t \gamma w_P - w_P \geq \bar{m}_h^{PPE}$, or equivalently

$$q_t \geq \frac{1}{\gamma} + \frac{\bar{m}_h^{PPE}}{\gamma w_P}$$

Substituting the threshold $\bar{m}_h^{PPE} = \frac{K_h}{\Delta \pi_h}$ gives

$$q_t \geq \frac{1}{\gamma} + \frac{K_h}{\gamma w_P \Delta \pi_h}$$

This is the central modification of the worker-parenting threshold in Acemoglu's framework. There, non-entrepreneur parents switch to the high-

return trait once $q_t > \frac{1}{\gamma}$. Here, because cultivating patience requires costly cognitive effort, the required employment probability is strictly higher by the term $\frac{K_h}{\gamma w_P \Delta \pi_h}$: the more costly or less effective is active parenting, the tighter the labour market must be before worker parents invest in patience.

Because the employment probability is increasing in patient entrepreneurship, $\frac{\partial q_t}{\partial \mu_t^e} > 0$, the switching condition can be expressed as a threshold share of patient entrepreneurs $\widetilde{\mu}_h^e$: worker parents choose active parenting if $\mu_t^e \geq \widetilde{\mu}_h^e$ and remain in the default if $\mu_t^e < \widetilde{\mu}_h^e$. The threshold is defined implicitly by

$$q\widetilde{\mu}_h^e = \frac{1}{\gamma} + \frac{K_h}{\gamma w_P \Delta \pi_h}$$

which, using the rationing equation, becomes

$$\frac{L_t^D w_P \mid E_t, \widetilde{\mu}_h^e - \left(1 - \mu_t^h\right)(1 - E_t)}{\gamma \mu_t^h (1 - E_t)} = \frac{1}{\gamma} + \frac{K_h}{\gamma w_P \Delta \pi_h}$$

The share of patient entrepreneurs thus determines whether worker parents cross the PPE threshold for active patience cultivation, which is the channel through which the two sides of the economy interact dynamically.

III.I. Parenting Dynamics and the Two Mobility Regimes

Entrepreneur households have direct access to entrepreneurship, and because patient entrepreneurs are more productive (Lemma 1) they always cultivate patience, $PPE_E(D_E) = d_E^P$. The patient share among entrepreneurs therefore evolves as

$$\mu_{t+1}^e = \pi_{PP}^e \mu_t^e + \pi_{IP}^e \left(1 - \mu_t^e\right)$$

where π_{PP}^e and π_{IP}^e are the probabilities that a patient and an impatient entrepreneur parent, respectively, raise a patient child under active parenting. The entrepreneur-side steady state is

$$\bar{\mu}^e = \frac{\pi_{IP}^e}{1 - \pi_{PP}^e + \pi_{IP}^e}$$

Worker households behave differently, because their parenting choice depends on the labour-market return and hence on μ_t^e, q_t, and the mobility environment. When $\mu_t^e \geq \widetilde{\mu}_h^e$ they cultivate patience, $PPE_W D_W = d_W^P$, and

the patient share follows $\mu^h_{t+1} = \pi^h_{PP}\mu^h_t + \pi^h_{IP}\left(1-\mu^h_t\right)$ with steady state $\bar{\mu}^h_P = \frac{\pi^h_{IP}}{1-\pi^h_{PP}+\pi^h_{IP}}$. When $\mu^e_t < \widetilde{\mu}^e_h$ they remain in the default, $PPE_W(D_W) = d^0_W$, and the share follows $\mu^h_{t+1} = \pi^h_{PI}\mu^h_t + \pi^h_{II}\left(1-\mu^h_t\right)$ with steady state $\bar{\mu}^h_I = \frac{\pi^h_{II}}{1-\pi^h_{PI}+\pi^h_{II}}$, and $\bar{\mu}^h_P > \bar{\mu}^h_I$. The horizontal worker dynamics therefore have two candidate rest points, the low-patience $\bar{\mu}^h_I$ and the high-patience $\bar{\mu}^h_P$, and which one obtains is selected by whether the vertical entrepreneur dynamics push the patient share above or below the switching threshold $\widetilde{\mu}^e_h$. Figures 6 displays the two resulting phase portraits in the $\left(\mu^h, \mu^e\right)$ plane, with the patient share of workers on the horizontal axis and of entrepreneurs on the vertical axis.

The low-mobility regime.

When mobility is low, barriers are high and labour demand is weak, so the entrepreneur-side steady state lies below the worker-parenting threshold, $\bar{\mu}^e < \widetilde{\mu}^e_h$. Worker households then stay in the default and the system converges to $\left(\mu^h, \mu^e\right) = \left(\bar{\mu}^h_I, \bar{\mu}^e\right)$, shown in Figure VI (left-hand).

The horizontal arrows point toward $\bar{\mu}^h_I$, because worker households do not actively cultivate patience and their dynamics are governed by the impatient transition probabilities. The vertical arrows point toward $\bar{\mu}^e$, because entrepreneur households keep cultivating patience; but since the entrepreneur-side share settles below $\widetilde{\mu}^e_h$, indicated by the vertical dotted line at $\bar{\mu}^h_{Impatience}$, labour demand never tightens enough to flip worker parenting. The economy is absorbed into the low-patience corner, the dynamic analogue of Proposition 2.

The high-mobility regime.

When mobility is high, barriers are low and patient worker-household children have credible routes into productive employment or entrepreneurship. The long-run entrepreneur-side share now rises above the worker-parenting threshold, $\widetilde{\mu}^R > \widetilde{\mu}^e_h$. Worker households switch to active parenting, and the system converges to $\left(\mu^h, \mu^e\right) = \left(\bar{\mu}^h_P, \widetilde{\mu}^R\right)$, shown in Figure VI (right-hand).

The horizontal arrows now point toward $\bar{\mu}^h_P$, because worker households instil patience and follow the patient transition probabilities, crossing the vertical dotted line at $\bar{\mu}^h_{Patience}$. The vertical arrows point toward $\widetilde{\mu}^R$, because low barriers allow patient worker-household children to enter entrepreneurship,

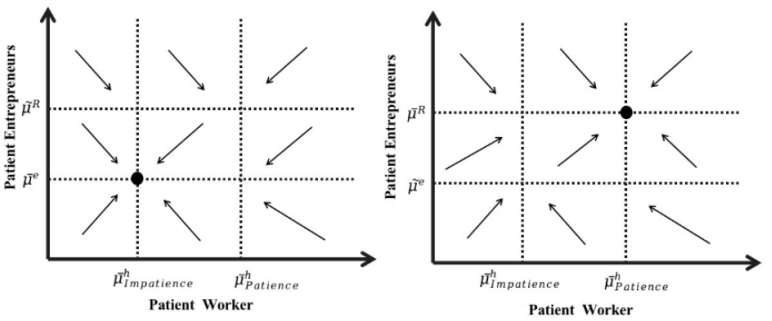


Figure VI
Dynamics of Fractions of Patient Entrepreneurs and Workers

raising the long-run patient entrepreneur share above $\bar{\mu}^e$. The economy is absorbed into the high-patience corner, the dynamic analogue of Propositions 2 and 3.

Proposition 5

The economy exhibits complementarity between patient entrepreneurship and worker-household patience cultivation. Since $\frac{\partial L_t^D}{\partial \mu_t^e} > 0$ and $\frac{\partial q_t}{\partial L_t^D} > 0$, we have $\frac{\partial q_t}{\partial \mu_t^e} > 0$; and since worker-household cultivation is increasing in the employment probability, $\frac{\partial \mu_{t+1}^h}{\partial q_t} > 0$, it follows that $\frac{\partial \mu_{t+1}^h}{\partial \mu_t^e} > 0$. A higher patient entrepreneur share raises labour demand, the employment probability of patient workers, and the expected return to patience, making active parenting more likely to be the PPE for worker households; in turn a higher patient share among workers enlarges the future pool of patient entrepreneurs, especially when barriers are low.

This two-way feedback supports the two steady-state regimes traced out above: the low-mobility regime $(\mu^h, \mu^e) = (\bar{\mu}_I^h, \bar{\mu}^e)$ of Figure VI (left) and the high-mobility regime $(\mu^h, \mu^e) = (\bar{\mu}_P^h, \bar{\mu}^R)$ of Figure VI (right).

III.J. Testable Implications

The model maps directly onto the cross-country patterns documented in Section II and sharpens them into predictions that can be taken to micro data.

Five implications follow from the comparative statics above; each pairs a mechanism with an observable counterpart already used in Figures 1 and 2.

Implication 1 : Mobility raises cultivated patience.

Because the expected return is increasing in mobility, $\frac{\partial m_{ht}^W}{\partial M_t} > 0$, and the PPE switch is crossed only once $m_{ht}^W \geq \bar{m}_h^{PPE}$, average patience should rise with social mobility across societies. This is the relationship in Panel A of Figure I, between the Global Preferences Survey patience measure and the Global Social Mobility Index, and the model predicts it is causal through the parenting channel rather than an artefact of income. The prediction is specifically about cultivated patience: the steepest gradient should appear where the return to patience operates through formal labour-market advancement.

Implication 2 : The gradient is conditional on occupational structure

The mobility channel works through the worker-market return $R_t^W(M_t, q_t, w_P)$, which is rewarded when advancement runs through education, training, and career progression. Where self-employment dominates, patience is rewarded through delayed returns regardless of mobility, so $\frac{\partial m_{ht}^W}{\partial M_t}$ is attenuated. The model therefore predicts a steeper mobility-patience slope in low-self-employment economies than in high-self-employment ones, the interaction shown in Panel A of Figure II. The same logic applies to the relational traits in Panel B: mobility should raise the value parents place on tolerance and respect only where wage employment, not self-employment, governs returns.

Implication 3 : Wealth is not the driver.

The switch depends on the expected return relative to the effort threshold $\bar{m}_h^{PPE} = \frac{K_h}{\Delta \pi_h}$, not on income levels. Holding the return channel fixed, variation in resources alone should not move cultivated patience. The model thus predicts no systematic patience gradient in GDP per capita once occupational structure is held constant, the placebo in Panel C of Figure II; a strong wealth effect would be evidence against the mechanism.

Implication 4 : Cognitive bandwidth and expertise shift the threshold.

Since $\frac{\partial \bar{m}_h^{PPE}}{\partial a_h} < 0$ and $\frac{\partial \bar{m}_h^{PPE}}{\partial \rho_h} < 0$, parents with greater cognitive bandwidth or better knowledge of mobility-relevant opportunities should cross into active parenting at lower expected returns. Empirically, at a given level of

local mobility the propensity to invest in children's patience should be higher among parents facing fewer competing demands on attention and among those with more information about education and labour-market pathways. This is a within-country, micro-level prediction that does not follow from a pure income or pure labour-market account.

Implication 5 : Preference reversal and a mobility threshold in transmission.

Because the impatient-parent threshold $\bar{m}_I^{PPE} = \frac{K_I}{\Delta \pi_I}$ exceeds the patient-parent threshold, the model predicts a threshold rather than smooth relationship between mobility and intergenerational transmission of impatience: once expectations cross $\bar{m}_I^{PPE}$, even impatient parents instill patience and the mechanical reproduction of impatience breaks. Intergenerational persistence of time preference should therefore weaken sharply above a mobility cutoff, and the decline should be concentrated among initially impatient households. This distinguishes the PPE mechanism from models in which transmission declines smoothly with opportunity.

Taken together, the five implications make the framework falsifiable along several independent margins: the sign and conditionality of the mobility-patience gradient (Implications 1-2), the absence of a wealth effect (Implication 3), the micro-level role of bandwidth and information (Implication 4), and the nonlinearity in transmission (Implication 5). The cross-country evidence in Section II is consistent with the first three; the latter two point to household-level tests using panel data on parenting investment, local mobility, and children's realized time preferences.

IV. Extension: Technology Shocks and Patience

The baseline model holds techology fixed. Productivity is given conditional on entrepreneurial type, so patience formation affects aggregate output only through the composition of the entrepreneurial pool, and anything that raises labour demand raises the employment probability q_t and hence the return to patience. We now relax that restriction. Motivated by Doepke and Zilibotti (2008), who study how the Industrial Revolution changed the economic returns to class-specific preferences, and by the task-based approach to automation developed by Acemoglu and Restrepo (2018, 2019) and Acemoglu and

Loebbing (2026), we ask how a technology shock that simultaneously raises firm productivity and displaces labour from automatable tasks propagates through the two margins the model has isolated: the employment probability of Proposition 1 and the accessible mobility ladder of Section III.D.

This extension matters for the argument of the paper because it separates two objects that the baseline treats as one. In the baseline, the scale of employment and the accessibility of advancement move together. Automation breaks that link. The productivity effect governs the scale of labour demand, whereas the displacement effect also governs the composition of the tasks that remain assigned to workers, and interior automation of the kind documented in the polarisation literature removes precisely the intermediate-complexity occupations that have historically served as rungs on the ladder (Autor, Levy, and Murnane, 2003; Autor and Dorn, 2013; Goos, Manning, and Salomons, 2014). It is therefore possible for aggregate employment to expand while the ladder along which children of worker households used to climb contracts. The framework developed above is well suited to making this distinction precise, because the return to patience already runs through two separate branches — production work and entry into entrepreneurship — and a task shock reaches them by different routes.

Three features of the exercise should be stated at the outset. First, it is a partial-equilibrium extension of the baseline model: when differentiating with respect to the automation state we hold fixed the efficiency wage, the number of firms, and the population shares μ_t^h and μ_t^e, so that long-run feedback through wages, entry, and the endogenous composition of households remains governed by the equilibrium mechanism of Section III.J. Second, we use a reduced-form task structure inspired by the assignment approach of Acemoglu and Loebbing (2026) rather than reproducing their full assignment model; this buys tractability at the cost of not solving for the equilibrium task threshold, and we do not separately solve for the optimal capital stock or the rental price of machines. Third, the current automation state is assumed to be persistent enough that parents expect it to affect the labour market their children will enter, which is what makes it relevant to the parenting decision at all. To keep the extension nested within the baseline, firm types continue to be indexed by the entrepreneur's time preference, and we write $A_H \equiv A_P$ and $A_L \equiv A_I$

where the high- and low-productivity labels aid readability; setting $\tau \equiv 1$ and $A_T(z_t) \equiv A_T$ recovers every expression of Section III.B exactly.

The exercise that follows replaces the earlier treatment of productivity as an exogenous contraction in the number of jobs with an explicit task structure. Automation is represented by an increase in machine productivity in automatable tasks, which can simultaneously increase firm scale and reduce the share of tasks performed by labour. Consequently the effect on labour demand, on the employment probability of patient workers, and on the parental return to cultivating patience is not signed *a priori*.

The analysis below is a partial-equilibrium extension of the baseline model. When differentiating with respect to the automation state, the efficiency wage, the number of firms, and the population shares are held fixed. The current automation state is also assumed to be sufficiently persistent that parents expect it to affect the labour market their children will enter. Long-run feedback through wages, entry, and the endogenous composition of households remains governed by the equilibrium mechanism in the baseline model.

IV.A. The technology shock

This productivity shock generates two opposing effects. The first is a **displacement effect**: as machines become more productive, the set of tasks assigned to capital expands, leaving fewer tasks to be performed by workers. The second is a **productivity effect**: cheaper and more productive capital raises aggregate output, which increases demand for labour in the tasks that remain assigned to workers.

This ambiguity feeds into the patience mechanism through the employment probability of patient workers. In the rationing regime, patient workers are employed only after impatient workers have been absorbed, so they are located at the employment margin. Consequently, a change in aggregate labour demand produces a particularly strong change in their employment probability $q(z)$. When displacement dominates, $q'(z) < 0$, reducing the production-work return to patience. When the productivity effect dominates, $q'(z) > 0$, and this channel may instead raise the return to patience. However, aggregate labour demand is only one part of the mechanism. Even when the productivity effect is strong enough to increase total employment, interior automation may still eliminate the intermediate-complexity tasks and middle-skill occupations

that provide relatively accessible routes of upward mobility for children from wage-worker families. Aggregate labour demand can therefore rise while the mobility ladder contracts. There is no contradiction: the productivity effect determines the scale of labour demand, whereas the displacement effect also changes the composition and accessibility of the tasks available to workers. Interior automation specifically pushes workers away from the middle of the task distribution and towards its lower and upper ends.

Let z_t denote machine or artificial-intelligence productivity in automatable tasks. After firms compare the effective costs of machine and labour inputs, let $\tau(z_t)$ denote the reduced-form measure of tasks that remain assigned to workers. The task-assignment margin is therefore summarised by

$$0 < \tau(z_t) \leq 1, \qquad \tau'(z_t) < 0.$$

The model does not separately solve for the optimal capital stock or for the rental price of machines. Instead, $A_T(z_t)$ denotes the net effect of machine productivity on the revenue productivity of the labour-complementary part of a type-T firm, after machine adoption, rental, and operating costs have been taken into account. To avoid confusing entrepreneurial productivity with worker patience, firms are indexed by $T \in \{H, L\}$, where H and L denote high- and low-productivity entrepreneurs. Assume

$$A'_T(z_t) > 0, \qquad T \in \{H, L\}.$$

An increase in z_t therefore operates through two opposing channels. A rise in $A_T(z_t)$ raises the productivity and scale of the firm, generating a productivity or scale effect. A fall in $\tau(z_t)$ reduces the measure of tasks requiring workers, generating a displacement effect. These two reduced-form effects capture the relevant tension without requiring a separate capital-market block.

Both restrictions are inherited from the task-based framework rather than imposed for convenience. That $\tau'(z_t) < 0$ — a rise in machine productivity in automatable tasks shifts the assignment margin against labour — is the displacement effect of Acemoglu and Restrepo (2018, 2019), and it is the formal content of the observation that automation reallocates tasks rather than uniformly scaling factor productivity. That $A'_T(z_t) > 0$ is the productivity

effect: cheaper and more capable machines raise the revenue productivity of the labour-complementary part of the firm, net of adoption, rental, and operating costs. Acemoglu and Restrepo (2019) show that these two forces have opposite signs for labour demand and that neither dominates as a matter of theory, which is why the effect on q_t below is not signed a priori. Restricting $\tau(z_t) \leq 1$ and A_T to be type-specific allows the shock to fall unevenly on high- and low-productivity entrepreneurs, which is what drives the wedge between conditions (6) and (7).

IV.B. Firm labour demand and net operating profit

Let l denote total employment in a firm. Conditional on the reduced-form task assignment, the firm allocates workers symmetrically across the $\tau(z_t)$ labour tasks. Net revenue from the labour-complementary block is

$$y_T(l, z_t) = \chi(\alpha)\,\tau(z_t)\,A_T(z_t)\left[\frac{l}{\tau(z_t)}\right]^{\alpha}, \qquad 0 < \alpha < 1, \qquad \chi(\alpha) \equiv \alpha^{-\alpha}(1-\alpha)^{-(1-\alpha)}. \tag{1}$$

Here y_T is written net of revenues and costs associated solely with the automated block. The formulation is equivalent to $y_T = A_T(z_t)\,\tau(z_t)^{1-\alpha}\,l^{\alpha}$. A firm facing wage w chooses l to maximise $y_T(l, z_t) - wl$. The first-order condition is

$$w = \alpha\chi(\alpha)A_T(z_t)\left[\frac{l}{\tau(z_t)}\right]^{\alpha-1},$$

which gives type-T labour demand

$$l_T(w, z_t) = \frac{\alpha}{1-\alpha}\,\tau(z_t)\left[\frac{A_T(z_t)}{w}\right]^{\frac{1}{1-\alpha}}. \tag{2}$$

Equation (2) nests the original labour-demand expression when $\tau \equiv 1$ and A_T is constant. The corresponding net operating profit is

$$\Pi_T(w, z_t) = \tau(z_t)\,A_T(z_t)^{\frac{1}{1-\alpha}}\,w^{-\frac{\alpha}{1-\alpha}}. \tag{3}$$

Under the normalisation of Section III.B the profit function is exactly the baseline profit scaled by the surviving task share, $\Pi_T(w, z_t) = \tau(z_t)\,\Pi_T(w)$, which makes the nesting transparent and leaves the elasticity in equation (7)

unchanged. This explicit profit expression matters because automation affects not only employment but also the entrepreneurship branch of the return to patience. Under this reduced-form specification, labour demand and net profit are governed by the same displacement-versus-productivity comparison for a given firm type.

Let E_t be the number of firms and let μ_t^e be the share run by high-productivity entrepreneurs. Aggregate labour demand is

$$\begin{aligned} L_t^D(w, z_t) = \frac{\alpha}{1-\alpha} E_t \tau(z_t) w^{-\frac{1}{1-\alpha}} \\ \times \left[\mu_t^e A_H(z_t)^{\frac{1}{1-\alpha}} + \left(1-\mu_t^e\right) A_L(z_t)^{\frac{1}{1-\alpha}} \right]. \end{aligned} \tag{4}$$

IV.C. Displacement versus productivity

Define the displacement elasticity and the type-specific productivity elasticities by

$$\varepsilon_\tau(z_t) \equiv -\frac{z_t \tau'(z_t)}{\tau(z_t)} > 0, \qquad \varepsilon_{A_T}(z_t) \equiv \frac{z_t A_T'(z_t)}{A_T(z_t)} > 0.$$

Let s_H denote the share of aggregate labour demand originating in high-productivity firms:

$$s_H \equiv \frac{\mu_t^e A_H^{\frac{1}{1-\alpha}}}{\mu_t^e A_H^{\frac{1}{1-\alpha}} + \left(1-\mu_t^e\right) A_L^{\frac{1}{1-\alpha}}}.$$

At a fixed wage and fixed firm composition, the elasticity of aggregate labour demand with respect to z_t is

$$z_t \frac{\partial ln L_t^D}{\partial z_t} = -\varepsilon_\tau(z_t) + \frac{s_H \varepsilon_{A_H}(z_t) + (1-s_H)\varepsilon_{A_L}(z_t)}{1-\alpha}. \tag{5}$$

Writing the employment-weighted productivity elasticity as $\bar{\varepsilon}_A \equiv s_H \varepsilon_{A_H} + (1-s_H)\varepsilon_{A_L}$, automation reduces aggregate labour demand if and only if

$$\varepsilon_\tau(z_t) > \frac{\bar{\varepsilon}_A(z_t)}{1-\alpha}. \tag{6}$$

The reverse inequality implies that the productivity effect dominates and aggregate labour demand rises. The factor $1/(1-\alpha)$ reflects the scale

response of a decreasing-returns firm: an increase in revenue productivity raises optimal employment by more than the direct percentage increase in A_T. The model therefore does not impose that automation mechanically reduces employment.

For an individual type-T firm, equation (3) implies

$$z_t \frac{\partial ln\Pi_T}{\partial z_t} = -\varepsilon_\tau (z_t) + \frac{\varepsilon_{A_T}(z_t)}{1-\alpha}. \tag{7}$$

The aggregate labour-demand condition in (6) and the profit condition for a high-productivity entrepreneur need not coincide when automation raises A_H and A_L at different rates. In the homogeneous benchmark $A_H = A_L = A$, both conditions collapse to $\varepsilon_\tau > \varepsilon_A/(1-\alpha)$.

IV.D. Transmission to the employment probability of patient workers

Consider the intermediate rationing regime of the baseline model, in which impatient workers are employed first and patient workers are the marginal group. Their employment probability is

$$q_t (z_t) = \frac{L_t^D (w_p, z_t) - (1-\mu_t^h)(1-E_t)}{\gamma \mu_t^h (1-E_t)}. \tag{8}$$

Here w_p is the relevant efficiency wage and $\gamma > 1$ is the efficiency supplied by a patient worker. Holding w_p, E_t, μ_t^e, and μ_t^h fixed, differentiation gives

$$\frac{\partial q_t}{\partial z_t} = \frac{1}{\gamma \mu_t^h (1-E_t)} \frac{\partial L_t^D}{\partial z_t}, \qquad sign(q_z) = sign\left(L_z^D\right). \tag{9}$$

Thus, within this partial-equilibrium exercise, the displacement-productivity comparison in (6) also determines the direction of the patient employment probability. Let $N_I \equiv (1-\mu_t^h)(1-E_t)$ denote the inframarginal employment of impatient workers. The elasticity of q_t is

$$\frac{z_t}{q_t}\frac{\partial q_t}{\partial z_t} = \frac{L_t^D}{L_t^D - N_I}\left(z_t \frac{\partial ln L_t^D}{\partial z_t}\right). \tag{10}$$

Small proportional movements in aggregate labour demand are magnified in the hiring probability of patient workers, because impatient employment is inframarginal and all the adjustment falls on the tail of the queue. The more binding the rationing (the closer $N_I$is to L^D), the stronger the amplification.

Patient workers are the residual claimant on labour demand. Impatient workers are hired first and their employment does not adjust within the intermediate rationing regime, so every unit of change in aggregate demand lands entirely on the patient tail of the queue.

The leading factor exceeds one whenever rationing is genuine. A given proportional change in aggregate labour demand therefore produces a larger proportional change in the hiring probability of patient workers, because they occupy the margin of the employment queue.

When the employment branch is the active source of the return to patience, the minimum patient employment probability required to cover the parental investment cost is

$$q^{crit} = \frac{1}{\gamma} + \frac{K_h}{\gamma \underline{w}_P \Delta\pi_h}. \tag{11}$$

This threshold is the automation-model counterpart of the worker-parent switching condition of Section III.I, and it uses the effort threshold in the form derived there, $\bar{m}_h^{PPE} = K_h/\Delta\pi_h$. It is a branch-specific threshold, not a general threshold for the complete model. Parents may still cultivate patience when $q_t < q^{crit}$ if the entrepreneurship branch yields a sufficiently high return. The employment branch can support patience investment only if $q^{crit} \leq 1$, equivalently $K_h/\Delta\pi_h \leq (\gamma - 1)\underline{w}_P$, which is the same parameter restriction stated in Section III.I. If $q_t\ (z_t)$ is strictly decreasing over a relevant interval and crosses q^{crit} once, a local critical value z^{crit} may be defined by $q_t\left(z^{crit}\right) = q^{crit}$. No global claim that all $z_t > z^{crit}$ eliminate patience investment is warranted unless monotonicity and branch dominance are maintained over that entire range.

IV.E. Aggregate labour demand and accessible mobility opportunities

Aggregate labour demand and the set of jobs that form an accessible upward-mobility ladder are distinct objects. Automation can expand total employment while removing middle-skill positions through which children from worker households could previously progress. To represent this distinction, define

$$M_t\ (z_t) = M_t^0 - \beta D\ (z_t) + \xi G\ (z_t)\,, \qquad \beta, \xi > 0. \tag{12}$$

Here $D(z_t)$ measures the accessible middle-skill or progression tasks displaced by machines, rather than all intermediate-complexity tasks, and $G(z_t)$ measures newly created high-value tasks that workers with foundational skills can realistically enter. Hence

$$M_z = -\beta D'(z_t) + \xi G'(z_t). \tag{13}$$

Purely substituting automation is characterised by $D' > 0$ and a small G', in which case $M_z < 0$. Worker-augmenting artificial intelligence may instead make previously expert tasks accessible to ordinary workers, raising G' and potentially making M_z positive. It is therefore possible to observe $L_z^D > 0$ and $M_z < 0$ at the same time: automation may increase aggregate firm scale while weakening the occupational ladder relevant for mobility.

Equation (12) is a modelling assumption rather than an implication of the assignment framework, and we state it as such. It requires that a subset of middle-complexity occupations constitutes the principal progression route available to worker households, which is the interpretation supported by the polarisation literature: employment growth has been concentrated at the top and bottom of the occupational skill distribution while the middle has hollowed out (Autor, Levy, and Murnane, 2003; Goos, Manning, and Salomons, 2014), and the middle-skill occupations that disappeared were disproportionately those offering internal promotion ladders and on-the-job training. The creation term $\xi G(z_t)$ captures the countervailing force emphasised by Acemoglu and Restrepo (2018) and by the recent evidence that generative technologies compress performance gaps by making previously expert tasks accessible to less experienced workers (Brynjolfsson, Li, and Raymond, 2025; Noy and Zhang, 2023). Whether M_z is positive or negative is therefore an empirical question about the balance of displacement and creation at accessible skill levels, not something the model can settle. Similarly, a change in the automation regime may create a kink in $M_t(z_t)$, but a discrete jump requires the additional assumption that a discrete block of accessible progression tasks disappears at the regime boundary; continuity or discontinuity should be stated as an assumption rather than attributed to the task-assignment framework.

IV.F. The return to patience

The worker household compares the return obtained through patient employment with the return obtained through entrepreneurial entry. Using H for a high-productivity entrepreneur and p for a patient worker, the return to cultivating patience is

$$R_t^W(z_t) = max\left\{q_t(z_t)\gamma w_p,\ \Pi_H\left(w_p, z_t\right) - b\left[M_t(z_t)\right]w_p\right\} - w_p. \quad (14)$$

Parents cultivate patience if and only if

$$R_t^W(z_t) \geq \bar{m}_h^{PPE} = \frac{K_h}{\Delta\pi_h}. \quad (15)$$

Suppose first that one branch is strictly active, so that the upper envelope is locally differentiable. Holding the wage fixed, the derivative on the employment branch is

$$\frac{\partial R_t^W}{\partial z_t} = \gamma w_p q_z \quad (16)$$

whereas the derivative on the entrepreneurship or mobility branch is

$$\frac{\partial R_t^W}{\partial z_t} = \Pi_{H,z} - b'(M_t)\, w_p M_z. \quad (17)$$

Because $b'(M_t) < 0$, a reduction in accessible mobility opportunities, $M_z < 0$, lowers the return on the mobility branch. However, that negative effect must be compared with the direct change in high-productivity entrepreneurial profit. The mobility branch falls if and only if

$$\Pi_{H,z} < \left|b'(M_t)\right| w_p \left|M_z\right| \qquad \text{when } M_z < 0. \quad (18)$$

The automation effect is therefore governed by three margins: the net effect on aggregate labour demand, the net effect on accessible mobility opportunities, and the direct effect on high-productivity entrepreneurial profit.

Sufficient condition. A convenient sufficient condition for automation to reduce the return to patience on both branches is

$$\varepsilon_\tau > \frac{\bar{\varepsilon}_A}{1-\alpha}, \qquad \varepsilon_\tau > \frac{\varepsilon_{A_H}}{1-\alpha}, \qquad M_z \leq 0. \quad (19)$$

The first inequality makes q_z negative through aggregate labour demand. The second makes $\Pi_{H,z}$ negative for the entrepreneur type relevant to the mobility branch. The third ensures that the change in accessible mobility does not offset that profit decline. In the homogeneous benchmark, the first two inequalities coincide. More generally, condition (18) can replace the second and third restrictions when automation raises entrepreneurial profit but the loss of accessible mobility opportunities is larger.

At points where the two branches are equal, R_t^W is an upper envelope and need not be differentiable. Statements about comparative statics at such points should use one-sided derivatives. Because the active branch may change as z_t varies, R_t^W need not be globally monotone, and multiple threshold crossings are possible. They are a possibility of the model, not a generic prediction.

IV.G. The switching locus and equilibrium interpretation

The baseline model may retain a low-patience and a high-patience steady state. The present extension shifts the set of entrepreneurial compositions for which worker parents find patience investment worthwhile. Let $\bar{\mu}^e(z_t)$ be defined locally by

$$R^W(\bar{\mu}^e, z_t) = \bar{m}_h^{PPE} = \frac{K_h}{\Delta\pi_h}. \tag{20}$$

If a single branch is strictly active, R^W is continuously differentiable, and $R_{\mu^e}^W > 0$, the implicit-function theorem gives

$$\frac{\partial \bar{\mu}^e}{\partial z_t} = -\frac{R_z^W}{R_{\mu^e}^W}. \tag{21}$$

Accordingly, if $R_z^W < 0$, the threshold $\bar{\mu}^e$ rises: a larger share of high-productivity entrepreneurs is required before worker parents cultivate patience. If $R_z^W > 0$, the threshold falls and patience investment becomes consistent with a wider set of entrepreneurial compositions.

Equation (21) is a comparative-static result for the switching locus. By itself, it does not establish that the basin of attraction of a steady state expands or that the economy is tipped from one steady state to another. Those stronger claims require an explicit dynamic law of motion for μ_t^h and μ_t^e, local-stability conditions, and proof that the switching locus is the relevant separatrix.

Without that additional dynamic analysis, the appropriate interpretation is that automation expands or contracts the region in which parents choose to cultivate patience.

V. Conclusion

There is growing interest in understanding how preferences shape economic outcomes, and how these preferences themselves respond to economic environments. A dominant approach in the literature has treated time preferences as relatively stable traits - formed early in life and slow to change - thereby placing greater emphasis on their effects on economic behaviour rather than on the economic forces that shape them. While this perspective allows for the co-evolution of preferences and outcomes, time preferences are often assumed to exert a more fundamental influence on economic outcomes than the reverse.

This paper advances an alternative perspective by emphasizing the endogeneity of time preferences to economic incentives. Building on recent work highlighting the fluidity of cultural traits, we have study how labour market equilibria and social mobility shape families' incentives to cultivate patience in their children. Motivated by substantial cross-country heterogeneity in time preferences, we develop a model in which patience and impatience arise as optimal responses to the expected returns to long-term investment, which are in turn shaped by labour demand, wage levels, and the scope for upward mobility.

In the model, patience is valuable when future-oriented investments, such as education and skill accumulation, yield sufficiently high expected returns. Because cultivating patience is itself a costly cognitive effort, parents adopt it only once the expected return clears a threshold, and loss aversion creates a band in which patience is feasible but not yet chosen. This generates four parenting regimes, from an informational trap in which patience is never contemplated, through a cognitive-effort trap in which it is recognised but not adopted, to expectation-driven regions in which it becomes the preferred course. When social mobility is high and labour market conditions are favourable, parents optimally invest in costly efforts to cultivate patience in their children. By contrast, when social mobility is low and wages are weak for

workers, impatience can function as an adaptive preference: impatient workers are more willing to accept existing wage offers and exert effort under limited incentives, whereas patient workers may be disadvantaged. Anticipating these constraints, parents may rationally refrain from investing in the cultivation of patience, allowing impatience to persist as an optimal response to limited economic opportunities.

A key implication of the framework is that the effects of labour market equilibria are not uniform across families. Differences in expected mobility generate systematic heterogeneity in parenting choices, which in turn feeds back into economic outcomes. When patience is costly to cultivate and yields low returns, parenting values adjust endogenously, reinforcing existing patterns of inequality. When mobility is high enough, however, the reverse occurs: even impatient parents come to cultivate patience, so that impatience is no longer mechanically reproduced across generations. Because a larger share of patient entrepreneurs raises labour demand and hence the return to patience for workers, while patient workers in turn enlarge the future pool of patient entrepreneurs, the two sides of the economy reinforce one another, and the system converges into either a low-patience or a high-patience regime. In this way, labour market conditions and parenting choices interact to produce persistent cross-country differences in time preferences, even in the absence of intrinsic differences in cultural traits.

The analysis also highlights several promising directions for future research. First, there is a clear need for more systematic empirical work on the causal relationship between labour market conditions, social mobility, and parenting in patience. While the evidence presented in this paper documents robust cross-country patterns, identifying the underlying causal mechanisms remains an important challenge.

Second, the paper has focused on vertical transmission of time preferences through parenting choices. An important extension would be to allow for horizontal transmission through schools, peer groups, and social environments. Differences in educational institutions or social segregation may play a key role in shaping the incentives to cultivate patience and could interact with labour market conditions in important ways.

Third, the model abstracts from within-family conflicts that may arise when patient children contribute less to within-family public goods in the

present. Incorporating such intra-household trade-offs could further illuminate why cultivating patience requires substantial effort and why some families optimally choose not to do so, even when patience is socially desirable.

Finally, following recent work emphasizing the fluidity of cultural traits, it would be interesting to allow time preferences to adjust over the life cycle in response to changing economic conditions, rather than being fixed during childhood. Such extensions could help bridge the gap between short-run economic shocks and long-run changes in preferences.

Overall, the framework developed in this paper highlights how differences in time preferences need not reflect fixed cultural traits or endowments but can instead emerge endogenously from economic incentives shaped by labour market equilibria and social mobility. By placing parenting choices at the centre of this process, the analysis offers a unified perspective on the interaction between economic structure, preference formation, and the persistence of inequality.

References

Acemoglu, Daron. 2022. "Obedience in the Labour Market and Social Mobility: A Socioeconomic Approach." *Economica* 89 (S2): S2-S37.

Acemoglu, Daron, and Jonas Loebbing. 2026. “Automation and Polarization.” Working Paper, MIT.

Acemoglu, Daron, and Pascual Restrepo. 2018. “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment.” A.E.R. 108 (6): 1488–1542.

Acemoglu, Daron, and Pascual Restrepo. 2019. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” J. Econ. Perspectives 33 (2): 3–30.

Agostinelli, Francesco, Matthias Doepke, Giuseppe Sorrenti, and Fabrizio Zilibotti. 2025. "A Stairway to Success: How Parenting Shapes Culture and Social Stratification." NBER Working Paper no. 33963.

Alan, Sule, and Seda Ertac. 2018. "Fostering Patience in the Classroom: Results from a Randomized Educational Intervention." *J.P.E.* 126 (5): 1865-1911.

Alesina, A., Stantcheva, S., & Teso, E. (2018). Intergenerational mobility and preferences for redistribution. *American Economic Review, 108*(2), 521–554. https://doi.org/10.1257/aer.20162015

Anger, Silke, and Guido Heineck. 2010. "Do Smart Parents Raise Smart Children? The Intergenerational Transmission of Cognitive Abilities." *J. Population Econ.* 23 (3): 1105-32.

Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." A.E.R. 103 (5): 1553–97.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." Q.J.E. 118 (4): 1279–1333.

Baland, Jean-Marie, and James A. Robinson. 2000. "Is Child Labor Inefficient?" *J.P.E.* 108 (4): 663-79.

Benjamin, Daniel J., Sebastian A. Brown, and Jesse M. Shapiro. 2013. "Who Is 'Behavioral'? Cognitive Ability and Anomalous Preferences." *J. European Econ. Assoc.* 11 (6): 1231-55.

Bisin, Alberto, and Thierry Verdier. 2000. "'Beyond the Melting Pot': Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits." *Q.J.E.* 115 (3): 955-88.

Bisin, Alberto, and Thierry Verdier. 2001. "The Economics of Cultural Transmission and the Dynamics of Preferences." J. Econ. Theory 97 (2): 298–319.

Björklund, Anders, and Markus Jäntti. 1997. "Intergenerational Income Mobility in Sweden Compared to the United States." *A.E.R.* 87 (5): 1009-18.

Bloom, Nicholas, and John Van Reenen. 2007. "Measuring and Explaining Management Practices across Firms and Countries." Q.J.E. 122 (4): 1351–1408.

Borgonovi, F., & Pokropek, A. (2025). Young people's social mobility expectations in an unequal world. *Science, 388*(6746), 599–601. https://doi.org/10.1126/science.adt1842

Bortolotti, Stefania, Thomas Dohmen, Hartmut Lehmann, Frauke Meyer, Norberto Pignatti, and Karine Torosyan. 2021. "Patience, Cognitive Abilities, and Cognitive Effort: Survey and Experimental Evidence from a Developing Country." American Behavioral Scientist 65 (11): 1512–30.

Brenøe, Anne Ardila, and Thomas Epper. 2022. "Parenting Values and the Intergenerational Transmission of Time Preferences." *European Econ. Rev.* 148: 104208.

Brown, Heather, and Marjon van der Pol. 2015. "Intergenerational Transfer of Time and Risk Preferences." *J. Econ. Psych.* 49: 187-204.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2025. "Generative AI at Work." Q.J.E. 140 (2): 889–942.

Chowdhury, Shyamal, Matthias Sutter, and Klaus F. Zimmermann. 2022. "Economic Preferences across Generations and Family Clusters: A Large-Scale Experiment in a Developing Country." *J.P.E.* 130 (9): 2361-2410.

Cunha, Flavio, and James J. Heckman. 2007. "The Technology of Skill Formation." *A.E.R.* 97 (2): 31-47.

Cunha, Flavio, Sofia Nihtianova, James Rood, and Anna-Lena van der Merwe. 2026. "Beliefs, Attention, and Investments in Early Childhood." NBER Working Paper no. 35150.

Dawson, C., & Failla, V. (2022). *Impatience and entrepreneurship*. SSRN. https://doi.org/10.2139/ssrn.4151432

DellaVigna, Stefano, and M. Daniele Paserman. 2005. "Job Search and Impatience." *J. Labor Econ.* 23 (3): 527-88.

Deming, David J. 2017. "The Growing Importance of Social Skills in the Labor Market." *Q.J.E.* 132 (4): 1593-1640.

Doepke, M., & Zilibotti, F. (2013). Culture, entrepreneurship, and growth. In P. Aghion & S. N. Durlauf (Eds.), *Handbook of economic growth* (Vol. 2A, pp. 1–48). Elsevier. https://doi.org/10.1016/B978-0-444-53538-2.00001-0

Doepke, Matthias, and Fabrizio Zilibotti. 2008. "Occupational Choice and the Spirit of Capitalism." *Q.J.E.* 123 (2): 747-93.

Doepke, Matthias, and Fabrizio Zilibotti. 2017. "Parenting with Style: Altruism and Paternalism in Intergenerational Preference Transmission." *Econometrica* 85 (5): 1331-71.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde. 2018. "Global Evidence on Economic Preferences." *Q.J.E.* 133 (4): 1645-92.

Falk, Armin, Anke Becker, Thomas Dohmen, David Huffman, and Uwe Sunde. 2016. "The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences." IZA Discussion Paper no. 9674.

Falk, Armin, Fabian Kosse, Pia Pinger, Hannah Schildberg-Hörisch, and Thomas Deckers. 2021. "Socioeconomic Status and Inequalities in Children's IQ and Economic Preferences." *J.P.E.* 129 (9): 2504-45.

Freeman, David J. 2019. "Expectations-Based Reference-Dependence and Choice under Risk." *Econ. J.* 129 (622): 2424-58.

Golsteyn, Bart H. H., Hans Grönqvist, and Lena Lindahl. 2014. "Adolescent Time Preferences Predict Lifetime Outcomes." *Econ. J.* 124 (580): F739-F761.

Goos, Maarten, Alan Manning, and Anna Salomons. 2014. "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring." A.E.R. 104 (8): 2509–26.

Haerpfer, Christian, Ronald Inglehart, Alejandro Moreno, Christian Welzel, Kseniya Kizilova, Jaime Diez-Medrano, Marta Lagos, Pippa Norris, Eduard Ponarin, and Bi Puranen. 2022. "World Values Survey Wave 7 (2017-2022) Cross-National Data-Set." World Values Survey Assoc.

Hai, Rong, and James J. Heckman. 2022. "The Causal Effects of Youth Cigarette Addiction and Education." NBER Working Paper no. 30304.

Heckman, James J., Jora Stixrud, and Sergio Urzua. 2006. "The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior." *J. Labor Econ.* 24 (3): 411-82.

Hirshleifer, David, Angie Low, and Siew Hong Teoh. 2012. "Are Overconfident CEOs Better Innovators?" *J. Finance* 67 (4): 1457-98.

Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991). Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives, 5*(1), 193–206. https://doi.org/10.1257/jep.5.1.193

Kőszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." Q.J.E. 121 (4): 1133–65.

Kőszegi, Botond, and Matthew Rabin. 2007. "Reference-Dependent Risk Attitudes." A.E.R. 97 (4): 1047–73.

Kwon, Seok-Woo, Colleen Heflin, and Martin Ruef. 2013. "Community Social Capital and Entrepreneurship." *Amer. Sociological Rev.* 78 (6): 980-1008.

List, John A., Julie Pernaudet, and Dana L. Suskind. 2021. "Shifting Parental Beliefs about Child Development to Foster Parental Investments and Improve School Readiness Outcomes." *Nature Comm.* 12: 5765.

Lucas, Robert E., Jr. 1978. "On the Size Distribution of Business Firms." Bell J. Econ. 9 (2): 508–23.

Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao. 2013. "Poverty Impedes Cognitive Function." Science 341 (6149): 976–80.

Michelman, Valerie, Joseph Price, and Seth D. Zimmerman. 2022. "Old Boys' Clubs and Upward Mobility among the Educational Elite." *Q.J.E.* 137 (2): 845-909.

Muñoz, Eduardo, and Roy Van der Weide. 2023. "Intergenerational Income Mobility around the World: A New Database." Policy Research Working Paper, World Bank Group.

Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." Science 381 (6654): 187–92.

Ryan, R., Kalil, A., Schneider, M. M., Destin, M., Silverman, D., & Hernandez, I. A. (2026). Economic mobility beliefs and parental investment: Experimental evidence. *Economics Letters, 267*, Article 113090. https://doi.org/10.1016/j.econlet.2026.113090

Saez-Marti, Maria, and Fabrizio Zilibotti. 2008. "Preferences as Human Capital: Rational Choice Theories of Endogenous Preferences and Socioeconomic Changes." *Finnish Econ. Papers* 21 (2): 81-94.

Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty, 1*(1), 7–59. https://doi.org/10.1007/BF00055564

Shapiro, Carl, and Joseph E. Stiglitz. 1984. “Equilibrium Unemployment as a Worker Discipline Device.” A.E.R. 74 (3): 433–44.

Sunde, Uwe, Thomas Dohmen, Benjamin Enke, Armin Falk, David Huffman, and Gerrit Meyerheim. 2022. "Patience and Comparative Development." *Rev. Econ. Stud.* 89 (5): 2806-40.

Sutter, Matthias, Martin G. Kocher, Daniela Glätzle-Rützler, and Stefan T. Trautmann. 2013. "Impatience and Uncertainty: Experimental Decisions Predict Adolescents' Field Behavior." *A.E.R.* 103 (1): 510-31.

Topel, R. H., & Ward, M. P. (1992). Job mobility and the careers of young men. *The Quarterly Journal of Economics, 107*(2), 439–479. https://doi.org/10.2307/2118478

Tversky, Amos, and Richard H. Thaler. 1990. "Anomalies: Preference Reversals." *J. Econ. Perspectives* 4 (2): 201-11.

Weiss, Andrew. 1980. “Job Queues and Layoffs in Labor Markets with Flexible Wages.” J.P.E. 88 (3): 526–38.

Wolfe, Marcus T., and Pankaj C. Patel. 2016. "Grit and Self-Employment: A Multi-Country Study." *Small Bus. Econ.* 47 (4): 853-74.