

Disadvantage and Beliefs

Patricio S. Dalton, Sayantan Ghosal & Damiano Turchet

January 2026

No: 1596

Warwick Economics Research Papers

ISSN 2059-4283 (online)

ISSN 0083-7350 (print)

Disadvantage and Beliefs*

Patricio S. Dalton[†]
Sayantan Ghosal[‡]
Damiano Turchet[§]

January 21, 2026

Abstract

We study how structural disadvantage (e.g., race, class, gender) shapes the formation of subjective beliefs about the returns to effort. We formalize and distinguish between psychological constructs such as locus of control, self-efficacy, and grit, and study their response to structural disadvantage and policy interventions. In our model, individuals share the same (unknown) innate ability, and beliefs about success can only be updated through effort. However, agents fail to internalize the dynamic feedback between effort and belief formation. Structural disadvantage raises the threshold of belief required to justify effort, increasing the likelihood of falling into a pessimistic low-effort trap. We characterize the conditions under which psychological interventions that bound beliefs from below and enhance grit improve welfare, and when such interventions must be complemented by policies that relax external constraints to be effective.

Keywords: Disadvantage; Self-Beliefs; Pessimism; Grit; Role Models; Aspirations.

JEL Classification Codes: D90, D91, I30.

*We thank Sule Alan, Prateek Bhan, Elena Cettolin, Martin Cripps, Johannes Haushofer, Jalnidh Kaur, Mahreen Mahmud, Maddie McKelway, Matthew Ridley and participants at the TiBER 2025 symposium for helpful comments and suggestions. All errors are our own.

[†]Tilburg University, Department of Economics and CentER, Warandelaan 2, 5037 AB Tilburg, the Netherlands.

[‡]Adam Smith Business School, University of Glasgow, Glasgow G12 8QQ, UK

[§]Department of Economics, University of Warwick, CV4 7AL Coventry, UK

1 Introduction

Noncognitive or *soft skills*—such as *locus of control*, *self-efficacy*, and *grit*—play a central role in economic behavior.¹ Following Heckman et al. (2006) and Heckman and Kautz (2012), economists have shown that these traits predict educational attainment, employment stability, health investment, and entrepreneurial behavior (Alan et al. 2019). Although often treated as stable, these psychological constructs vary substantially between individuals and environments. The work of Heckman et al. (2006) and Duckworth et al. (2007) documents large heterogeneity in noncognitive skills, and growing evidence links an external locus of control, weaker self-efficacy, and reduced grit to socioeconomic disadvantage and repeated exposure to adversity. Yet little is known about whether and how structural disadvantage shapes such noncognitive skills or about the distinct ways in which these different noncognitive skills influence perceived returns to effort. Answering these questions is key to designing effective policies that reduce persistent inequality and foster resilience in the face of adversity.

This paper proposes a model of adaptive belief formation to explain how structural disadvantage shapes perceived returns to effort, ultimately leading to persistent gaps in noncognitive skills and economic outcomes. We define structural disadvantage as arising from factors beyond an individual’s control—such as race, gender, caste, class, or inherited wealth—that systematically affect economic opportunities. In our model, individuals share the same (unknown) objective probability of success $q \in [0, 1]$ when exerting effort, but differ in their circumstances. Effort choices are guided by current subjective beliefs about the probability of success p_t , which can only be updated by exerting effort.² However, individuals exhibit a form of myopia: they fail to internalize how today’s effort improves future beliefs, and thus future incentives. Since effort is costly, it is only chosen when the perceived probability of success exceeds a belief threshold. A higher level of disadvantage raises this threshold, reducing the likelihood of exerting effort in face of failure and increasing the risk of falling into a **pessimistic low-effort trap**. Over time, this discourages persistence and reinforces inaction. **A lack of grit and an external locus of control emerge endogenously**, as individuals abandon effort before beliefs have the chance to improve.

The model generates two long-term outcomes. One is an optimistic, high-effort path in which individuals persist and their subjective beliefs eventually converge, in expectation, to the true probability of success. The other is a pessimistic, low-effort trap in which beliefs fall below the threshold and remain there. These long-term outcomes are welfare-ranked. When $q = 1$ —that is, when agents always succeed in exerting effort—then

¹Locus of control captures whether individuals attribute outcomes to their own actions rather than external forces (Rotter 1966; Cobb-Clark 2015), and self-efficacy reflects beliefs about one’s ability to carry out goal-directed actions (Bandura 1977). Grit is defined as the ability to maintain effort in the face of setbacks to achieve long-term goals (Duckworth et al. 2007).

²Subjective beliefs about the probability of success p_t conditional on exerting effort is interpreted as an individual’s **locus of control**. These beliefs are distinct from prior beliefs over q , a point we return to below which we interpret as **self-efficacy** beliefs. For a conceptual distinction between locus of control and self-efficacy, see Ajzen 2002.

there is convergence to the optimistic high-effort or the pessimistic low-effort outcome, depending on their initial beliefs. In this case, the optimistic high-effort outcome welfare-dominates the pessimistic low-effort one. If initial beliefs are drawn from the same underlying probability distribution for all agents, those facing higher levels of disadvantage are more likely to converge to the pessimistic, low-effort outcome.³

Such multiple welfare-ranked outcomes also arise in models of behavioral poverty traps, such as those based on aspirations (Dalton et al. 2016). However, when success is stochastic, that is, when $0 < q < 1$, our model departs sharply from these frameworks. In this case, there is only one long-term outcome: a *pessimistic, low-effort absorbing state*. The optimistic, high-effort path is fragile and sustained only if individuals repeatedly succeed or receive persistent support. Unlike in standard Bayesian learning models, agents here do not have a prior belief in the true value of the probability of success q . Instead, they form and update subjective beliefs about success conditional on effort p_t based solely on past outcomes. Because effort is costly and belief updating requires effort, structural disadvantage makes it harder to justify continued investment in the face of failure.

This mechanism generates a key implication: even when individuals share the same fundamentals—preferences, objective probability of success, initial beliefs, and behavioral bias—those facing greater structural disadvantage are more likely to fall into a pessimistic low-effort trap. The model does not rely on differences in preferences or inherent motivation. Instead, it shows how external circumstances distort incentives, raising the belief threshold needed to justify exerting effort. This makes it harder for disadvantaged individuals to persist long enough for their beliefs to improve. As a result, otherwise identical individuals can diverge sharply in outcomes purely due to differences in structural constraints. This pessimistic low-effort trap is a **suboptimal** outcome, in which individuals stop trying despite the fact that effort would eventually lead to better beliefs and outcomes. Psychologically, it resembles the phenomenon of **learned helplessness** (Seligman 1972), where individuals, after repeated failures or negative feedback, become convinced that they have no control over the results and thus cease exerting effort, even when success is possible. In our setting, structural disadvantage induces a similar dynamic: it undermines motivation not by reducing ability but by distorting perceptions of success conditional on effort, thereby making learned helplessness and external locus of control an endogenous consequence of structural barriers.

Extensions. We extend this baseline model along four dimensions to test the robustness of our results and to explore related mechanisms. In Appendix B we consider agents who are **farsighted** and internalize how current effort affects future beliefs. We show that farsightedness preserves a key feature of the short-term decision protocol, namely that once the decision-maker switches to a low-effort, pessimistic outcome, they remain there in all subsequent time periods. This feature of our short-term decision protocol is key to establishing a pessimistic low-effort trap and is preserved even when a decision-maker is farsighted. Moreover, even forward-looking individuals can choose zero effort

³In the extreme case where agents never succeed (i.e., $q = 0$), or more generally when q falls below a critical threshold, the optimistic, high-effort outcome becomes welfare-dominated. Continued effort yields no perceived gains and only repeated failure, making inaction a rational choice.

in all periods if expected future returns remain too low relative to the cost of exerting effort. However, in contrast to decision outcomes with a short-term decision protocol, farsighted decision outcomes cannot be welfare-ranked.

In Appendix C, we incorporate **aspirations** as an additional psychological determinant of effort. Aspirations lower the belief threshold required to justify effort and can encourage persistence. However, if aspirations evolve slowly or are repeatedly unmet, they may result in frustration or collapse altogether—leading to distinct long-term traps of either **frustration** or **fatalism**. In Appendix D, we modify the model to allow disadvantage to reduce the **objective probability of success**, rather than the return to success. The core dynamics and predictions remain: disadvantaged individuals are still more likely to fall into pessimistic traps due to an increased likelihood of failure and downward belief revisions .

Finally, Appendix E extends the model to allow for **Bayesian learning**, where agents hold a prior distribution over q , representing their own ability. This extension provides a natural way to incorporate **self-efficacy**, the belief in one’s own capacity to succeed through effort (Bandura 1977), into the framework. In this setting, changes in prior beliefs about ability, denoted by P_0 , reflect evolving assessments of innate talent. This contrasts with the baseline model, where agents form subjective beliefs p_t about the probability of success conditional on effort, which we interpret as capturing an individual’s **locus of control**, that is, whether they believe success is contingent on their own actions.

These two belief types are conceptually and formally distinct. While p_t is not a Bayesian prior, it evolves non-parametrically based on past experiences (as shown in Proposition 1), and in the limit converges to the true q if effort is consistently exerted. In contrast, priors over q are updated via Bayes’ rule and similarly converge to the true ability level under sustained effort. The non-parametric nature of p_t allows for a more flexible modeling of belief formation, particularly relevant when agents face uncertainty in complex or unfamiliar environments.⁴

Across all four extensions, the key insight remains: grit, beliefs, and disadvantage interact dynamically to shape effort decisions over time, and small differences in structural constraints can compound into persistent disparities in outcomes.

Policy Implications. The analysis suggests three key policy implications. First, *psycho-*

⁴It is possible to combine beliefs about ability and environmental responsiveness in a unified framework—for example, by letting $q = a \cdot r$, where $a \in [0, 1]$ represents beliefs about one’s own ability (self-efficacy) and $r \in [0, 1]$ captures beliefs about how strongly effort translates into outcomes (locus of control). In this paper, we abstract from this interaction and assume that all agents share the same objective ability. Our focus is not on heterogeneity in q itself, but on how beliefs about the returns to effort evolve under disadvantage. More generally, modeling both dimensions jointly would require a parametric, Bayesian framework with beliefs defined over multiple latent traits. Such an approach relies on the “small worlds” assumption—that agents can specify correct priors and likelihoods over all relevant states (see Binmore 2008). This assumption may be inappropriate in contexts of structural disadvantage, where individuals face complex environments and limited feedback, making structured probabilistic reasoning difficult. A non-parametric, experience-based updating process may therefore provide a more realistic description of belief formation in these settings. Whether disadvantaged individuals rely more on experiential or Bayesian learning remains an important question for future research.

logical or informational interventions that help individuals maintain or increase their beliefs about success—conditional on exerting effort—can improve welfare, but only when the true probability of success q exceeds the threshold level of belief required to justify effort. In such cases, maintaining optimistic beliefs is critical to ensure persistence. These interventions can range from providing information on returns to effort to growth mindset training, coaching and mentoring, connection to successful peers (role models), and cognitive behavioral therapy (CBT), all of which can support motivation in the face of adversity. In Section 3, we discuss empirical evidence supporting the effectiveness of interventions targeting these mechanisms.

Second, *material interventions*, such as cash transfers, educational subsidies, increased access, or health programs that directly mitigate the negative effects of structural disadvantage, can lead to welfare gains. They work by lowering the belief threshold required for action, either through increasing the objective probability of success (e.g., with affirmative action or reservation policies) or by reducing the cost of effort (e.g., with cash transfers). However, these interventions may fail when the beliefs of the individuals are already too low to motivate effort. This limitation has been documented in cash-based anti-poverty and development programs. For example, Blattman et al. (2014) find that cash grants increased earnings for some Ugandan youth but had limited effects for others, potentially due to differences in psychological readiness. Wollburg et al. (2023) conduct a meta-analysis showing that cash transfers can alleviate depression and anxiety in low- and middle-income countries, especially in the short term. However, they find limited and mixed evidence on the persistence of these effects after financial support ends, underscoring the importance of complementary interventions that address underlying psychological constraints.

Third, *combining both* approaches yields the strongest effect. Psychological and material supports are complementary to shifting individuals to the high-effort optimistic path. Evidence from multifaceted ‘graduation’ programs supports this view. Banerjee et al. (2015) show that the combination of asset transfers, training, and mentoring produces sustained gains in income and well-being in six countries. Similarly, Bryan et al. (2010) provide evidence that combining aspirations-based messaging with capital grants significantly boosts investment and productivity.

Related Literature. This paper contributes to several strands of literature. First, it contributes to the literature on *behavioral economics* by providing a unified, formal framework that connects several core psychological constructs such as locus of control, grit, self-efficacy, and aspirations to economic decision-making. Although these concepts have been widely discussed in psychology (Duckworth et al. 2007), to the best of our knowledge, our model is the first to explicitly incorporate them into a formal economic framework. It also provides a tractable representation of how disadvantage interacts with subjective beliefs, persistence, and disengagement over time. Moreover, the pessimistic low-effort trap that emerges in our model offers a microfoundation for the psychological concept of *learned helplessness* (Seligman 1972), where agents stop exerting effort despite having the ability to succeed.

Within the broad literature of behavioral economics, this paper contributes to the theoret-

ical foundations of *behavioral poverty traps*, which explain persistent disadvantage through internal constraints such as aspirations, self-control, or cognitive capacity (see Kremer et al. 2019 for a review). Our model shares with Dalton et al. 2016 the idea that disadvantage (poverty in their case) affects behavior through distorted internal states: aspirations in their model and beliefs about the return to effort in ours (with aspirations included as an extension). However, the mechanisms differ in important ways. Aspirations models are based on reference-dependent preferences, generating multiple steady states based on individuals' goals and constraints. In contrast, our model centers on a dynamic feedback loop between effort and belief formation, where beliefs evolve only through sustained effort. This results in a distinct kind of trap—one not driven by adaptive preferences but by biased belief updating under structural constraints. Furthermore, our framework is more general, since it extends beyond the economic domain of poverty to capture structural disadvantage in forms such as race, gender, and class. Our framework also relates to models linking poverty and distorted beliefs, such as de Quidt and Haushofer (2016) and Treurniet and Lensink (2022). Although these studies emphasize exogenous shocks or poorly designed interventions that weaken agency, we show how structural disadvantage alone can endogenously erode beliefs and effort. Compared to models of self-control (Bernheim et al. 2015), hope (Lybbert and Wydick 2018) or cognitive bandwidth (Mani et al. 2013), we focus on the endogenous evolution of beliefs and its interaction with material constraints.

Second, our model contributes to the literature on *distorted beliefs and learning-by-doing*. A growing body of work explores how beliefs can be systematically biased due to external shocks (e.g. Malmendier and Nagel 2011), stigma (Ghosal et al. 2022), or limited experience with relevant outcomes (e.g. Kuchler and Zafar 2019; Malmendier and Nagel 2016). We complement this literature by proposing a mechanism in which beliefs become distorted not through misperception or shocks, but through a form of dynamic inattention bias: the fact that individuals do not internalize how today's effort improves future beliefs. We build on rational learning models where, under constant effort, beliefs converge in expectation to the objective probability of success q , such as Veldkamp (2005) and Van Nieuwerburgh and Veldkamp (2006). In this literature, we introduce behavioral friction. The model also departs from traditional learning-by-doing frameworks by focusing on subjective beliefs rather than productivity accumulation and by highlighting how effort choices are shaped more by prior circumstances than by technological complementarities.

Our approach to modeling distorted beliefs is related to the literature on *bandit problems*, in which agents face a trade-off between exploiting known options and exploring uncertain ones to learn about payoffs. As in a one-armed bandit, agents in our model can learn the return to effort only by exerting it, an action that is both costly and risky (see Piketty 1995 for an application in a political economy context). However, unlike classical formulations based on farsighted Bayesian decision-making (e.g., Bellman 1956) or minimax procedures (Robbins 1952), our agents follow a short-term decision protocol and do not fully internalize the informational value of exerting effort.⁵ This boundedly rational

⁵Our decision protocol is similar to that introduced by Mahalanobis (1940), often credited with the first formulation of the bandit problem. Moreover, unlike standard bandit models, our baseline framework does

behavior generates multiple long-term, welfare-ranked outcomes: some agents fall into low-effort pessimism, while others persist and succeed. This multiplicity contrasts with standard bandit models, which typically converge to a unique outcome. Our approach also differs from Bénabou and Tirole (2002, 2006), where distorted beliefs arise from selective memory or quasi-hyperbolic preferences, but agents are otherwise farsighted.

Third, we contribute to the literature on *discrimination and the persistence of inequality*. Structural inequality based on race, gender, caste, or class persists not only through differential treatment in markets and institutions but also by shaping expectations, beliefs, and behavior. Our framework introduces a novel behavioral channel through which disadvantage—such as that arising from discrimination—can sustain itself: by influencing the evolution of subjective beliefs about the returns to effort. We complement a growing body of work that emphasizes the psychological dimensions of inequality (e.g., Bohnet 2016; Bertrand and Mullainathan 2004; Hoff and Pandey 2004) In particular, our framework departs from approaches such as Loury (2019), which treat preferences and beliefs as fixed. In such settings, persistent inequality arises through self-selection, implying limited scope for belief-based interventions—such as grit training or exposure to role models—to alter long-term outcomes. In contrast, by allowing beliefs and effort to evolve endogenously, our model identifies when and how such interventions can be effective.

The remainder of the paper is structured as follows. Section 2 introduces the baseline model and summarizes its key extensions. Section 3 explores the policy implications of the framework and Section 4 concludes. Appendices A contain the proofs of all propositions, while Appendices B, C, D, and E provide detailed discussions of the model extensions.

2 Disadvantage, Beliefs and the Returns to Effort

In this section, we introduce a model that links structural disadvantage, measured by exogenous circumstances of an agent, to beliefs and effort exertion.

2.1 Model Set-up

Time is indexed by $t = 0, 1, \dots, T$, where $T > 0$ is a large but finite number. We consider an agent whose fixed external circumstance is denoted by $\omega \geq 0$. This circumstance encompasses a broad set of factors that are exogenous to the individual’s effort but shape the returns to exerting it. These include protected characteristics such as inherited wealth, gender, caste, or race, as well as contextual or structural factors such as access to social networks, neighborhood quality, education systems, or credit markets. A lower value of ω reflects a higher level of structural disadvantage. Although ω is fixed at the individual level, the disadvantage it represents can, in principle, be alleviated through policy interventions.⁶

not assume a prior over the true value of q , although we explore this extension in Appendix E.

⁶This assumption is supported by evidence showing that structural factors such as caste, gender or parental education are highly persistent and largely unaffected by short-term outcomes such as school

The agent begins with a prior subjective belief p_0 about the return to effort $e \in \{0, 1\}$ in undertaking a project or task, where the outcome is given by $\theta \in \{0, 2e\omega\}$. The return to effort is proportional to ω , and the proportionality factor is determined by the effort. The interpretation is that effort generates a lower return for an agent who is relatively more disadvantaged than another agent. This specification captures, in a parsimonious way, the effect of structural disadvantage on the return to effort.⁷

Although θ and ω could lie in the same domain, for example education, where ω could be interpreted as parental education and θ as child education outcomes, they need not necessarily do so. For example, ω may denote fixed social circumstances such as caste or gender, while θ may capture results such as academic performance in school. A key assumption we make is that there is no feedback from θ to ω ; the agent's circumstances are fixed and unaffected by individual effort choices. Continuing with the example above, we assume that there is no feedback from child education outcomes to parental education.

In each time period $t > 0$, the agent has a subjective distribution over the returns θ_t from exerting effort $e_t = 1$, given by:

$$\theta_t = \begin{cases} 2\omega & \text{with probability } p_{t-1} \\ 0 & \text{with probability } 1 - p_{t-1} \end{cases}$$

where p_{t-1} is the prior probability inherited from the preceding period. If $e_t = 0$, then, of course, the outcome is 0 with probability one.

Subjective beliefs about success conditional on exerting effort do not directly depend on the circumstances of an agent, e.g. two agents facing different structural conditions may begin life with identical beliefs. Such subjective beliefs are updated over time according to the following rule:

$$p_t(e_t) = \begin{cases} \frac{p_{t-1}+1}{2} & \text{if } e_t = 1 \text{ and } \theta_t = 2\omega \\ \frac{p_{t-1}}{2} & \text{if } e_t = 1 \text{ and } \theta_t = 0 \\ p_{t-1} & \text{if } e_t = 0 \end{cases}$$

That is, subjective beliefs prior to the event are updated into posterior beliefs based on the outcome achieved in that period. These updated beliefs serve as the prior when the agent chooses to exert effort in the subsequent period. Notice that belief updating occurs only when the agent exerts effort. Depending on the outcome, beliefs can increase or decrease. If the agent does not exert effort, the beliefs remain unchanged.

performance (see Hertz et al. 2008; Narayan et al. 2018). These characteristics function as fixed attributes at the individual level, even if they evolve slowly at the social level, which allows us to isolate the effects of belief formation and effort while still acknowledging that structural disadvantage can be addressed through policy.

⁷The particular functional form adopted here simplifies both the exposition and the assumptions underlying our results; In general, any function of effort and circumstance that is increasing and mutually complementary in both variables suffices.

We differentiate between subjective beliefs p_t and the objective probability of success $q \in [0, 1]$, which can be interpreted as an individual’s true talent or potential.⁸ When $e_t = 1$, the objective distribution of the outcome in each t for $0 \leq q \leq 1$ is as follows:

$$\theta_t = \begin{cases} 2\omega & \text{with probability } q \\ 0 & \text{with probability } 1 - q \end{cases}$$

We assume that q is **unknown** to the agent, consistent with empirical evidence showing widespread misperceptions of returns to effort.⁹ Moreover, agents **do not hold a prior** over q , nor do they assume that the true value of q lies within the support of any such prior. Instead, we treat p_0 as a primitive—a subjective belief about success conditional on effort. As agents observe the results of their own effort, they update p_t through a heuristic process that does not follow the Bayes’ rule. Beliefs evolve on the basis of the empirical frequency of success, much like in **nonparametric learning**. In this sense, the evolution of p_t resembles a kernel-based estimate of conditional probabilities, where the “sample” consists of the observed sequence of successes and failures conditional on exerting effort (see, e.g., Parzen 1962).¹⁰

This allows us to interpret p_t as the agent’s **locus of control**—whether they perceive success as dependent on their own actions (internal LoC) or as beyond their control (external LoC). Agents with a more internal LoC will tend to have higher p_t values. However, since p is not derived from a belief in q (i.e., the agent’s own ability or potential), it does not directly represent **self-efficacy** (Bandura 1977) —the belief in one’s own ability to succeed. In Section 2.5 and Appendix E, we relax this assumption and consider an extension in which p_t is derived from a prior distribution over q . In that Bayesian formulation, p_t naturally reflects self-efficacy and allows us to connect more directly to that psychological construct.

We further assume that q is identical between individuals and **independent of their circumstances** ω , capturing the idea that all agents have equal potential conditional on effort.¹¹ Since q does not depend directly on ω , any systematic link between the circumstance of an agent and the evolution of their subjective beliefs and effort choices cannot be a direct consequence of the level of disadvantage that is exogenous to the agent. Instead,

⁸For ease of interpretation, let $0 \leq \epsilon \leq 1$ be the probability that an agent makes an error conditional on exerting effort. In our model, $q = 1 - \epsilon$, so the inherent talent or potential is the probability that the agent will perform the task at hand without making systematic errors. We assume that ϵ does not depend on ω .

⁹For example, Jensen (2010) finds that schoolboys in the Dominican Republic underestimate the returns to secondary education. Talsma et al. (2019) show that students’ self-confidence often diverges from actual performance and predicts future outcomes. Nikolov and Jimi (2018) document large belief gaps in Tanzania, especially among disadvantaged groups, and Bobba and Frisncho (2022) find that feedback correcting misbeliefs in Mexico improves educational choices.

¹⁰What is important here is that when the agent exerts effort, conditional on success, we have $p_t > p_{t-1}$, while conditional on failure $p_t < p_{t+1}$ for each $t \geq 0$. Moreover, as we will show in Proposition 1 below, as $t \rightarrow \infty$, $\lim_{t \rightarrow \infty} E(p_t) = q$; that is, in the limit p_t is a consistent estimator of q . Our results are robust to any process of updating subjective beliefs p_t that satisfies these two properties.

¹¹In Section 2.5 and Appendix D we examine the case where q is allowed to depend on ω and in Section 3 we explore policies that directly affect q .

any such link, if it exists, must arise endogenously as beliefs and effort choices adapt over time to their circumstance.

The following proposition shows how subjective beliefs p_t and objective probability q are linked: in the limit, as the number of time periods becomes infinite, the expected value of p_t converges to q provided that effort is exerted in each period.

Proposition 1. *For every value of p_0 , $\lim_{t \rightarrow \infty} E(p_t) = q$ if $e_t = 1$ for all $t \geq 1$. In the special case, where $q = 1$ or $q = 0$, $\lim_{t \rightarrow \infty} p_t = q$.*

The underlying intuition of Proposition 1 is as follows. First, consider the special case where $q = 1$. In any period t , whenever the agent exerts effort, the updated belief satisfies $p_t > p_{t-1}$. The proposition shows that, in the limit, p_t monotonically converges to 1. In contrast, when $q = 0$, each effort instance leads to $p_t < p_{t-1}$, and p_t monotonically converges to 0 over time.¹²

In the general case where $0 < q < 1$, if the agent continues to exert effort in every period, the sequence of beliefs p_t converges in expectation to the true probability of success q . That is, over time, the influence of the initial belief p_0 vanishes, and the belief update process produces an asymptotically unbiased estimate of q . However, if the agent stops exerting effort after a finite number of periods, belief updating ceases. In that case, subjective belief p_t remains fixed and can retain bias, as it continues to reflect initial belief p_0 rather than the true probability of success q . Thus, sustained effort is necessary not only for success but also for accurate belief formation.

2.2 Short-term Decisions

We begin by specifying the instantaneous utility of the agent from the outcomes realized and the effort in the period t , given by:

$$u(\theta_t, e_t) = v(\theta_t) - c(e_t),$$

where $v(0) = 0$. We assume that $v(\cdot)$ is strictly increasing and continuous and that the cost function satisfies $c(0) = 0$ and $c(1) = c > 0$.¹³

In each $t \geq 1$, the individual solves a short-term decision problem where they choose to exert effort if and only if:

$$E(u|e_t = 1, p_{t-1}) = p_{t-1}v(2\omega) - c \geq 0 = E(u|e_t = 0, p_{t-1}).$$

The agent's expected utility preferences adapt indirectly through their previous effort choices, as the results achieved update their beliefs, which then influence future effort decisions. In a short-term decision problem, the agent internalizes the immediate consequences of exerting effort, that is, the return it generates within the current period. How-

¹²The assumption that effort is exerted in every period $t \geq 1$ can be relaxed to require that effort is exerted in a countably infinite number of periods.

¹³Note that the cost of effort does not depend directly on the extrinsic circumstance ω .

ever, we assume that the agent does not internalize how current effort choices influence the evolution of their future beliefs and, consequently, future effort decisions.¹⁴

This behavioral bias allows the model to generate multiple long-term outcomes, each consistent with individually rational short-term behavior, but differing in welfare. Crucially, all agents—regardless of their circumstances—share this bias. It resembles what educational psychologists term metacognitive ignorance (Flavell 1979): the failure to recognize or reflect on one’s own learning processes. In our context, this takes the form of not realizing that effort is necessary to improve beliefs about the likelihood of success. More specifically, Bandura (1977) highlights how people underestimate the role of effort in achieving mastery and updating beliefs, while Zimmerman (2000) shows that in self-regulated learning, people often do not anticipate how repeated effort can improve self-efficacy and future performance.

In a short-term decision problem, $e_t(p_{t-1}) = 1$ iff

$$p_{t-1} \geq \hat{k}(\omega) = \frac{c}{v(2\omega)}.$$

In other words, the individual will exert effort if and only if p_{t-1} is (weakly) greater than a threshold value $\hat{k}(\omega)$. As $v(\cdot)$ is an increasing function of θ_t , and hence an increasing function of ω , it follows that $\hat{k}(\omega)$ decreases in ω . Hence, an agent with a higher level of disadvantage - characterized by a lower value of ω — requires a higher level of prior beliefs than a less disadvantaged agent to choose to exert effort. A short-term decision is optimal if (a) $e_t(p_{t-1}) = 1$ and both $p_{t-1} \geq \hat{k}(\omega)$, $q \geq \hat{k}(\omega)$, or (b) $e_t(p_{t-1}) = 0$ and both $p_{t-1} < \hat{k}(\omega)$, $q < \hat{k}(\omega)$.¹⁵

Figure 1 illustrates the choice of effort by an agent with circumstance ω and belief p_{t-1} , assuming $0 < q < 1$. The figure features two key functions: the horizontal line in q , which does not depend on circumstances and represents the true probability of objective success; and the curve $\hat{k}(\omega)$, which represents the threshold belief required to exert effort and decreases with ω . At each level of ω , comparing q with $\hat{k}(\omega)$ tells us whether exerting effort is optimal: if $q \geq \hat{k}(\omega)$, effort is optimal; if $q < \hat{k}(\omega)$, effort is not optimal. The horizontal line at q also separates the space into areas where the agent’s belief p_{t-1} is below or above the true success probability q , helping to visualize the potential with or without confidence.

In the figure, regions **A** and **D** represent outcomes that are *optimal*: **A** corresponds to optimal low effort (when both p_{t-1} and q are not higher than $\hat{k}(\omega)$), for example, a student

¹⁴This is a particular instance of the notion of boundedly rational decision-making with adaptive preferences, specifically in our case, adaptive beliefs (see e.g., Von Weizsäcker 1971, Hammond 1976, Pollak 1978 and Dalton and Ghosal 2012)

¹⁵The assumption that effort is binary is made for simplicity of exposition of both the models and the results. If $e \in [0, 1]$, then, of course, both the objective probability of success and the subjective probability of success would be maps from $[0, 1]$ to itself. If both these maps, as well as the cost of effort, are linear in effort, then generically corner solutions will arise in the short-term decision problem we study. In that case, the analysis would be formally equivalent to our two-effort model.

with low innate ability and low confidence who rationally chooses not to study hard. The region **D** corresponds to an optimal high effort (when both p_{t-1} and q exceed $\hat{k}(\omega)$), such as a talented and confident student who invests effort and achieves strong results. Regions **B** and **C** represent *suboptimal choices*, where the agent's effort decision does not match what would maximize the welfare given q . In **B**, the agent exerts effort when it is not justified ($p_{t-1} \geq \hat{k}(\omega)$ but $q < \hat{k}(\omega)$), like a student who works very hard despite lacking innate ability and ultimately gains little. In **C**, the agent fails to exert effort despite having a high innate ability ($p_{t-1} < \hat{k}(\omega)$ but $q > \hat{k}(\omega)$), as when a capable student wrongly doubts their ability and does not invest in studying.

This partition highlights how differences in circumstances ω , which determine the value of the threshold $\hat{k}(\omega)$ through their effect on potential returns to effort $v(2\omega)$, interact with belief distortions to shape effort choices. When comparing two students, one more disadvantaged (with a lower ω) and one more advantaged (with a higher ω), the disadvantaged student faces a higher threshold $\hat{k}(\omega)$ and therefore requires stronger beliefs to justify exerting effort. As a result, they are more likely to end up choosing suboptimal low-effort—particularly in region **C**—even when their objective potential q is high.¹⁶

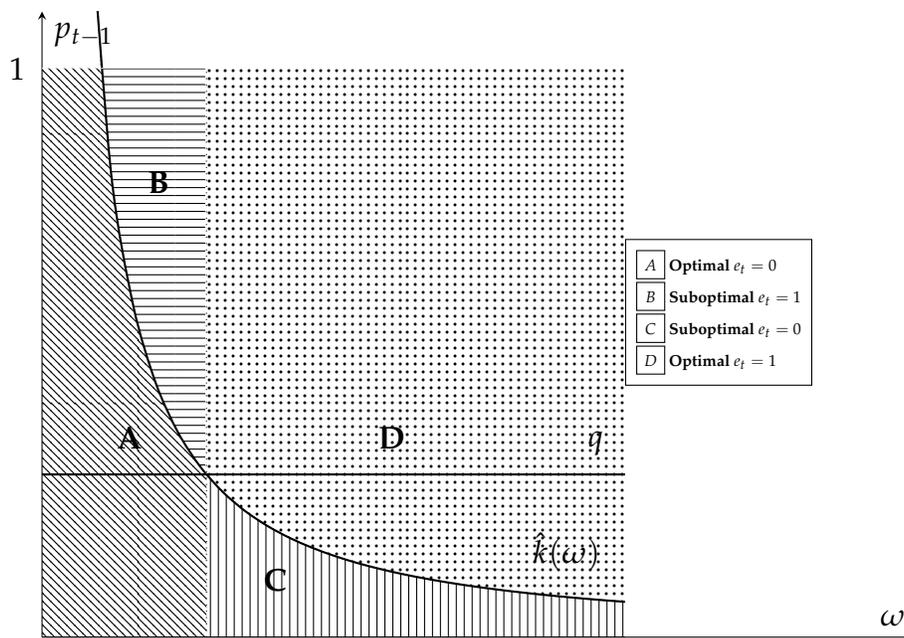


Figure 1: Short-run outcomes when $0 < q < 1$.

Note that if $v(\cdot)$ is unbounded above, then $\lim_{\omega \rightarrow \infty} v(2\omega) = \infty$ and $\lim_{\omega \rightarrow \infty} \hat{k}(\omega) = 0$. Since $\hat{k}(\omega)$ is decreasing in ω , at very high levels of advantage, we have $\hat{k}(\omega) \approx 0$. Thus, unless the agent is extremely pessimistic (that is, $p_{t-1} < \epsilon$ for some ϵ close to

¹⁶Recall that the threshold $\hat{k}(\omega)$ at each ω is determined by the ratio of effort cost to the potential return, $\hat{k}(\omega) = c/v(2\omega)$. Policies or interventions that reduce c (e.g. subsidies, reduced barriers) or increase $v(2\omega)$ (e.g., higher returns to success) can shift the threshold downward, making the effort more attractive even for agents with pessimistic beliefs. We discuss these policies in detail in Section 3.

zero), the highly advantaged agent will exert effort at each $t \leq T$. In other words, when the advantage is large, the marginal cost of effort becomes negligible compared to the potential gain, making effort worthwhile unless the agent strongly doubts their chances of success. At the other extreme, it is possible that $v(2\omega) < c$, implying that $\hat{k}(\omega) > 1$, in which case the agent will exert effort at each $t \leq T$. When $0 < \hat{k}(\omega) < 1$, the decision to exert effort depends on the agent's level of confidence, that is, whether or not $p_{t-1} \geq \hat{k}(\omega)$.

In general, for a given level of prior belief p_{t-1} , agents' effort choices vary systematically with their circumstances. Intuitively, an agent faced with high effort costs or very limited potential gains (i.e., low ω) will require a high belief in their likelihood of success to justify the effort. In contrast, when the returns to effort are potentially high (i.e., large ω), even relatively pessimistic agents may find it worthwhile to try. For example, consider two students who are equally uncertain about their abilities. One attends a school with poor resources (low ω), while the other benefits from strong institutional support (high ω). Although both students are identical in all respects, only the student in the more advantaged environment may find it worthwhile to study, simply because the returns to effort are greater in their context.

2.3 Long-term Outcomes

To characterize the predictions and welfare implications of our framework, in this subsection we study the long-term outcomes of the short-term decisions described in the preceding section.

Given ω , let the pair (e_t, p_t) , $t > 0$ denote the outcome of a series of short-term decisions. A long-term outcome is a pair (e_T, p_T) such that $\lim_{t \rightarrow T} e_t(p_{t-1}) = e_T$ and $\lim_{t \rightarrow T} p_t(e_t) = p_T$. Intuitively, long-term outcomes reflect how effort and beliefs co-evolve through repeated short-term choices.

Long-term outcomes may depend on the value of T . Therefore, when characterizing long-term results, we do so for all $T > 0$ or in the limit as $T \rightarrow \infty$. For example, consider a student who decides each year whether to study hard. The student does not internalize how today's effort might improve their future confidence or learning. However, from an external perspective (e.g., a planner or analyst looking at the outcomes over time), the long-term result depends on the time horizon T . If T is short (e.g., focusing only on immediate exam performance this year), the student's beliefs might be sustained throughout, as no discouragement might occur during one year. On the other hand, if T is long (e.g., considering the entire schooling career), it is more likely that the student will incur in a short sequence of setbacks that will eventually dissuade them from studying, possibly resulting in a drop off.

In addition, the specific values of the innate talent, q , shape the dynamics and stability of the long-term outcomes. In extreme cases, when $q = 1$ (always succeed with effort) or $q = 0$ (never succeed with effort), the sequence of short-term decisions becomes deterministic. In these cases, the question of local stability reduces to whether there exist sequences

of short-term decisions that converge to a given long-term outcome. However, when $0 < q < 1$, the sequence of short-term decisions becomes stochastic over the state space $(e_t, p_t) \in \{0, 1\} \times [0, 1]$. In this case, the relevant notion of stability is whether a long-term outcome constitutes an absorbing state, i.e. whether there exist sequences that converge to it and, once reached, the process remains there with probability one, with zero probability of transitioning to another state.

The next proposition formally characterizes long-term outcomes, their welfare, and their stability properties.

Proposition 2. *Given ω , the long-term outcomes, their welfare ranking, and stability are as follows.*

1. **Long-term outcomes:** *For all $T > 0$, there are two long-term outcomes, that is, (i) High effort ($e_T = 1$) and optimistic beliefs $p_T \geq \hat{k}(\omega)$, and (ii) Low effort ($e_T = 0$) and pessimistic beliefs $p_T < \hat{k}(\omega)$;*
2. **Welfare ranking:** *For all $T > 0$, if $q > \hat{k}(\omega)$, the first long-term outcome with $e_T = 1, p_T > \hat{k}(\omega)$ welfare dominates the second long-term outcome with $e_T = 0, p_T < \hat{k}(\omega)$, that is, $E(u|e_T = 1, q) > E(u|e_T = 0, q)$.*
3. **Stability:**
 - *When $q = 1$, if $p_0 \geq \hat{k}(\omega)$, the sequence converges to the first long-term outcome (high effort, optimistic beliefs); if $p_0 < \hat{k}(\omega)$, it converges to the second outcome (low effort, pessimistic beliefs).*
 - *When $q = 0$, the sequence converges to the second long-term outcome (low effort, pessimistic beliefs).*
 - *When $0 < q < 1$, the second long-term outcome (low effort, pessimistic beliefs) is the unique absorbing state: $\lim_{T \rightarrow \infty} e_T = 0, \lim_{T \rightarrow \infty} p_T < \hat{k}(\omega)$.*

The first part of the proposition shows that only two types of long-term outcomes are consistent with the dynamics of short-term decision problems: either the agent ultimately settles into exerting effort, supported by sufficiently optimistic beliefs, or the agent ultimately stops exerting effort, sustained by sufficiently pessimistic beliefs. The intuition is that effort and beliefs reinforce each other over time: exerting effort increases the (perceived) likelihood of success, which sustains or boosts optimistic beliefs, encouraging further effort. In contrast, not exerting effort leads to no success, which reinforces pessimism and further inaction. Intermediate or oscillating patterns are not sustained because any deviation either converges toward persistent effort with high beliefs or collapses into no effort with low beliefs. Thus, the dynamics converges to one of these two outcomes. In terms of Figure 1 above, and replacing p_{t-1} with p_T and e_t with e_T , long-term results will be located only in areas **D** (optimistic high-effort outcome) and **C** (pessimistic low-effort outcome).

The second part of the proposition shows that the two long-term outcomes can be ranked on welfare. The key comparison is between the expected utility of the agent, evaluated

using the *objective* success probability q , under each outcome. When $q > \hat{k}(\omega)$, the optimistic high-effort outcome yields the expected utility higher than the pessimistic, no-effort one; the reverse holds when $q < \hat{k}(\omega)$. Hence, any long-term outcome located in area C is long-term dominated by an outcome located in area D.

The third part of the proposition concerns the stability of the two long-term outcomes. Here, there are three cases to consider. First, when $q = 1$, that is, effort always results in success. If the initial belief satisfies $p_0 \geq \hat{k}(\omega)$, the agent has an incentive to exert effort, and the sequence of short-term decisions converges to the long-term optimistic result with high effort. In contrast, if $p_0 < \hat{k}(\omega)$, the agent never exerts effort, the beliefs remain unchanged ($p_T = p_0$), and the result is a persistent low-effort pessimistic outcome.

Since $\hat{k}(\omega)$ is decreasing in ω , agents facing greater disadvantage are more likely to fall into the welfare-dominated outcome. For example, a highly capable student ($q = 1$) will succeed whenever they study. If they start to feel confident enough ($p_0 \geq \hat{k}(\omega)$), success reinforces effort and optimism. But if they start too pessimistic, they will not try, they will not succeed, and they will remain stuck in low effort despite their innate ability. Figure 2 (left panel) illustrates the stability result when $q = 1$, with p_t on the x-axis and p_{t+1} on the y-axis. As long as $p_0 \geq \hat{k}(\omega)$, p_t converges to 1; when $p_0 < \hat{k}(\omega)$, $p_t = p_0$ remains constant. Thus, $[\hat{k}(\omega), 1]$ is the attraction basin for the long-term result with $e_T = 1, p_T = 1$, while $[0, \hat{k}(\omega))$ is the attraction basin for the long-term result with $e_T = 0, p_T = p_0$. Note that the map from p_t to p_{t+1} has an S-shaped shape, typical of models that generate poverty traps; however, here it is the pessimistic low-effort trap that is the primary object of interest.

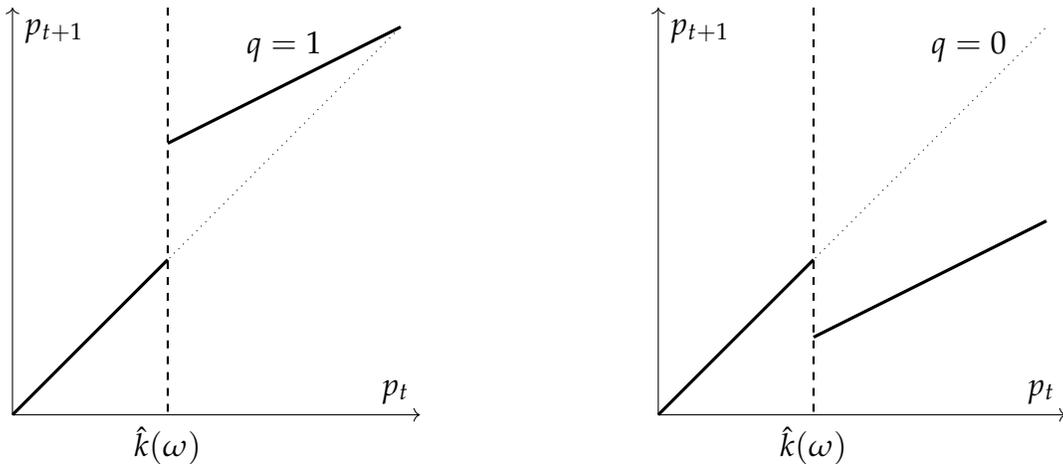


Figure 2: Belief dynamics and stability for $q = 1$ (left) and $q = 0$ (right).

Second, when $q = 0$, effort never leads to success. Even if the initial belief satisfies $p_0 \geq \hat{k}(\omega)$, there exists some period t' such that $p_{t'} < \hat{k}(\omega)$, after which the agent stops exerting effort. Consequently, the agent converges to the pessimistic long-term outcome with little effort. In this case, the convergence to the pessimistic result occurs regardless of the agent's circumstance ω , as exerting effort does not produce a return. Continuing with the example of the student, if they initially believe that they are capable and study

hard, but their true ability is zero ($q = 0$), the effort never results in success. Over time, repeated failure leads them to revise their beliefs downward and ultimately give up entirely, regardless of their initial circumstance. Figure 2 (right panel) illustrates the stability result when $q = 0$. The key difference between $q = 0$ and $q = 1$ is that the map from p_t to p_{t+1} is Z-shaped, implying that there is only one long-term stable global outcome: $e_T = 0, p_T < \hat{k}(\omega)$.

When $0 < q < 1$, there are several cases to consider. In extremes, an agent with a very high level of advantage, for whom $p_0 \geq \hat{k}(\omega) \approx 0$, will always exert effort and converge to an optimistic long-term result of high-effort. In contrast, at extreme levels of disadvantage, where $\hat{k}(\omega) > 1$, the agent will never find it optimal to exert effort and instead converge to a long-term pessimistic low-effort result. Figure 3 illustrates the long-term results for extreme advantage (left) and extreme disadvantage (right).

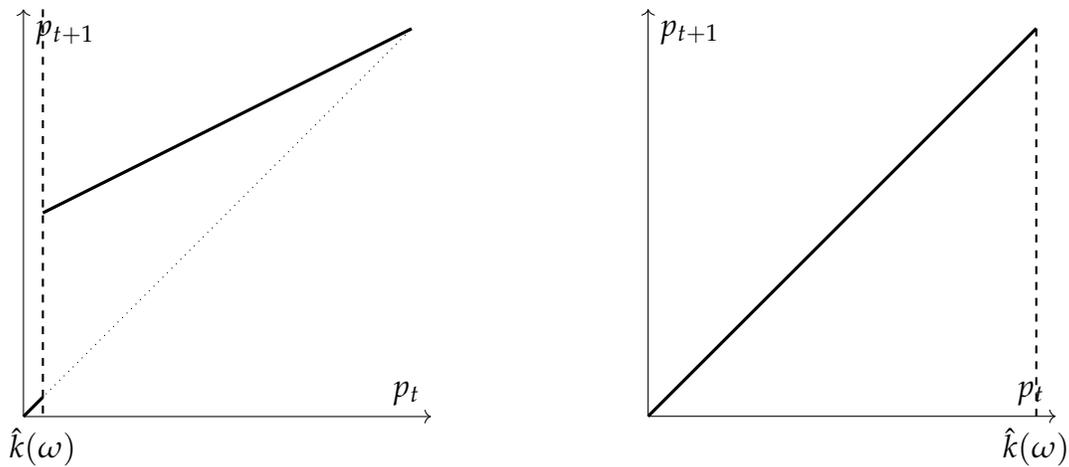


Figure 3: Belief dynamics when $0 < q < 1$, for extreme advantage (left) and extreme disadvantage (right).

In both cases, the agent converges to the welfare-dominant outcome, given their underlying potential $0 < q < 1$. For example, a highly talented student (q close to 1) with strong support (ω high) will exert effort and succeed. Meanwhile, a severely disadvantaged student (ω low), regardless of their belief in success and talent, will not try, and this inaction is consistent with their true chances of succeeding.¹⁷

However, for intermediate levels of disadvantage, when $0 < \hat{k}(\omega) < 1$, the pessimistic low-effort outcome is the only absorbing state, even if it is welfare-dominated. If at any time the agent's belief falls below the threshold $\hat{k}(\omega)$, they stop exerting effort, and this decision becomes self-reinforcing: without effort, the beliefs stagnate or decline, and the agent never relinquishes.

In contrast, convergence to the optimistic, high-effort long-term outcome requires the

¹⁷We can also use Figure 1 to illustrate the stability result when $0 < q < 1$. Areas B and C are the only two absorbing states. When $q > \hat{k}(\omega)$, then C, with the long-term welfare-dominated outcome, is the absorbing state.

agent to exert effort in every period, which in turn requires their belief to remain above the threshold indefinitely. But when $0 < q < 1$, belief updating is stochastic and with positive probability $1 - q$, a failed effort leads to $p_{t+1} < p_t$. Thus, there is always a chance that beliefs fall below the threshold, triggering permanent disengagement. As a result, the high-effort outcome cannot be an absorbing state.

Continuing with the above example, a student with moderate potential ($0 < q < 1$) and an intermediate level of disadvantage may be confident enough to study. But if early efforts fail, something that can happen by chance, their belief drops below the threshold, and they stop trying altogether. Since effort ceases, beliefs no longer recover, and the student becomes trapped in a self-fulfilling cycle of low effort and pessimism.

2.4 Pessimistic Low-effort Traps

A pessimistic low-effort trap is a scenario in which an agent acts suboptimally by not exerting effort. This corresponds precisely to region **C** in Figure 1, where the agent does not exert effort despite having a high objective probability of success ($q > \hat{k}(\omega)$). Since this region induces low effort, it is also an absorbing state and therefore a long term outcome.

Formally, the trap is defined as follows:

Definition 1. *Given ω and $T > 0$, a pessimistic low-effort trap is a long-term outcome (e_T, p_T) such that $e_T = 0$ and $p_T < \hat{k}(\omega) < q$.*

A pessimistic low-effort trap is closely related to what the psychological literature refers to as **learned helplessness**, which describes a state in which individuals stop trying because they believe that effort has no effect, often as a result of repeated failures (Seligman 1972). The beliefs about the probability of success conditional on effort can only be updated if the agent constantly exerts effort, which, in turn, requires the beliefs to be optimistic enough in the first place. However, the agent's subjective belief only matters when $0 < \hat{k}(\omega) < 1$: if $\hat{k}(\omega) \approx 0$, the agent will always exert effort, and if $\hat{k}(\omega) \geq 1$, the agent will never exert effort although in this case the absence of effort is welfare dominated.

The following result demonstrates the link between structural disadvantage and learned helplessness:

Proposition 3. *The relationship between disadvantage and pessimism is as follows. Consider two agents i, j with intermediate levels of disadvantage where agent i is more disadvantaged than agent j , that is, $\omega_i < \omega_j$ and a common value of initial beliefs $p_{0,j} = p_{0,i}$. In a long-term outcome, for all $T > 0$, the likelihood that agent i is stuck in a pessimistic low-effort trap is greater than that of agent j .*

Consider two agents such that the first agent has a relatively higher level of advantage. Assume that both agents have the same prior beliefs about the probability of success conditional on effort p_0 . To begin with, assume that $q = 1$. As $\hat{k}(\omega)$ decreases in ω , whenever $\hat{k}(\omega_j) < p_0 < \hat{k}(\omega_i)$, agent i will not exert effort, but agent j will. Hence, if p_0 is chosen according to a common probability distribution among all agents, the higher

the level of disadvantage, the higher the probability that the agent will converge to the pessimistic trap of low effort.

To gain intuition when $0 < q < 1$, suppose that both agents initially exert effort, but the return obtained is zero. In the next period, both agents will be more pessimistic. However, such pessimism will have a more detrimental effect on the relatively disadvantaged agent. This is because the threshold level of beliefs below which the relatively disadvantaged agent will choose not to exert effort is higher than the threshold level of beliefs below which the other agent will choose not to exert effort. Hence, over time, the likelihood that the disadvantaged agent will be stuck in low effort, pessimistic trap will be higher. Thus, the higher the level of disadvantage, the higher the likelihood that an agent will fail to believe in their probability of success, leading them to underinvest in effort even when their innate ability is high.

2.5 Extensions

In this section, we discuss four extensions of the baseline model, which are further expanded in the appendix.

First, in Appendix B, we examine the case where agents are **farsighted** and internalize the impact of the current effort on future outcomes. We consider various degrees of partial farsightedness, where agents anticipate how future beliefs and effort choices depend on current decisions. We show that partial farsightedness, by itself, is insufficient to fully restore the incentives to exert effort: without simultaneously raising beliefs, the agent will not exert effort.

When we move to the case of full farsightedness, where agents anticipate the entire stream of future utilities, we find that a key feature of the short-term protocol is preserved: once a decision-maker switches to a pessimistic, low-effort path, they remain on that path in all subsequent periods. Even fully forward-looking agents follow a threshold strategy for effort, which means that they exert no effort in all periods if their initial belief is too low. We state a proposition that highlights the role of the degree of farsightedness. On the one hand, agents who internalize more future periods are more likely to exert effort, as the threshold belief required is lower. However, agents with a lower discount factor (i.e., more impatience) are less likely to invest effort. A direct implication is that early interventions to raise initial beliefs can generate larger welfare gains by nudging agents onto a high-effort path before pessimistic traps take hold.

Second, in Appendix C, we extend the baseline model by explicitly incorporating **aspirations** as an additional psychological driver of effort. As in Dalton et al. (2016), agents do not internalize how their effort affects the evolution of aspirations, mirroring their bias in belief updating. Aspirations may be exogenous or shaped by realized outcomes. We show that aspirations interact dynamically with beliefs and effort: higher aspirations lower the threshold belief needed to justify exerting effort, thereby encouraging persistence. However, when aspirations evolve slowly relative to beliefs, frustration can arise if high aspirations are repeatedly unmet. Importantly, we identify two distinct traps: **frustra-**

tion, where aspirations remain high but unmet and effort eventually stops, and **fatalism**, where aspirations collapse alongside beliefs and effort.

Third, in Appendix D, we consider a different implication of structural disadvantage: rather than lowering the returns to effort, it reduces the **objective probability of success**, q . In this case, the disadvantage acts as a barrier to success, while the rewards from success remain unchanged. We show that our main results are robust to this formulation. When disadvantage directly lowers the likelihood of success, it increases the frequency of failure, causing beliefs to be revised downward more often. As a result, disadvantaged agents are more likely to fall below the belief threshold required to justify the effort. Since the belief updating process is the same, the structure of the long term outcomes mirrors the baseline model: for intermediate levels of disadvantage (i.e., ones for which the corresponding probability of success is neither 0 or 1) the only stable long term outcome is the pessimistic, low effort trap.

Finally, in Appendix E we consider the case where the agent begins with a **prior** distribution over q , which we interpret as **self-efficacy beliefs**. Suppose $q \in \{q_L, q_H\}$ with $0 \leq q_L < q_H \leq 1$, and let P_0 denote the prior $q = q_H$. Beliefs are updated via **Bayes' rule** whenever effort is exerted. We show that, as in the baseline model, multiple long-term welfare-ranked outcomes can arise. For intuition, consider the special cases $q_L = 0$ and $q_H = 1$. If the agent exerts effort, the beliefs jump to 0 or 1, depending on whether the outcome is a failure or a success. However, if the initial self-efficacy P_0 lies below the same threshold \hat{k} identified in the baseline model,¹⁸ the agent never finds it worthwhile to exert effort and becomes trapped in a welfare-dominated low-effort outcome. More generally, when $0 < q_L < q_H < 1$, Bayesian learning recovers the true q only if the agent continuously exerts effort; when $q_L < \hat{k} < q_H$, even high-ability individuals ($q = q_H$) may stop exerting effort prematurely and fall into a pessimistic trap. Appendix E connects our analysis to the literature on Bayesian decision-making in one-armed bandit problems when agents are farsighted. A key result (Lemma 4.1 in Bradt 1956) in this literature is that once the known arm of a bandit is chosen by a farsighted agent, it continues to be chosen in all subsequent time periods, a feature consistent with our short-term decision protocol. When agents are farsighted, outcomes cannot be welfare-ranked as they can when they use our short-term decision protocol. However, it is still possible for the agent to be stuck with the overly pessimistic beliefs, and hence choose a welfare dominated outcome.

Extending the model in this way reinforces the central link between *disadvantage and beliefs*: structural disadvantage can erode self-efficacy beliefs (low P_0), increasing the probability of a pessimistic low-effort trap even when the objective potential is high. This is consistent with a large empirical literature documenting systematic gaps in self-assessed ability between socioeconomic groups. For example, Guyon and Huillery (2020) show that low-SES students in France underestimate their academic ability relative to equally performing high-SES peers. Gender gaps in confidence appear even to maintain the performance constant: Barber and Odean (2001) and Buser et al. (2014) document lower self-beliefs among women in math, science, and financial decisions. Similar patterns arise

¹⁸Appendix E shows that the relevant threshold for P_0 coincides with the baseline threshold.

in caste contexts: Hoff and Pandey (2006) find that making caste prominent depresses performance for low-caste students. Related work shows that women systematically underestimate their abilities and avoid competition (Niederle and Vesterlund 2007; Bordalo et al. 2019; Coffman 2014). Together, these findings highlight how structural disadvantage shapes self-beliefs and thereby increases the risk of falling into pessimistic low-effort traps. Moreover, even forward-looking individuals can choose zero effort in all periods if $P_0 < \frac{c}{v(2\omega)}$, the condition under which the agent never exerts effort under myopia. So, if $q > \frac{c}{v(2\omega)}$, under full farsightedness, in the unique outcome, the agent may choose a welfare dominated outcome because they get stuck with the wrong beliefs. The earlier the initial beliefs are changed, the greater the welfare gains.

3 Policy Implications

In this section, we study policies and interventions that can reduce the risk of falling into a pessimistic, low-effort trap or facilitate escaping from such a trap. We discuss the effectiveness of interventions that boost grit, increase aspirations and subjective beliefs such as locus of control or self-efficacy, and reduce structural disadvantages. In addition, we analyze when each type of policy may work in isolation, when they must be combined to be effective, and when they may fail to achieve the desired outcomes.

For policy analysis, we assume that a benevolent policy-maker has more information than the decision-maker about the underlying ability of the agent q and external circumstances ω , so can assess whether $q > \hat{k}(\omega)$ is or not. This means that the policy-maker knows whether the agent would benefit from exerting effort, given their true potential and structural disadvantage. This information asymmetry is plausible in many real-world settings. For example, a coach or teacher with years of experience may have observed a wide range of individuals with varying levels of innate ability and background constraints. Drawing on this experience, they can form a reliable estimate of the objective probability of success of a student q , and also a reasonable assessment of the structural disadvantage of the student ω , for example, by observing access to resources, parental involvement, or previous opportunities. This allows them to infer whether the true success potential of the agent exceeds the belief threshold $\hat{k}(\omega)$, even if the agent itself lacks that knowledge. Similarly, a program evaluator or mentor might have access to richer data—e.g., results from a pilot program or longitudinal cohort—which enables them to form more accurate judgments about who is likely to benefit from sustained effort.

3.1 Boosting Beliefs: Locus of Control, Self-Efficacy and Grit

Our model highlights the central role of subjective beliefs about the probability of success conditional on effort. In the baseline model, p_t captures an individual’s perceived control over outcomes, a notion closely related to *locus of control*. In the Bayesian extension introduced in Appendix E, where P_t is derived from a prior over an unknown ability parameter q , beliefs take on a more natural interpretation as *self-efficacy*. In both cases, interventions that raise beliefs or prevent them from falling below the belief threshold $\hat{k}(\omega)$

can promote sustained effort.

Informational interventions offer a direct route to the formation of beliefs. When credible and contextualized, information on the returns to effort, such as those studied in Ersoy (2023) and Rury and Carrell (2023), can shift beliefs upward. If these interventions succeed in raising p_0 above $\hat{k}(\omega)$ and maintaining this margin over time, they ensure that the agents will persist in exerting effort, thus improving long-term outcomes provided that $q > \hat{k}(\omega)$.

Another way to shift subjective beliefs is through exposure to **role models**—individuals from similar backgrounds who have succeeded through effort. These figures can signal that success is achievable despite shared structural constraints, prompting upward belief revisions. For example, Nguyen (2008) shows that in rural Madagascar, students and parents significantly underestimated the returns on education. When some villages received visits from local adults who had succeeded in education, the students updated their beliefs, increased school attendance, and achieved higher test scores. This evidence suggests that role models can change both perceived returns and effort decisions. Similarly, LaFerrara et al. (2012) show that exposure to aspirational characters in Brazilian soap operas influenced viewers' behavior, including fertility and divorce decisions, suggesting the broad potential of relatable success narratives to shape expectations. In a context in which aspirations matter, Di (2025) finds that exposure to role models with a similar background increases aspirations, beliefs, and effort among middle school students in China. Moreover, the paper documents complementarity between outcomes and aspirations and shows that frustration may arise when increases in aspirations are not matched by corresponding increases in confidence about achieving them.

Importantly, role model exposure may not occur naturally—especially if the role model bears costs from mentoring or visibility—so policy interventions may be required to facilitate these connections. At the same time, not all exposure is informative: observing someone succeed can only productively shift beliefs if the agent correctly infers that their own ability q exceeds their threshold $\hat{k}(\omega)$. Otherwise, the signal can be misleading, raising unrealistic expectations or leading to eventual disillusionment. This underscores the value of targeted contextualized interventions—ideally implemented by an informed and benevolent planner—to ensure that social learning produces beneficial belief updates.

Formally, this is how a role model can be effective. Let the circumstance of the role model be indicated by $\omega_k \geq \omega$, with $\omega_k - \omega < \epsilon$ for some small $\epsilon > 0$, ensuring that the role model is sufficiently similar to the agent. We assume that the role model communicates their successful experience of exerting effort to the agent, which the agent incorporates into their belief updating process. Under this mechanism, the agent's beliefs evolve taking into account external success evidence. As a result, the agent no longer relies solely on their own experience but also draws on external evidence from the role model's success. It is *as if* the agent is updating beliefs based on two signals: their own outcome and an additional positive signal from the role model. Formally, belief updating proceeds as

follows:

$$p_t = \begin{cases} \frac{p_{t-1}+2}{3} & \text{if } e_t = 1 \text{ and } \theta_t = 2\omega \\ \frac{p_{t-1}+1}{3} & \text{if } e_t = 1 \text{ and } \theta_t = 0 \\ \frac{p_{t-1}+1}{2} & \text{if } e_t = 0 \end{cases}$$

By construction, beliefs are eventually bounded below, since $p_t \geq \frac{1}{2}$, for all $t > 1$. To see this, take the extreme case where $p_0 = 0$: then, $p_1 = \frac{1}{2}$, as exposure to a role model increases beliefs. From the following period onward, it must be $p_t \geq \frac{1}{2}$. Therefore, if the belief threshold satisfies $\hat{k}(\omega) < \frac{1}{2}$, role model interventions will ensure sustained high effort over time.¹⁹

It is important to distinguish these policy interventions designed by a benevolent and informed planner, from informal signals that agents can encounter naturally (e.g., via social media, peers, or community networks). Individuals might observe others who succeed through effort and infer that exerting effort is worthwhile. However, such observational learning can be misleading because the agent typically lacks two key pieces of information: (i) whether the observed individual’s true success potential q exceeds their own belief threshold $\hat{k}(\omega)$, and (ii) whether their *own* q exceeds $\hat{k}(\omega)$. Without this knowledge, observed success could lead to unjustified optimism, followed by disillusionment if a similar effort does not produce success. This highlights the value of informed and beneficial interventions: by tailoring signals or feedback to the underlying potential and structural context of the agent, such interventions avoid the risk of false inferences and can reliably shift behavior toward welfare-enhancing paths.

Psychological interventions aimed at **increasing self-efficacy**, such as cognitive behavioral therapy (CBT) or mindset workshops, offer another route to increase self-beliefs. These interventions work by directly increasing the agent’s belief in their own ability to succeed if effort is exerted. For example, McKelway (2025) shows that a light-touch workshop designed to enhance self-beliefs among women led to persistent increases in economic outcomes. The treated participants displayed an S-shaped pattern in the outcomes over time which, in the lens of our model, is consistent with a setting where $q = 1$.²⁰

Similarly, Kaur (2025) studies an intervention that targets teachers’ perceived behavioral control, the belief that effort can influence student outcomes. This treatment led to persistent increases in teacher effort and improved student performance. In our framework,

¹⁹Note that the role model enters the belief-updating rule implicitly through an additional positive signal. Relative to the baseline updating rule, the numerator is inflated by one unit to reflect the role model’s successful experience, while the denominator is adjusted to reflect the larger set of signals being averaged. For instance, following success under effort, beliefs update as $p_t = (p_{t-1} + 2)/3$ rather than $(p_{t-1} + 1)/2$: one unit reflects the agent’s own success, and the additional unit captures the role model’s success. Likewise, following failure or no effort, beliefs still increase because the positive role-model signal is assumed to be continuously present. Similarity between the agent and the role model (i.e., $\omega_k \approx \omega$) ensures that this external signal is perceived as relevant and therefore fully incorporated. We abstract from explicitly modeling the role model’s characteristics or actions and instead represent their influence in reduced form as a persistent positive belief signal.

²⁰McKelway (2025) has a theoretical model which follows Benabou and Tirole (2002), leading to a different interpretation of this result.

such interventions increase the perceived return to effort of teachers p_t , which in turn leads to sustained investment and improved long-term results for students. These examples highlight how low-cost psychological nudges can sustainably increase perceived efficacy, particularly when q is high and the main constraint is pessimistic beliefs.

Interventions that strengthen grit. When $q = 1$, raising p_0 above $\hat{k}(\omega)$ is sufficient to guarantee convergence to the optimistic high-effort outcome. However, when $0 < q < 1$, the update of beliefs becomes stochastic and a one-time increase in p_0 is not enough. In such cases, *grit-enhancing interventions* that support persistence even in the face of setbacks are essential to keep beliefs p_t above the threshold $\hat{k}(\omega)$ over time. As quoted by Banerjee et al. (2019, p. 970), “life in 1.90 a day is emotionally draining, laced with perennial frustrations and the occasional crushing setback. Why then do we make the sacrifices needed to earn a few dollars each day and save for an abstract and uncertain better life decades away? To do so would take remarkable grit.” A key insight from their review is that interventions for the extreme poor may need to explicitly foster this kind of psychological resilience to generate sustained economic improvements.

An approach that explicitly targets grit is the “graduation model,” which combines cash, productive assets, life-skills training, and intensive mentoring to gradually build both economic capabilities and self-confidence. By supporting individuals closely over time, these programs help beneficiaries build a sense of control and persistence—reinforcing beliefs even in the face of temporary failure.

Another way to strengthen grit is by designing targeted psychological interventions. For example, Alan et al. (2019) develop and evaluate a structured training program that teaches primary-school children to persevere through difficulty and learn from failure. They find that treated children exhibit significantly greater consistency in effort, with long-term effects that suggest that beliefs remain bounded away from the pessimistic threshold even after setbacks. This kind of belief reinforcement—whether through mentoring, structured training, or psychological exercises—can help ensure that the effort persists in the presence of uncertainty.

Similarly, Campos et al. (2017) show that teaching “personal initiative” to microentrepreneurs in Togo increased firm profits. Their intervention emphasized proactive behavior and persistence in overcoming obstacles, traits closely aligned with grit. A key mechanism, they argue, is the ability to “better overcome setbacks” (p. 3), which reinforces the belief in continued success.

In a related domain, Bettinger et al. (2018) study perseverance in math learning. They find that a light-touch intervention that increases students’ willingness to stick to difficult problems significantly improves test performance. This shows that even small nudges to grit can help students stay engaged when challenges might otherwise lead to disengagement.

Interventions that foster a *growth mindset*—the belief that ability is malleable—also contribute to grit. Yeager et al. (2020) emphasize that the success of such interventions is critically dependent on the school environment. They find that growth mindset interven-

tions are most effective when peer norms support student’ persistence through difficulty or when schools provide formal resources that make effort more likely to be rewarded. In our framework, these environments help keep beliefs p_t above $\hat{k}(\omega)$, supporting sustained effort and long-term welfare improvements.

Finally, Ashraf et al. (2024) study an intervention that uses episodic future thinking to help economically vulnerable individuals better imagine how current actions can lead to future benefits. Participants who participated in guided visualization exercises exhibited improved economic behaviors and greater willingness to invest in their futures. This intervention can be understood as correcting distortions in memory or mental simulation (what the authors call “mental experiencing”) that otherwise make it harder for agents to perceive the link between effort and success. In our model, this can be interpreted as helping agents avoid the collapse of beliefs p_t after failure, thus preventing a fall below $\hat{k}(\omega)$. Alternatively, one might interpret the intervention as increasing farsightedness. However, as we show in Appendix B, farsightedness itself (interpreted as an increase in the discount factor δ) might not increase effort and therefore welfare if not accompanied by an increase in beliefs. Looking ahead without optimism will not lead to greater engagement. Hence, the welfare gains observed in Ashraf et al. (2024) likely stem from helping agents maintain sufficiently high beliefs, rather than inducing farsightedness on its own.

Crucially, the strength required for such interventions depends on the disadvantage parameter ω . More disadvantaged agents (i.e., those with a higher $\hat{k}(\omega)$) need stronger or more persistent belief reinforcement to escape pessimistic traps. Thus, belief-based interventions must be carefully tailored to the context. In addition, these interventions are most relevant for agents located in Region C of Figure 1, where beliefs are marginal and outcomes are sensitive to small shocks. In Region A, where $q < \hat{k}(\omega)$, even optimistic beliefs do not justify effort. In those cases, structural changes to returns or costs are required. In Region D, agents already exert effort and maintain high beliefs, making belief-boosting interventions redundant. Recognizing the region-specific effectiveness of these interventions is essential for appropriately targeting policy.

3.2 Shifting $\hat{k}(\omega)$ Downward

A disadvantaged agent has a higher probability of becoming stuck in a welfare-dominated, pessimistic, low-effort trap compared to a more advantaged agent. For agents with a very high degree of advantage, such that $\hat{k}(\omega) \approx 0$ - no intervention, or only minimal intervention, is needed to sustain effort. In this sense, belief- and grit-enhancing policies are largely redundant for the most advantaged.

In contrast, for severely disadvantaged agents with $\hat{k}(\omega) \geq 1$ and $q < 1$, belief-boosting interventions alone are insufficient. In such cases, material interventions that directly address the disadvantageous circumstances are necessary to shift the threshold $\hat{k}(\omega)$ downward. We distinguish two types of policies that reduce $\hat{k}(\omega)$: policies that reduce the effective cost of effort c and policies that raise the value of success, thus improving the returns to effort $v(2\omega)$.

Reducing the cost of effort c . Policies that lower the effective cost of exerting effort make it more feasible for agents to persist, even when beliefs are fragile. For example, **unconditional cash transfers** can ease liquidity constraints, allowing students to afford school materials, cover transportation, or reduce working hours thereby making it easier to invest time and energy in studying. In our framework, such interventions reduce the cost of effort c , which shifts the belief threshold $\hat{k}(\omega)$ downward. This, in turn, decreases the likelihood that agents fall into pessimistic, low-effort traps. Empirically, these traps may correspond to psychological states such as hopelessness or depression. Supporting this interpretation, unconditional cash transfers have been shown to improve psychological well-being (Haushofer and Shapiro (2016), Crosta et al., 2025)) and reduce depression (Christian et al. 2024). For a comprehensive meta-analysis on the impact of cash transfers on mental health and subjective well being, see McGuire et al. (2022).

Increasing the value of success $v(2\omega)$. Other policies operate by raising the actual rewards from success conditional on effort. **Conditional cash transfers**, such as Brazil’s Bolsa Família or Mexico’s PROGRESA link financial benefits to effort-driven behaviors like school attendance, thereby increasing the value of success on both margins: reducing the cost of effort and increasing its payoff.

Changing ω . Although the policies discussed above lower $\hat{k}(\omega)$ and make effort more attractive, they do not necessarily alter the underlying structural circumstance ω , which in our model represents factors like background endowments, such as family environment, social capital, or early-life health. However, in some cases ω itself can be directly improved. In the context of **discrimination**, policies aimed at reducing bias or correcting stereotypes can permanently improve the situation of disadvantaged groups. For example, Balafoutas and Sutter (2012) show that affirmative action increases women’s participation in competitive environments, effectively reducing gender-based disadvantage. In the context of **disability**, policies that facilitate labor market inclusion—such as remote work or accessibility accommodations—directly lower ω . Bloom et al. (2024) demonstrate that increases in remote work explain a significant share of the increase in disability employment in the US.

Temporary and permanent effects. Another important insight from our model is the distinction between temporary and permanent effects of policies. In their review of works related to the psychology of poverty, Park et al. (2024) state that “more work is also needed to understand whether belief-change interventions can have persistent effects in the long run”. Our model offers a framework to establish whether policies affecting beliefs can have long lasting impacts.

Suppose a policy that temporarily reduces the threshold to $\bar{k}(\omega) < \hat{k}(\omega)$. The long-term impact of this policy is critically dependent on both initial beliefs and the objective probability of success q . When $q = 1$, and if beliefs satisfy $\bar{k}(\omega) < p_0 < \hat{k}(\omega) < \frac{p_0+1}{2}$, the temporary policy can have a permanent effect: it induces effort in the short run, shifts beliefs above the threshold, and ultimately ensures convergence to the optimistic high-effort equilibrium even after the transfer ends.

In contrast, when $0 < q < 1$, Proposition 3 implies that a temporary policy cannot sus-

tain long-term effort and will only have a short-term effect. In this case, a permanent intervention, such as a **universal basic income** (UBI), is necessary to continuously reduce disadvantage and maintain a lower threshold $\bar{k}(\omega)$, to achieve lasting impacts on effort and welfare. As Banerjee et al. (2019) argue in their review, introducing measures such as UBI could alleviate psychological constraints and help poorer agents escape what they describe as a "pessimism trap". Such policies can induce effort both among agents who would have never felt optimistic enough to try in the first place, as well as among those who had stopped trying after experiencing failure.

3.3 Raising aspirations

Interventions that simultaneously sustain aspirations and beliefs (e.g., through role models or psychological support) are especially effective in fostering grit and avoiding pessimistic traps. Aspirations and grit interact in meaningful ways: when aspirations are raised such that the true success probability q exceeds the action threshold, and a grit intervention ensures a lower bound on belief updating, their combined effect can support people who would not benefit from either intervention in isolation. In addition, both *frustration* (when aspirations are exogenous) and *fatalism* (when aspirations evolve endogenously) are consistent with pessimism and low effort. However, temporary episodes of frustration can strengthen grit and lead to long-term welfare gains. In the words of Duckworth (2017, p. 109) aspirations and grit together create "a deep interest, a ready appetite for constant challenge, and an evolved sense of purpose".

These predictions are consistent with empirical evidence and can be more tightly interpreted through the lens of our model. Bernard et al. (2025) show that exposing rural Ethiopians to videos of successful role models raises aspirations and leads to increased investments in education and savings in the short- and long-run - consistent with our finding that higher aspirations can lower the belief threshold needed to trigger effort in a sustained way. Similarly, Beaman et al. (2012) find that repeated exposure to female leaders in Indian village councils changes gender attitudes and raises career aspirations among adolescent girls, suggesting that role models can reshape both perceived possibility and perceived payoff from effort.

However, McKenzie et al. (2022) find that raising ambitions among microentrepreneurs in the Philippines increased initial goal setting, but it also led to discouragement and lower economic performance when these goals were not met. This is consistent with our prediction that raising aspirations without sufficient support or belief updating can lead to frustration and reduced effort. Similarly, Genicot and Ray (2017) emphasize that aspirations must fall within an 'aspiration window' to motivate rather than demoralize, echoing the complementarity in our model between aspirations, belief formation, and grit-enhancing interventions. Together, these findings underscore the importance of coordinating aspiration-raising policies with mechanisms that support learning and resilience.

3.4 Changing $q(\omega)$

Changes in $q(\omega)$ are conceptually distinct from changes in the threshold $\hat{k}(\omega)$. Interventions that target $q(\omega)$ typically involve skill-building programs, training, or other ability-enhancing measures that improve the agent’s objective probability of success, rather than modifying incentives or constraints. While $q(\omega)$ may partially reflect an agent’s innate ability, in our framework it also captures structural and institutional factors, such as access to quality education, mentoring, or fair evaluation systems, that affect the likelihood of success conditional on effort. In such a setting, policies such as **affirmative action** can directly increase q by improving access to opportunities, thereby making effort more worthwhile and shifting long-run outcomes. For example, a student from a historically marginalized caste may be more motivated to study for a competitive university entrance exam if reservation policies improve their chances of admission conditional on effort. Similarly, targeted **mentoring** or **sponsorship programs** can raise q by providing disadvantaged students with access to networks and opportunities that would otherwise be unavailable.

Quota systems or **diversity hiring** initiatives in labor markets can further increase q by reducing evaluation biases or offering guaranteed positions, thereby incentivizing investment in skills and job preparation. Additional interventions such as guaranteed **internships** or **apprenticeships** for underrepresented youth create clear pathways into competitive fields. Finally, **bridge scholarships** or **conditional admission guarantees** in higher education can improve the likelihood of success given effort, reinforcing the perceived and actual returns to investing in human capital.

4 Conclusion

Our paper develops a simple yet rich dynamic framework to study how disadvantage and belief formation interact. The model uniquely integrates different types of policy interventions, material, psychological, and structural, and provides a unified lens to evaluate their nuances and effectiveness under various conditions. It can account for a wide range of interventions, from those that address discrimination and disability to those that boost self-confidence, grit, and aspirations, as well as material support, such as cash transfers. Our framework thus offers clear insights into when and why these policies may permanently address the structural effects of disadvantage (broadly defined, beyond just poverty) or only yield temporary improvements. The richness of these policy implications sets the foundation for future empirical research.

To the best of our knowledge, our work is the first to formally incorporate and distinguish psychological constructs such as locus of control, self-efficacy, grit, and aspirations into a single economic model. Crucially, we do not study these constructs in a vacuum: we highlight their role in reducing structural inequality in a stochastic world and clarify under which conditions distinctive interventions are welfare improving. Our framework also provides clear predictions about when role models and psychological interventions that increase subjective beliefs about the returns to effort will or will not be effective. In

addition, to our knowledge, we are the first to formally differentiate between fatalism and frustration, to emphasize the distinction between temporary and permanent impacts of both psychological and material interventions on effort, beliefs, and welfare, to show why psychological interventions tend to be more effective when introduced earlier and to clarify how farsightedness improves long-term welfare only when beliefs remain sufficiently optimistic over time.

Our framework opens many avenues for future research. Empirically, it motivates testing the predicted interactions between grit, beliefs, aspirations, and structural disadvantage, for example, using field experiments that can track belief and effort dynamics over time. Theoretically, extensions could explore general equilibrium effects, such as social learning, inequality dynamics, or collective action problems. Another promising direction is to study how policy interventions that target grit and beliefs interact with networks and peer effects, potentially amplifying or dampening individual responses. Furthermore, one could analyze the political economic implications of grit-enhancing interventions—for example, how improving individual agency and effort perceptions might shape demand for redistribution or support for different types of social policy. Finally, understanding how grit can benefit not only individual welfare, but also collective outcomes, such as public good provision, group resilience, or societal innovation, offers an interesting path for future work.

Beyond these directions, our framework could also be extended to incorporate insights from psychology on the emotional consequences of decision-making. In particular, it provides a natural structure for analyzing different forms of regret, such as regret from action versus inaction, as discussed by Gilovich et al. (1995). While we do not model regret explicitly, our threshold structure suggests a clear pathway: agents who fall below the effort threshold stop trying, yet may experience a lingering regret from not knowing what might have happened had they persisted. This counterfactual uncertainty may be more psychologically painful than the disappointment of trying and failing. Whether individuals anticipate this asymmetry and how it shapes their initial decisions could be a fruitful area of investigation.

References

- Ajzen, Icek. 2002. "Perceived behavioral control, self-efficacy, locus of control, and the theory of planned behavior." *Journal of applied social psychology* 32(4) : 665-683.
- Alan, Sule, Teodora Boneva, and Seda Ertac. 2019. "Ever Failed, Try Again, Succeed Better: Results from a Randomized Educational Intervention on Grit." *Quarterly Journal of Economics*, 134(2): 1121–1162.
- Ashraf, Nava , Gharad Bryan, Alexia Delfino, Emily A. Holmes, Leonardo Iacovone, Christian Johannes Meyer, and Ashley Pope. 2024. "Learning to see a world of opportunities." Working Paper.
- Balafoutas, Loukas, and Matthias Sutter. 2012. "Affirmative action policies promote women and do not harm efficiency in the laboratory". *Science*, 335(6068), 579-582.
- Bandura, Albert 1977. "Self-Efficacy: Toward a Unifying Theory of Behavioral Change." *Psychological Review*, 84(2): 191–215.
- Banerjee, Abhijit, Esther Duflo, Nathanael Goldberg, Dean Karlan, Robert Osei, William Parienté, Jeremy Shapiro, Bram Thuysbaert, and Christopher Udry. 2015. "A Multifaceted Program Causes Lasting Progress for the Very Poor: Evidence from Six Countries." *Science*, 348(6236): 1260799.
- Banerjee, Abhijit, Paul Niehaus, and Tavneet Suri. 2019. "Universal Basic Income in the Developing World.", *Annual Review of Economics*, 11:959-983.
- Barber, Brad M. and Odean, Terrance 2001. "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment." *The Quarterly Journal of Economics*, 116(1): 261–292.
- Beaman, Lori, Esther Duflo, Rohini Pande and Petia Topalova. 2012. "Female leadership raises aspirations and educational attainment for girls: A policy experiment in India." *Science*, 335(6068), 582-586.
- Bellman, Richard. 1956. "A Problem in the Sequential Design of Experiments." *Sankhyā: The Indian Journal of Statistics, Series B*, 16(3): 221–229.
- Bénabou, Roland, and Jean Tirole. 2002. "Self-confidence and Personal Motivation." *The Quarterly Journal of Economics* 117.3: 871-915.
- Bénabou, Roland, and Jean Tirole. 2006. "Belief in a Just World and Redistributive Politics." *The Quarterly Journal of Economics* 121.2 (2006): 699-746.
- Bernard, Tanguy, Stefan Dercon, Giulio Schinaia, Kate Orkin, and Alemayehu Taffesse. 2025. "The Future in Mind: Short and Long Run Impact of an Aspirations Intervention in Rural Ethiopia." *Quarterly Journal of Economics*, conditional accepted.
- Bertrand, Marianne, and Sendhil Mullainathan. "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination." *American Economic review* 94.4 (2004): 991-1013.

- Bernheim, B. Douglas, Debraj Ray and Sevin Yeltekin 2015. "Poverty and Self-Control." *Econometrica*, 83(5): 1877–1911.
- Bettinger, Eric, Sten Ludvigsen, Mari Rege, Ingeborg F. Solli and David Yeager. 2018. "Increasing perseverance in math: Evidence from a field experiment in Norway." *Journal of Economic Behavior & Organization*, 146, 1-15.
- Binmore, Ken. 2008. *Rational Decisions*. Princeton University Press.
- Blattman, Christopher, Nathan Fiala, and Sebastian Martinez. 2014. "Generating Skilled Self-Employment in Developing Countries: Experimental Evidence from Uganda." *Quarterly Journal of Economics*, 129(2): 697–752.
- Bloom, Nicholas, Gordon B. Dahl, and Dan-Olof Rooth. "Work from home and disability employment". No. w32943. National Bureau of Economic Research, 2024.
- Bobba, Matteo, and Verónica Frisancho. 2022. "Self-Perceptions about Academic Achievement: Evidence from Mexico City." *Journal of Econometrics*, 231(1): 58–73.
- Bohnet, Iris. "How to take the bias out of interviews." *Harvard Business Review* 16 (2016): 1-5.
- Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Schleifer 2019. "Beliefs about gender." *American Economic Review* 109.3 (2019): 739-773.
- Bradt, Russell N., S. M. Johnson, and Samuel Karlin. 1956. "On sequential designs for maximizing the sum of n observations". *The Annals of Mathematical Statistics* 27.4 1060-1074.
- Bryan, Gharad, Dean Karlan, and Scott Nelson. 2021. "Commitment Devices" *Annual Review of Economics*, 2:671–98.
- Buser, Thomas, Muriel Niederle, and Hessel Oosterbeek. 2014. "Gender, Competitive-ness, and Career Choices." *The Quarterly Journal of Economics*, 129(3): 1409–1447.
- Campos, Francisco, Michael Frese, Markus Goldstein, Leonardo Iacovone, Hillary C. Johnson, David McKenzie and Mona Mensmann. 2017. "Teaching personal initiative beats traditional training in boosting small business in West Africa." *Science*, 357(6357), 1287-1290.
- Christian, Cornelius, Lukas Hensel, and Christopher Roth. 2019. "Income shocks and suicides: Causal evidence from Indonesia." *Review of Economics and Statistics* 101.5 : 905-920.
- Cobb-Clark, Deborah A. 2015. "Locus of Control and the Labor Market." *IZA Journal of Labor Economics*, 33(4): 667–700.
- Coffman, Katherine Baldiga. 2014 "Evidence on self-stereotyping and the contribution of ideas." *The Quarterly Journal of Economics* 129.4 : 1625-1660.
- Crosta, Tommaso, Dean Karlan, Finley Ong, Julius Rüschenpöhler, and Christopher R. Udry. 2024. "Unconditional Cash Transfers: A Bayesian Meta-analysis of Randomized

Evaluations in Low and Middle Income Countries." National Bureau of Economic Research. (No. w32779).

Dalton, Patricio S., and Sayantan Ghosal. 2012. "Decisions with Endogenous Frames." *Social Choice and Welfare* 38(4): 585–600.

Dalton, Patricio S., Sayantan Ghosal, and Anandi Mani. 2016. "Poverty and Aspirations Failure." *Economic Journal*, 126(590): 165–188.

Di, Aurelia. 2025. "The Power of Words: How Role Models Influence Beliefs and Behaviors." Available at SSRN 4749301 (2025).

Duckworth, Angela L., Christopher Peterson, Michael D. Matthews, and Dennis R. Kelly. 2007. "Grit: Perseverance and Passion for Long-Term Goals." *Journal of Personality and Social Psychology*, 92(6): 1087–1101.

Duckworth, Angela. 2016. *Grit: The Power of Passion and Perseverance*. Scribner.

Flavell, John H. 1979. "Metacognition and Cognitive Monitoring: A New Area of Cognitive-Developmental Inquiry." *American Psychologist*, 34(10): 906–911.

Genicot, Garance, and Debraj Ray. 2017. "Aspirations and Inequality." *Econometrica*, 88(4): 1401–1435.

Ghosal, Sayantan, Smarajit Jana, Anandi Mani, Sandip Mitra, and Sanchari Roy 2022. "Sex workers, Stigma, and Self-image: Evidence from Kolkata Brothels." *Review of Economics and Statistics*, 104(3), 431–448.

Gilovich, Thomas and Victoria Husted Medvec. 1995. "The Experience of Regret: What, When, and Why." *Psychological Review*, 102(2): 379–395.

Guyon, Nina, and Huillery, Elise. 2020. "Biased Aspirations and Social Inequality at School: Evidence from French Teenagers." *The Economic Journal*, 130(631): 1673–1717.

Hammond, Peter J. 1976. "Endogenous Tastes and Stable Long-Run Choice." *Journal of Economic Theory*, 13(2): 329–340.

Haushofer, Johannes, and Jeremy Shapiro. 2016. "The short-term Impact of Unconditional Cash Transfers to the Poor: Experimental Evidence from Kenya." *The Quarterly Journal of Economics* 131.4: 1973–2042.

Heckman, James J., Jora Stixrud, and Sergio Urzua. 2006. "The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior." *Journal of Labor Economics*, 24(3): 411–482.

Heckman, James J., and Tim Kautz. 2012. "Hard Evidence on Soft Skills." *Labour Economics*, 19(4): 451–464.

Hertz, Tom, Tamara Jayasundera, Patrizio Piraino, Sibel Selcuk, Nicole Smith, and Alina Verashchagina. 2008. "The Inheritance of Educational Inequality: International Comparisons and Fifty-Year Trends." *The BE Journal of Economic Analysis & Policy*, 7(2): 1775.

- Hoff, Karla Ruth, and Priyanka Pandey. *Belief systems and durable inequalities: An experimental investigation of Indian caste*. Vol. 3351. World Bank Publications, 2004.
- Hoff, Karla, and Priyanka Pandey. 2006. "Discrimination, social identity, and durable inequalities." *American Economic Review* 96(2) : 206-211.
- Jensen, Robert 2010. "The (Perceived) Returns to Education and the Demand for Schooling." *The Quarterly Journal of Economics*, 125(2): 515–548.
- Kaur, Jalnidh 2023. "How much do I matter? Teacher Self-beliefs, Effort, and Student Learning."
- Kremer, Michael, Guatama Rao, and Frank Schilbach. 2019. "Behavioral Development Economics." In *Handbook of Behavioral Economics: Applications And Foundations*, Volume 3, edited by Douglas Bernheim, Stefano DellaVigna, and David Laibson, 1–81. Amsterdam: Elsevier.
- Kuchler, Theresa, and Basit Zafar. 2019. "Personal Experiences and Expectations about Aggregate Outcomes." *The Journal of Finance* 74.5, 2491-2542.
- La Ferrara, Eliana, Chong, A., and Suzanne Duryea 2012. "Soap Operas and Fertility: Evidence from Brazil". *American Economic Journal: Applied Economics*, 4(4), 1–31.
- Kahneman, Daniel, and Amos Tversky 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 47(2), 263–291.
- Lybbert, Travis J., and Bruce Wydick. 2018. "Poverty, Aspirations, and the Economics of Hope." *Economic Development and Cultural Change* 66.4 : 709-753.
- Loury, Glenn C. 2019. "Why Does Racial Inequality Persist? Culture, Causation, and Responsibility." *Manhattan Institute Essays*, 1-16.
- Mahalanobis, Prasanta Chandra.1940, "A Sample Survey of Acreage under Jute in Bengal", *Sankhyā*, Vol. 4, pp. 511-530.
- Malmendier, Ulrike, and Stefan Nagel. 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?" *The Quarterly Journal of Economics*, 126(1): 373–416.
- Malmendier, Ulrike, and Stefan Nagel. 2016. "Learning from Inflation Experiences." *The Quarterly Journal of Economics*, 131(1): 53–87.
- Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao. 2013. "Poverty Impedes Cognitive Function." *Science*, 341(6149): 976–980.
- McKelway, Madeline. 2025. "Women's Self-Efficacy and Economic Outcomes: Experimental Evidence from India." Available at SSRN 5158816.
- McGuire, Joel, Caspar Kaiser, and Anders M. Bach-Mortensen. 2022. "A systematic review and meta-analysis of the impact of cash transfers on subjective well-being and mental health in low-and middle-income countries." *Nature Human Behaviour* 6.3 : 359-370.
- McKenzie, David, Aakash Mohpal, and Dean Yang. 2022. "Aspirations and Financial Decisions: Experimental Evidence from the Philippines." *Journal of Development Economics*,

156:102846.

Narayan, Ambar, Roy Van der Weide, Alexandru Cojocaru, Christoph Lakner, Silvia Redaelli, Daniel Gerszon Mahler, Rakesh Gupta N., and Stefan Thewissen. 2018. "Fair Progress? Economic Mobility across Generations around the World." *World Bank Publications*.

Niederle, Muriel, and Lise Vesterlund. 2007. "Do women shy away from competition? Do men compete too much?." *The Quarterly Journal of Economics* 122.3 (2007): 1067-1101.

Nikolov, Plamen and Nusrat Jimi 2020. "What Factors Drive Individual Misperceptions of the Returns to Schooling in Tanzania? Some Lessons for Education Policy." *Applied Economics*, 59: 1063–1096.

Nguyen, T. 2008. "Information, Role Models and Perceived Returns to Education: Experimental Evidence from Madagascar." Unpublished manuscript.

Park, Ye Ren, Yuen Ho, Kristina Hallez, Supreet Kaur, Mahesh Srinivasan, and Jiaying Zhao 2025. "The Psychology of Poverty: Current and Future Directions". *Current Directions in Psychological Science*, 34(1), 21-28.

Parzen, E. 1962. "On Estimation of a Probability Density Function and Mode." *The Annals of Mathematical Statistics*. 33 (3): 1065–1076.

Piketty, Thomas. 1995. "Social mobility and redistributive politics." *The Quarterly Journal of Economics*. 110.3: 551-584.

Pollak, Robert A. 1978. "Endogenous Tastes in Demand and Welfare Analysis." *American Economic Review*, 68(2): 374–379.

de Quidt, Jonathan, and Johannes Haushofer. 2016. "Depression for Economists." *Mimeo*.

Rotter, Julian B. 1966. "Generalized Expectancies for Internal versus External Control of Reinforcement." *Psychological Monographs: General and Applied*, 80(1): 1–28.

Robbins, Herbert. 1952. "Some Aspects of the Sequential Design of Experiments." *Bulletin of the American Mathematical Society*, 58(5): 527–535.

Rury, Derek, and Scott E. Carrell. 2023. "Knowing what it takes: The effect of information about returns to studying on study effort and achievement." *Economics of Education Review* 94 : 102400.

Seligman, Martin E. P. 1972. "Learned helplessness." *Annual Review of Medicine*. 23, 407-412

Talsma Kate, Benhamin Schuz and Kimberley Norris 2019. "Miscalibration of Self-Efficacy and Academic Performance: Self-Efficacy \neq Self-Fulfilling Prohecy." *Learning and Individual Differences*, 63: 109–121.

Treurniet, Mark, and Robert Lensink. 2022. "Belief-based poverty traps and the effects of material and psychological development interventions." *SOM Research Reports* Vol. 2022007-EEF.

Van Nieuwerburgh, Stijn, and Laura Veldkamp. 2006. "Learning Asymmetries in Real Business Cycles." *Journal of Monetary Economics*, 53(4): 753–772.

Veldkamp, Laura. 2005. "Slow Boom, Sudden Crash." *Journal of Economic Theory*, 124(2): 230–257.

Von Weizsäcker, Carl Christian. 1971. "Notes on Endogenous Change of Tastes." *Journal of Economic Theory*, 3(4): 345–372.

Wollburg, Claudia, Janina I. Steinert, Aaron Reeves, and Edward Nye. 2023. "Do cash transfers alleviate common mental disorders in low- and middle-income countries? A Systematic Review and Meta-Analysis." *PLOS ONE*, 18(2): e0281283.

Yeager, David S., et al. "A national experiment reveals where a growth mindset improves achievement." *Nature* 573.7774 (2019): 364-369.

Zimmerman, Barry J. 2000. "Attaining Self-Regulation: A Social Cognitive Perspective." In *Handbook of Self-Regulation*, edited by Monique Boekaerts, Paul R. Pintrich, and Moshe Zeidner, 13–39. San Diego: Academic Press.

Appendix A Proofs of Propositions

Proof of Proposition 1. Suppose $e(p_t) = 1$ for all $t > 1$. Then $E[p_t | p_{t-1}, e_t = 1] = q \frac{p_{t-1}+1}{2} + (1-q) \frac{p_{t-1}}{2} = \frac{p_{t-1}}{2} + \frac{q}{2}$. By the law of iterated expectations, this is equivalent to $E[p_t] = \frac{E[p_{t-1}]}{2} + \frac{q}{2}$. By recursive substitution, we obtain $E[p_t] = \frac{p_0}{2^t} + \sum_{k=1}^t \frac{q}{2^k} = \frac{p_0}{2^t} + (1 - \frac{1}{2^t})q$. Therefore, $\lim_{t \rightarrow \infty} E[p_t] = \lim_{t \rightarrow \infty} [\frac{p_0}{2^t} + q(1 - \frac{1}{2^t})] = q$. If $q = 0$, it is easy to see that, conditional on effort, $p_t = \frac{p_{t-1}}{2} = \frac{p_0}{2^t}$, which converges to 0, which is equal to q . In the same way, if $q = 1$, $p_t = \frac{p_{t-1}+1}{2} = \frac{p_0}{2^t} + (1 - \frac{1}{2^t})$, which converges to 1, which is equal to q . ■

Proof of Proposition 2. Consider an agent with initial circumstances denoted by ω .

1) Suppose $e_T = \lim_{t \rightarrow T} e_t(p_{t-1}) = 1$, then it must mean that for each $0 < t \leq T$, $p_t(e_t) \geq \hat{k}(\omega)$. Then, there exists a subsequence $\{p_{k_t}\}$ which converges to a point p_T such that $\lim_{t \rightarrow T} p_t(e_t) = p_T \geq \hat{k}(\omega)$. Suppose that $p_T = \lim_{t \rightarrow T} p_t(e_t) > \hat{k}(\omega)$. Then, $p_t > \hat{k}(\omega)$ for each $0 < t \leq T$. But then it must be $e_t = 1$ each $0 < t \leq T$ and therefore $\lim_{t \rightarrow T} e_t(p_{t-1}) = e_T = 1$. Hence, there is a long-term outcome with a high effort ($e_T = 1$) and optimistic beliefs $p_T > \hat{k}(\omega)$.

Suppose now that $\lim_{t \rightarrow T} e_t(p_{t-1}) = 0$. Then, it must mean that there exists a time $t_n < T$ such that $e_t = 0$. But then $p_{t_n} < \hat{k}(\omega)$ and therefore $p_t = p < \hat{k}(\omega)$ for all $t > t_n$, and it must be $\lim_{t \rightarrow T} p_t(e_t) = p < \hat{k}(\omega)$. Suppose that $\lim_{t \rightarrow T} p_t(e_t) = p < \hat{k}(\omega)$. Then, there exists a time $t_n < T$ for which $p_t < \hat{k}(\omega)$, for each $t_n < t \leq T$, which implies $e_t = 0$ for all $t_n < t \leq T$ and so $\lim_{t \rightarrow T} e_t(p_{t-1}) = 0$. Hence, there is a long-term outcome with low effort ($e = 0$) and pessimistic beliefs $p < \hat{k}(\omega)$.

2) Suppose that $q > \hat{k}(\omega) = \frac{c}{v(2\omega)}$, then $qv(2\omega) - c > 0$. The expected utility for the long-term, high effort outcome is $E(u|e_T = 1) = qv(2\omega) - c$, while the expected utility for long-term, low effort outcomes is $E(u|e'_T = 0) = 0$. Hence, if $q > \hat{k}(\omega) = \frac{c}{v(2\omega)}$, $E(u|e_T = 1) > E(u|e'_T = 0)$.

3) Suppose that $q = 1$. Assume that $p_0 \geq \hat{k}(\omega)$. Then $e_1 = 1$. But then $p_t = \frac{p_{t-1}+1}{2} > p_{t-1}$, for each $t > 1$, and so $p_t > p_0 > \hat{k}(\omega)$. Hence, the sequence converges to the long-term outcome of high-effort optimistic belief since $\lim_{T \rightarrow \infty} e_T(p_{t-1}) = 1$ and $\lim_{T \rightarrow \infty} p_T(e_T) = 1 > \hat{k}(\omega)$ as $e_t = 1$ for each $t > 0$ and $\lim_{T \rightarrow \infty} p_T(e_T) = 1 = q$, by Proposition 1. On the other hand, if $p_0 < \hat{k}(\omega)$, $e_t = e_1 = 0$, for each t . Hence, the sequence converges to the long-term result of low-effort pessimistic belief as $\lim_{T \rightarrow \infty} e_T(p_{t-1}) = 0$ and $\lim_{T \rightarrow \infty} p_t(e_t) = p_0 < \hat{k}(\omega)$ as $e_t = 0$ for each $t > 0$.

Suppose instead that $q = 0$. Then $p_t < p_{t-1}$ whenever $e_t = 1$. Therefore, there exists a period $t_n < T$ such that $p_{t_n} < \hat{k}(\omega)$. Hence, the sequence will converge to the long-term result of low-effort pessimistic belief as $\lim_{T \rightarrow \infty} e_T(p_{t-1}) = 0$, as $e_t = 0$ for each $t > t_n$, and $\lim_{T \rightarrow \infty} p_T(e_t) = p_{t_n} < \hat{k}(\omega)$.

Finally, consider the case when $0 < q < 1$. Suppose that there exists a $t_n < T$ in which $e_{t_n} = 0$. Then it must mean $p_{t_n-1} < \hat{k}(\omega)$. Therefore, $p_t = p_{t_n-1} < \hat{k}(\omega)$ and $e_t = 0$ for each $t > t_n$. Hence, $\lim_{T \rightarrow \infty} e_T(p_{t-1}) = 1$ and $\lim_{t \rightarrow T} p_t(e_t) = p > \hat{k}(\omega)$ only if $e_t(p_t) = 1$

and $p_t(e_t) > \hat{k}(\omega)$ for all t . Therefore, the low effort, pessimistic belief long-term outcome is an absorbing state.

Suppose $p = \lim_{T \rightarrow \infty} p_T(e_t) > \hat{k}(\omega)$ and there exists a $t_n < T$ such that $p_{t_n} < \hat{k}(\omega)$. But then, this would imply that $e_{t_n} = 0$ and, therefore, $p_t = p_{t_n} < \hat{k}(\omega)$ for each $t > t_n$, a contradiction. Therefore, if $p = \lim_{T \rightarrow \infty} p_T(e_t) > \hat{k}(\omega)$, then $p_t > \hat{k}(\omega)$ for all t , and so $e_t = 1$ for all t .

Therefore, $\lim_{T \rightarrow \infty} e_T = 1, \lim_{T \rightarrow \infty} p_T = p > \hat{k}(\omega)$ if and only if $e_t(p_{t-1}) = 1, p_t(e_t) > \hat{k}(\omega)$ for all $t > 0$. Suppose that $e_t = 1$ for all t . Then, there exists a finite number m such that $\frac{p_t}{2^m} < \hat{k}$, i.e. if the agent fails m consecutive times, then the agent stops exerting effort. Since we let $T \rightarrow \infty$, the probability of m consecutive failures converges to 1, by the second Borel-Cantelli Lemma. Hence, the high effort, optimistic belief long-term outcome is not an absorbing state. ■

Proof of Proposition 3. Fix $p_{0,i} = p_{0,j} = p_{0,q}$ and let $\omega_i < \omega_j$. Then $\hat{k}(\omega_i) = \frac{c}{v(2\omega_i)} \leq \frac{c}{v(2\omega_j)} = \hat{k}(\omega_j)$, as $v(\cdot)$ is an increasing function. Each agent can reach one of the two long-term outcomes and, in particular, will enter a pessimistic belief low-effort trap as soon as their belief falls below the threshold, as in Proposition 2. There can be different cases depending on where the prior belief lies with respect to the threshold for the disadvantaged agent. Clearly, if $\hat{k}(\omega_i) > p_0$, then it is clear that the disadvantaged agent will be stuck in pessimistic belief low-effort trap since $t = 0$. Then, the disadvantaged agent is (weakly) more likely to enter a pessimistic belief low-effort trap. Then, let us consider the case in which $\hat{k}(\omega_i) < \hat{k}(\omega_j) < p_0$. Then, at any t ,

$$Pr(p_{t,i} \geq \hat{k}(\omega_i)|q) \geq Pr(p_{t,j} \geq \hat{k}(\omega_j)|q)$$

Hence, a more disadvantaged agent is more likely to enter a pessimistic belief low-effort trap. ■

Appendix B The Role of Farsightedness

In this appendix, we examine the implications of farsightedness for our analysis. With forward looking behavior, the key change is that now the agent will take into account the implications of current effort choice e_t on next period's beliefs p_t . We allow for the possibility that the agent may discount future payoffs. Let $0 \leq \delta \leq 1$ denote the agent's discount factor.

We start by considering a one-step farsighted agent. In the myopic case, discussed in the paper, we provided the condition for exerting effort, i.e. $e_t = 1$ iff $p_{t-1} \geq \frac{c}{v(2\omega)}$. Then, the agent will also take into account the utility derived from the effort decision in $t + 1$. In this case, the decision in $t + 1$ will be made according to the myopic decision protocol. First, we can easily see that if a partially farsighted agent does not exert effort at time t , then they anticipate that they will also not exert effort at $t + 1$ and therefore the expected utility of not exerting effort at t is equal to zero. On the other hand, the expected utility of exerting effort is the following.

$$\begin{aligned} E_1[u(e_t = 1, p_{t-1})] &= p_{t-1}(v(2\omega) + \delta E_0[u(e_{t+1} = 1, p_t)] \mathbb{1}_{\frac{p_{t+1}}{2} > \frac{c}{v(2\omega)}}) \\ &\quad + (1 - p_{t-1})\delta E_0[u(e_{t+1} = 1, p_t)] \mathbb{1}_{\frac{p_t}{2} \geq \frac{c}{v(2\omega)}} - c \end{aligned}$$

where we denote by $E_1[u]$ the expected utility of a one step farsighted agent and by $E_0[u]$ the myopic expected utility. The agent will decide to exert effort at time t whenever this expression is greater than 0. This will induce a threshold strategy, meaning that the partially farsighted agent will also exert effort when $p_{t-1} > k_1$, where $k_1 = \frac{2+\delta}{2\delta}(\sqrt{1 + \frac{8c\delta}{v(2\omega)(2+\delta)^2}} - 1)$ is the threshold for the one step farsighted agent. We can also notice that $k_1 < k_0 = \hat{k}$, as by simple computation we have $E_1[u(e_t = 1, p_{t-1} = k_0)] > 0$, and by the monotonicity of the expected utility in beliefs.

Moreover, we can also derive a different result using the same argument. Take a belief $p_{t-1} < k_0$, that is, using this belief a myopic agent will not exert effort. By continuity in δ , there exists a value $\bar{\delta}$ such that whenever $\delta < \bar{\delta}$, $E_1[u(e_t = 1, p_{t-1})] < 0$ and therefore also a one step farsighted agent will not exert effort.

We now move to the full farsighted case, in which the agent will discount the utility of all T periods. First, notice that the expected utility has a recursive representation as follows

$$\begin{aligned} E_T[u(e_1 = 1, p_0)] &= p_0(v(2\omega) + \delta E_{T-1}[u(e_1 = 1, p_1)] \mathbb{1}_{\frac{p_0+1}{2} > k_{T-1}}) \\ &\quad + (1 - p_0)\delta E_{T-1}[u(e_1 = 1, p_1)] \mathbb{1}_{\frac{p_0}{2} \geq k_{T-1}} - c \end{aligned}$$

Where k_t is the belief threshold for an agent who considers the next t periods (a t -step farsighted agent).

We characterize the fully farsighted case and compare the decisions of a farsighted agent versus a myopic agent. We state a result replicating the previous argument, showing both how the threshold moves in the case of (full) farsightedness and the combined role of beliefs and the discount factor.

Proposition 4. Consider an agent who is fully farsighted, i , and a myopic agent, j . Then

- (i) The threshold for the fully farsighted agent is below the myopic threshold, $k_T < \hat{k}$
- (ii) The relation between thresholds, for any partial farsighted agent, is such that $k_{n-1} > k_n$, for each $0 < n \leq T$.
- (iii) For any belief $p_0 < \hat{k} = \frac{c}{v(2\omega)}$, there exists $\bar{\delta}$ such that $e_t^i = e_t^j = 0$ for each $1 \leq t \leq T$ for the farsighted agent.

Proof. (i) Let $p_0 = \frac{c}{v(2\omega)} = \hat{k}$. We prove by induction that $E_T[u(e_1 = 1, \hat{k})] > 0$. We already showed, by explicit computation, that $E_1[u(e_1 = 1, \hat{k})] > 0$ (and therefore $k_1 < k_0$), where $E_1[u(e_1 = 1, \hat{k})]$ is the expected utility of a one step farsighted agent by exerting effort in period 1 when their initial beliefs are exactly equal to the myopic threshold. Suppose now that the expected utility, for a $n - 1$ -step farsighted agent (with $1 < n \leq T$), from exerting effort in period 1 when their beliefs are equal to the myopic threshold is positive. This first implies that the threshold $k_{n-1} < \hat{k}$, by monotonicity of the expected utility in beliefs. Then, the expected utility for the n -step farsighted agent is

$$E_n[u(e_1 = 1, \hat{k})] = \hat{k}v(2\omega) - c + \delta(\hat{k}E_{n-1}[u(e_2 = 1, p_1)]\mathbb{1}_{\frac{\hat{k}+1}{2} > k_{n-1}} + (1 - \hat{k})E_{n-1}[u(e_2 = 1, p_1)]\mathbb{1}_{\frac{\hat{k}}{2} \geq k_{n-1}})$$

In this expression, we can see that $\hat{k}v(2\omega) - c = 0$, by definition of the myopic threshold. Moreover, since we assumed that $E_{n-1}[u(e_2 = 1, \hat{k})] > 0$, then it must be that also $E_{n-1}[u(e_2 = 1, p_1)] > 0$ as by succeeding today, the belief will increase, so that $p_1 > p_0 = \hat{k} > k_{n-1}$. But then,

$$E_n[u(e_1 = 1, \hat{k})] > \delta\hat{k}E_{n-1}[u(e_2 = 1, p_1)] > 0$$

Hence, by monotonicity of the expected utility in beliefs, $k_n < \hat{k}$. Therefore, for all n , $k_n < \hat{k}$, and in particular the threshold of the fully farsighted agent, k_T , is such that $k_T < k_0$.

(ii) First, we show that $k_n > 2k_{n-1} - 1$. The expression for the expected utility for an n -step farsighted agent is

$$E_n[u(e_1 = 1, p_0)] = p_0v(2\omega) - c + p_0\delta E_{n-1}[u(e_1 = 1, p_1)]\mathbb{1}_{\frac{p_0+1}{2} > k_{n-1}} + (1 - p_0)\delta E_{n-1}[u(e_1 = 1, p_1)]\mathbb{1}_{\frac{p_0}{2} \geq k_{n-1}}$$

Computing this expression at k_n and suppose that $k_n < 2k_{n-1} - 1$. Then, $\frac{k_n+1}{2} < k_{n-1}$ and, therefore, all continuation utilities are 0. Moreover, $k_nv(2\omega) - c < 0$ as $k_n < \hat{k}$, from (i). But then $E_n[u(e_1 = 1, k_n)] < 0$, a contradiction. Therefore, if the agent exerts effort in the first period and succeeds, then they continue exerting effort.

We can now rewrite the expected utility from exerting effort for a farsighted agent of the n -step in the following way:

$$E_n[u(e_1 = 1, p_0)] = E_{n-1}[u(e_1 = 1, p_0)] + \delta^n \sum_{\{p_{n-1} > \hat{k}\}} (\prod_{t=0}^{n-2} p_t) E_0[u(e_1 = 1, p_{n-1})]$$

That is, the expected utility at step n is the sun of the expected utility at the $(n - 1)^{th}$ -step plus the instantaneous utility in the last period, whenever the belief at the last stage is above the myopic threshold. Notice that since $k_t < \hat{k}$, from (i), we do not need to specify that the expected utility in the last stage is realized conditional on exerting effort in all previous periods.

Now compute the farsighted utility expected in n -step when $p_0 = k_{n-1}$. By definition, we have $E_{n-1}[u(e_1 = 1, p_0)] = 0$. Moreover, we are sure that there is an instance for which $p_{n-1} > \hat{k}$, which is when the agent always succeeds, from the previous point. But then,

$$\delta^n \sum_{\{p_n > \hat{k}\}} (\prod_{t=0}^{n-1} p_t) E_0[u(e_1 = 1, p_{n-1})] > 0$$

Hence, $k_n < k_{n-1}$, as $E_n[u(e_1 = 1, p_0)] > 0$ and by monotonicity of the expected utility in the belief.

(iii) Let $p_0 < \frac{c}{v(2\omega)}$. Then, any myopic agent will be stuck in the low effort trap. Moving to the farsighted agent, we have that

$$\begin{aligned} E_T[u(e_1 = 1, p_0)] &= p_0 v(2\omega) - c + \delta(p_0 E_{T-1}[u(e_2 = 1, p_1)] \mathbb{1}_{\frac{p_0+1}{2} > k_{T-1}} \\ &\quad + (1 - p_0) E_{T-1}[u(e_2 = 1, p_1)] \mathbb{1}_{\frac{p_0}{2} \geq k_{T-1}}) \end{aligned}$$

Since we assumed $p_0 < \frac{c}{v(2\omega)}$, it is easy to see that there must exist a value $\bar{\delta}$ such that $E_T[u(e_1 = 1, p_0)] = 0$, by the continuity of the expected utility in δ . Therefore, for each $\delta < \bar{\delta}$, the farsighted agent will not exert effort in the first period and, therefore, will not exert effort in any period subsequently afterward. Hence, for each $p_0 < \frac{c}{v(2\omega)}$, there exists a threshold level on the discount factor $\bar{\delta}$, both the myopic and farsighted agents will never exert effort for $\delta < \bar{\delta}$. ■

The preceding proposition shows that introducing farsightedness makes agents more likely to exert and sustain effort, although the extent to which they do so is linked both to their level of discounting and to their beliefs. Finally, the result continues to hold for any partially farsighted agent.

Appendix C The Role of Aspirations

In this appendix, we extend our model to explicitly incorporate aspirations as follows. Let $g_t \in \{0, 2\omega\}$ denote the agent's aspirations for outcomes in period $t \geq 0$.²¹ We assume that all agents start with the same initial level of aspirations, denoted by g_0 .

We consider two possible dynamics for aspirations. In the first case, aspirations are static: $g_t = g_0$ for all $t > 1$, which means that they do not evolve in response to the outcomes realized. This can represent, for example, an agent who sets a fixed goal early on, such as passing an exam or securing a stable job, and does not revise it regardless of their experiences. In the second case, aspirations are adaptive: $g_t = \theta_{t-1}$, which means they adjust instantaneously to match the result achieved in the previous period. This captures agents who continuously update their goals based on recent performance, such as a student who raises their aspirations after good grades or lowered them after disappointing results.²²

Note that beliefs and aspirations evolve differently over time. Recall that subjective beliefs evolve according to the following rule:

$$p_t(e_t) = \begin{cases} \frac{p_{t-1}+1}{2} & \text{if } e_t = 1 \text{ and } \theta_t = 2\omega \\ \frac{p_{t-1}}{2} & \text{if } e_t = 1 \text{ and } \theta_t = 0 \\ p_{t-1} & \text{if } e_t = 0 \end{cases}$$

Thus, beliefs depend jointly on the agent's effort choice and the realized outcome, while aspirations depend only on outcomes. If the agent does not exert effort, their subjective beliefs remain unchanged, but their aspirations can still adjust downward after poor results. For example, consider a student who decides not to study for an exam (that is, $e_t = 0$). In this case, their subjective belief of success remains unchanged, as they did not actively engage or receive feedback on the success of putting effort. However, if the student receives a poor grade (low θ_t), they may revise their aspirations downward, perhaps shifting from aiming to attend a top university to setting a more modest academic goal, even though their beliefs of success conditional on effort have not been updated.

The instantaneous utility function for an agent is now specified as follows: $U(e_t, \theta_t, g_t) = v(\theta_t) + r(\theta_t - g_t) - c(e_t)$, where following Kahneman and Tversky (1979), $r(x)$ denotes a value function where $r(x) = ax$ if $x \geq 0$ and $r(x) = bx$, if $x < 0$. Evidently, if $a = b = 0$, the aspirations are irrelevant and we revert to the model studied in Section 2. If $0 < a < b$, an assumption we make in this throughout this appendix, aspirations will matter and the value function will satisfy loss aversion.²³

²¹In effect, we require aspirations to be a feasible outcome.

²²Dalton et al. (2016) implicitly make the assumption that $g_t = \theta_{t-1}$. One could allow partially adaptive aspirations of the form $g_t = \lambda g_{t-1} + (1 - \lambda)\theta_{t-1}$, where $0 \leq \lambda \leq 1$. Here, λ represents the weight the agent places on past aspirations versus past outcomes. Our analysis focuses on the extreme cases: $\lambda = 1$ (static) and $\lambda = 0$ (fully adaptive). However, in our two outcome setting, for intermediate cases with $0 < \lambda < 1$, aspirations may not be feasible.

²³The value function does not satisfy diminishing sensitivity. This assumption is made for simplicity of exposition. Allowing for diminishing sensitivity will change the formal exposition of our results without

As before, we consider a short-term decision problem in which the agent chooses to exert effort at time t if

$$E(U(e_t = 1, p_t, g_t)) > E(U(e_t = 0, p_t, g_t)),$$

where

$$EU(e_t = 1, p_t, g_t) = p_{t-1}[v(2\omega) + r(2\omega - g_t)] + (1 - p_{t-1})r(0 - g_t) - c,$$

and

$$EU(e_t = 0, p_t, g_t) = r(0 - g_t).$$

This condition implies a threshold value, denoted by $\bar{k}(\omega, g_t)$, such that the agent exerts effort only if

$$p_{t-1} > \bar{k}(\omega, g_t) = \frac{c}{v(2\omega) + [2a\omega + g_t(b - a)]}.$$

Intuitively, when $b > a$, the agent places more psychological weight on falling short of aspirations than on exceeding them. This increases the effective cost of failure, raising the threshold $\bar{k}(\omega, g_t)$ and making the agent less likely to exert effort when pessimistic.

According to our assumptions — namely that $a, b > 0$ with $b > a$ and that both ω and g_t are strictly positive — we have $2a\omega + g_t(b - a) > 0$. It follows that $\bar{k}(\omega, g_t) \leq \hat{k}(\omega)$. Thus, the additional term capturing the impact of aspirations on preferences effectively lowers the threshold belief needed for the agent to find exerting effort worthwhile. Moreover, $\bar{k}(\omega, g_t)$ is decreasing in both ω and g_t , meaning that higher circumstances or higher aspirations reduce the belief required to justify effort. Keeping everything else constant, an agent with a higher level of disadvantage (lower ω) will face a higher threshold than a more advantaged agent. However, having higher (yet realistic) aspirations reduces this threshold, thereby encouraging effort even when initial beliefs are moderate. For example, a student with higher aspirations — such as aiming to enter a top university — will require a lower belief threshold to justify studying hard compared to a student with lower aspirations who only aims to pass exams, all else equal. The interaction between aspirations and disadvantage shapes the agent's willingness to exert effort over time. We now turn to the dynamic sequence of short-term decisions that jointly determine long-term outcomes, examining how effort, beliefs, and aspirations coevolve and settle into stable patterns or traps.

Given ω , let $(e_t(p_{t-1}, g_t), p_t(e_t), g_t)$, for $t > 0$, denote the sequence of short-term decisions and the resulting update of beliefs and aspirations. The corresponding long-term result is indicated by (e_T, p_T, g_T) , where $\lim_{t \rightarrow T} e_t = e_T$, $\lim_{t \rightarrow T} p_t = p_T$, and $\lim_{t \rightarrow T} g_t = g_T$. As before, when characterizing long-term outcomes, we state the results for all $T > 0$ or as $T \rightarrow \infty$ when $q = 1$ or $q = 0$. When $0 < q < 1$, our focus is on the absorbing states.

changing their underlying intuition. The most general statement of assumption required on $r(\cdot)$ is as follows: (a) $r(\cdot)$ is continuously differentiable over $x > 0$ and over $x < 0$, with $r'(x) > 0$ for both $x > 0$ and $x < 0$, (b) $r''(x) \leq 0$ for $x > 0$, (c) $r''(x) \geq 0$ for $x < 0$, (d) $\lim_{x \rightarrow 0^+} v'(x) > \lim_{x \rightarrow 0^-} v'(x)$ (loss aversion), (e) $2r'(x) > r'(y) > r'(x)$, $x > 0, y < 0$.

Note that when $g_t = g_0$ for all $t > 1$, we are in a setting similar to that of Section 2, in which only beliefs evolve over time while aspirations remain fixed. In this case, replacing $\hat{k}(\omega)$ with $\bar{k}(\omega, g_0)$ in Proposition 2 ensures that its characterization of long-term outcomes continues to hold. Specifically, there will be two possible long-term outcomes: one with high effort and optimistic beliefs, and another with low effort and pessimistic beliefs. The former predominates the latter whenever $q > \bar{k}(\omega, g_0)$. When $q = 1$, if $p_0 \geq \bar{k}(\omega, g_0)$, then the high-effort outcome is stable; the opposite holds if the inequality is reversed. When $q = 0$ or $0 < q < 1$, the low-effort pessimistic outcome becomes the unique absorbing state.

We now examine two key implications of introducing fixed aspirations $g_0 = 2\omega$ into the baseline model. First, aspirations can expand the region of initial beliefs from which optimistic, high-effort long-run outcomes emerge. Formally, the belief threshold becomes $\bar{k}(\omega, g_0) \leq \hat{k}(\omega)$, with strict inequality whenever $g_0 > 0$. This means that having positive aspirations strengthens the short-term incentive to exert effort. Consequently, when $q = 1$, the basin of attraction leading to long-run success is strictly larger.

However, these aspirations can generate *short-term frustration* if outcomes fall short of expectations. Suppose $g_t = g_0 = 2\omega$ for all $t \geq 0$, and the agent exerts effort but observes $\theta_t = 0$. Then the agent fails to meet their aspirations despite trying, leading to frustration. For instance, a student from a disadvantaged background may aim to attend a top university and study diligently, but fail early tests—creating emotional or motivational costs even when effort is exerted.

Second, when $g_0 = 2\omega$ and $0 \leq q < 1$, the low-effort, pessimistic outcome becomes the only absorbing state in the long run. In this case, even if the agent starts out with high aspirations and exerts effort initially, they will eventually give up, leading to *long-term frustration*. The resulting equilibrium is characterized by $e_T = 0$, $p_T < \bar{k}(\omega, g_0)$, and $g_T = g_0$, with unfulfilled aspirations. The higher the structural disadvantage (ω), the more severe this long-term frustration becomes. For example, a highly motivated student from an extremely disadvantaged background may exert effort early on but eventually drop out due to repeated failure, despite their initial optimism and ambition.

Next, we consider the case where aspirations are adaptive, specifically when $g_t = \theta_{t-1}$. We now present results that are direct analogs of Propositions 2 and 3, adapted to this dynamic aspirations setting. These results characterize the long-run outcomes under adaptive goals, assess their welfare implications, and examine their local stability properties.

Proposition 5. *Given ω , the long-term results, their welfare ranking, and stability are as follows.*

1. *Long-term outcomes: For all $T > 0$, there are two long-term outcomes, that is, (i) High effort ($e_T = 1$), optimistic beliefs and fluctuating aspirations with $p_T > \bar{k}(\omega, 0)$ and $g_T \in \{0, 2\omega\}$ and (ii) Low effort ($e_T = 0$), pessimistic beliefs and low aspirations $p_T < \bar{k}(\omega, g_T)$, $g_T = 0$;*
2. *Welfare ranking: For all $T > 0$, if $q > \hat{k}(\omega, 0)$, the first long-term result with $e_T = 1$, $p_T > \bar{k}(\omega, 0)$ and $g_T \in \{0, 2\omega\}$ welfare dominates the second long-term result with $e_T = 0$, $p_T < \bar{k}(\omega, g_T)$, $g_T = 0$, that is, $E(U|e_T = 1, q, 0) > E(U|e_T = 0, q, 0)$.*

3. Stability:

- When $q = 1$, if $p_0 > \bar{k}(\omega, g_0)$, the sequence converges to the first long-term outcome (high effort, optimistic beliefs, and high aspirations); if $p_0 < \bar{k}(\omega, g_0)$, it converges to the second outcome (low effort, pessimistic beliefs, and low aspirations).
- When $q = 0$, the sequence converges to the second long-term outcome (low effort, pessimistic beliefs, and low aspirations).
- When $0 < q < 1$, the second long-term outcome (low effort, pessimistic beliefs, and low aspirations) is the unique absorbing state:

$$\lim_{T \rightarrow \infty} e_T = 0, \quad \lim_{T \rightarrow \infty} p_T < \bar{k}(\omega, 0), \quad \lim_{T \rightarrow \infty} g_T = 0,$$

unless effort is exerted and beliefs remain above the threshold in all T , in which case

$$\lim_{T \rightarrow \infty} e_T = 1, \quad \lim_{T \rightarrow \infty} p_T \geq \bar{k}(\omega, 0), \quad g_T \in \{0, 2\omega\}.$$

Proof of Proposition 4. Consider an agent with initial circumstances denoted by ω . Recall that for any ω , $\bar{k}(\omega, 2\omega) < \bar{k}(\omega, 0)$, and $g_t = \theta_{t-1}$.

1) Suppose $\lim_{t \rightarrow T} e_t(p_{t-1}) = 1$, then it must mean that, for all t , $p_t(e_t) > \bar{k}(\omega, g_t)$. Then, there exists a subsequence $\{p_{k_t}\}$ which converges to a point p such that $\lim_{t \rightarrow T} p_t = p > \bar{k}(\omega, 0)$. Suppose that $p = \lim_{t \rightarrow T} p_t(e_t) > \bar{k}(\omega, 0)$. Then, there exists a period $t_n < T$ such that $p_t > \bar{k}(\omega, 0)$, for each $t > t_n$. But then it must be $e_t = 1$ for each $t > t_n$ and therefore $\lim_{t \rightarrow T} e_t = 1$. Hence, there is a long-term outcome with high effort ($e = 1$), optimistic beliefs $p > \bar{k}(\omega, 0)$ and fluctuating aspirations $g_t \in \{0, 2\omega\}$.

Suppose now that $\lim_{t \rightarrow T} e_t(p_{t-1}) = 0$. Then, it must mean that there exists a time $t_n < T$ such that $e_t = 0$ for each t . This implies that $g_t = 0$ and $p_t = p < \bar{k}(\omega, 0)$ for all $t_n < t \leq T$, and it must be $\lim_{t \rightarrow T} p_t(e_t) = p < \bar{k}(\omega, 0)$. Suppose that $\lim_{t \rightarrow T} p_t(e_t) = p < \bar{k}(\omega, 0)$. Then, there exists a time t_n for which $p_t < \bar{k}(\omega, 0)$, for each $t_n < t \leq T$, which implies $e_t = 0$ for all $t_n < t \leq T$ and so $\lim_{t \rightarrow T} e_t(p_{t-1}) = 0$. Hence, there is a long-term outcome with low effort ($e = 0$), low aspirations ($g_t = 0$) and pessimistic beliefs $p < \bar{k}(\omega, 0)$.

2) Suppose that $q > \bar{k}(\omega, 0)$. Then $q(v(2\omega) + r(2\omega - g_t)) - c > q(v(2\omega) + r(2\omega)) - c > 0$. The expected utility of long-term results with high effort is $E(U|e_T = 1) = q(v(2\omega) + r(2\omega - g_t))$, while the expected utility of long-term results with low effort is $E(U|e'_T = 0) = 0$. Hence, $q > \bar{k}(\omega, 0)$ implies $E(U|e_T = 1) > E(U|e'_T = 0)$.

3) Suppose $q = 1$. Assume $p_0 \geq \bar{k}(\omega, g_0)$. Then $e_1 = 1$. But then $p_t = \frac{p_{t-1} + 1}{2} > p_{t-1}$, for each $t > 1$, and so $p_t > p_0 > \bar{k}(\omega, g_0) \geq \bar{k}(\omega, 2\omega)$. Hence, the sequence converges to the long-term result of high-effort optimism and high aspirations since $\lim_{T \rightarrow \infty} e_T = 1$ and $\lim_{T \rightarrow \infty} p_T = 1 > \bar{k}(\omega, 2\omega)$ as $e_t = 1$ for each $t > 0$, $\lim_{T \rightarrow \infty} p_T = 1 = q$ similar to Proposition 1 and $g_t = 2\omega$, for each $t > 0$. On the other hand, if $p_0 < \bar{k}(\omega, g_0)$, $e_t = e_1 = 0$, for each t . Hence, the sequence converges to the low-effort pessimistic belief and the long-term result of low aspirations such as $\lim_{T \rightarrow \infty} e_T = 0$, $\lim_{T \rightarrow \infty} p_t = p_0 < \hat{k}(\omega)$ as $e_t = 0$ for each $t > 0$, and $g_t = 0$ for each $t > 0$.

Now suppose that $q = 0$. Then $g_{t+1} = 0$ and $p_t < p_{t-1}$ whenever $e_t = 1$. Therefore, there

exists a time t' such that $p_{t'} < \bar{k}(\omega, 0)$. Hence, the sequence will converge to the long-term result of low-effort pessimistic belief and low aspirations such as $\lim_{T \rightarrow \infty} e_T(p_{t-1}) = 0$, $e_t = 0$ for each $t > t'$, $\lim_{T \rightarrow \infty} p_T(e_t) = p'_t < \bar{k}(\omega, 0)$ for each $t > t'$ and $g_t = 0$ for each $t > t'$.

Finally, take the case where $0 < q < 1$. Suppose that there exists a t' in which $e_{t'} = 0$. Then it must mean $p_{t'-1} < \bar{k}(\omega, g_t) < \bar{k}(\omega, 0)$. Therefore, $p_t = p_{t'-1} < \bar{k}(\omega, 0)$, $e_t = 0$ and $g_t = g_T$ for each $t > t'$. Hence, $\lim_{T \rightarrow \infty} e_t = 1$ and $\lim_{t \rightarrow T} p_t(e_t) = p > \bar{k}(\omega, g_t)$ only if $e_t(p_{t-1}, g_t) = 1$ and $p_t(e_t, g_t) > \bar{k}(\omega, g_t)$ for all t .

Suppose $p = \lim_{T \rightarrow \infty} p_T(e_t) > \bar{k}(\omega, g_t)$ and there exists a t such that $p_t < \bar{k}(\omega, g_t) < \bar{k}(\omega, 0)$. But then, this would imply that $p_T = p_t < \bar{k}(\omega, 0)$ for each $T > t$, a contradiction. Therefore, if $p = \lim_{t \rightarrow T} p_t(e_t) > \bar{k}(\omega, 0)$, then $p_t > \bar{k}(\omega, 0)$ for all t , and so $e_t = 1$ for all t . Therefore, $\lim_{T \rightarrow \infty} e_t = 1$, $\lim_{T \rightarrow \infty} p_T > \bar{k}(\omega)$ if and only if $e_T = 1$, $p_T > \bar{k}(\omega, 0)$ for all $T > 0$. ■

As in Proposition 2 above, Proposition 4 shows that there are only two possible long-term outcomes. In the first, the agent exerts effort, maintains sufficiently optimistic beliefs, and maintains aspirations that may fluctuate over time. In the second, the agent stops exerting effort altogether, holds pessimistic beliefs, and settles into the lowest possible aspiration level. Effort, aspirations, and beliefs evolve together in a mutually strengthening way. When the agent exerts effort, the results may vary, causing the aspirations to fluctuate rather than converge. This dynamic can result in periods where aspirations are not fully met, even with effort. In contrast, once the agent stops exerting effort, both beliefs and aspirations stagnate at low levels permanently. The first long-term outcome welfare dominates the second whenever the objective probability q is sufficiently high (that is, $q > \bar{k}(\omega, 0)$), although it is also less stable.

As before, the stability properties of the two long-term outcomes depend on the value of q . Although intuition is similar to that of Proposition 2, it is worth noting that in the special case where $q = 1$, Proposition 4 exactly generalizes the failure of poverty and the aspirations result in Dalton et al. (2016). In this case, there are two regions separated by $\bar{k}(\omega, 0)$. If $p_0 < \bar{k}(\omega, 0)$, then $p_T = p_0$, $g_T = 0$, and $e_T = 0$ for all $T > 0$. In contrast, if $p_0 \geq \bar{k}(\omega, 0)$, then $\lim_{T \rightarrow \infty} p_T = 1$, $g_T = 2\omega$, and $e_T = 1$ for all $T > 0$. Thus, depending on the initial configuration of beliefs, an individual may either be trapped in a pessimistic, low-aspiration, low-effort state or escape to the welfare-dominant, high-effort outcome. Moreover, the likelihood of being stuck in the trap is higher, the more adverse their initial circumstances. The intuition for the general case where $0 < q < 1$ is similar is omitted here for brevity.

To formalize the idea of agents getting stuck in the undesirable long-term state described above, we define the notion of a *pessimistic, fatalistic, low-effort trap*.

Definition 2. Given ω and $T > 0$, a *pessimistic, fatalistic low-effort trap* is a long-term outcome (e_T, p_T, g_T) such that $e_T = 0$, $\lim_{T \rightarrow \infty} g_T = 0$, and $p_T < \bar{k}(\omega, 0) < q$.

In a pessimistic, fatalistic low-effort trap, the agent acts suboptimally by not exerting effort. Beliefs about the returns to high effort can only be updated - and aspirations can

avoid getting stuck at a low level (i.e., fatalism) — if the agent persistently exerts high effort. However, such a sustained effort requires sufficiently optimistic beliefs and the willingness to occasionally endure frustration when aspirations are not immediately met.

The second result of this section highlights how greater disadvantage increases the likelihood that an agent becomes trapped in a pessimistic, fatalistic, low-effort equilibrium.

Proposition 6. *The relationship between disadvantage, pessimism, and fatalism is as follows. Consider two agents i and j , where agent i is disadvantaged relative to agent j , that is, $\omega_i < \omega_j$. Given any $T > 0$, the probability that agent i becomes trapped in a pessimistic, fatalistic, low-effort trap is strictly greater than that of agent j .*

Proof of Proposition 5. Fix $p_{0,i} = p_{0,j} = q$, an initial common aspiration level g_0 . Then $\bar{k}(\omega_i, g_0) = \frac{c}{v(2\omega_i) + [2a\omega_i + g_0(b-a)]} \leq \frac{c}{v(2\omega_j) + [2a\omega_j + g_0(b-a)]} = \bar{k}(\omega_j, g_0)$, as $v(\cdot)$ is an increasing function. Each agent can reach one of the two long-term outcomes and, in particular, will enter a pessimistic belief low-effort trap as soon as their belief falls below the threshold, as in Proposition 3. Like in the case of Proposition 3, there can be different cases depending on where the prior belief lies with respect to the threshold for the disadvantaged agent. Clearly, if $\bar{k}(\omega_i, g_0) < p_0$, then it is clear that the disadvantaged agent will be stuck in pessimistic belief low-effort trap since $t = 0$. In this case, it is clear that the disadvantaged agent is (weakly) more likely to enter a pessimistic belief low-effort trap. Then, we can focus on the case in which $\bar{k}(\omega_j, g_0) < \bar{k}(\omega_i, g_0) < p_0$. Given that the two agents share the same initial prior, the same objective probability of success and the same level of aspiration, it is easy to see that, at any t ,

$$Pr(p_{t,i} \geq \bar{k}(\omega_i, g_t) | q) \geq Pr(p_{t,j} \geq \bar{k}(\omega_j, g_t) | q)$$

for any $g_t \in \{0, 2\omega\}$. ■

Hence, a more disadvantaged agent is more likely to enter a pessimistic belief low-effort trap.

The reasoning is similar to that noted for Proposition 3 above, with the additional observation that aspirations and beliefs interact in mutually complementary ways with effort choices; for brevity, we do not repeat the detailed argument here. Thus, the higher the level of disadvantage, the greater the likelihood that the individual will neither believe in nor aspire to their own potential.

Building on this, it is important to note that both frustration (when $g_t = g_0$ for all $t > 1$) and fatalism (when $g_t = \theta_{t-1}$) are long-term outcomes consistent with pessimistic beliefs and low effort. Hence, a lack of grit is closely related to both frustration and fatalism in the long run.

When aspirations are exogenous and lead to long-term frustration, interventions that boost grit (along the lines discussed in Section 3) have the added benefit of not only increasing persistence but also reducing frustration. In addition, exogenously raising aspirations can make it more likely that the agent will exert effort, although it may also increase the risk of long-term frustration if aspirations are not met.

When aspirations evolve over time, an intervention that increases grit by establishing a lower bound on beliefs - together with an intervention that maintains a high level of aspirations — can be particularly effective in sustaining effort over time. Formally, since $\bar{k}(\omega, g)$ decreases in both ω and g , an intervention that ensures that aspirations always remain high (so that $g_t = 2\omega$ for all $t \geq 0$) and sets a lower limit on beliefs ($p_t > \bar{k}(\omega, 2\omega)$ for all $t \geq 0$) will lead to welfare gains whenever $q > \bar{k}(\omega, 2\omega)$. Notice that this condition on q - that is, $q > \bar{k}(\omega, 2\omega)$ - is weaker than the condition $q > \bar{k}(\omega, 0)$ in Proposition 5. This is because, in Proposition 5, aspirations could fluctuate and potentially drop to zero, whereas here the intervention ensures that aspirations remain fixed at their highest possible value. In this case, the relevant threshold is $\bar{k}(\omega, 2\omega)$ rather than $\bar{k}(\omega, 0)$.

A concrete illustration of this idea is connecting the agent with a role model whose circumstances are only slightly more advantageous (so that $\omega_k - \omega < \epsilon$) and who maintains high aspirations ($g_{t,k} = 2\omega_k$). This interaction can help anchor the agent's own aspirations at a high level ($g_t = 2\omega$ for all $t \geq 0$). Under these conditions, the agent will continue to choose effort even after failures or setbacks. Over time, this guarantees convergence to the optimistic, high-effort long-term outcome, which welfare dominates the pessimistic, low-effort alternative. Although temporary episodes of frustration may occur (since aspirations are not always immediately achieved), these episodes act as motivators rather than leading to resignation. In this sense, interventions that maintain high aspirations and beliefs actively guard against *fatalism* — the state where individuals stop aspiring and stop trying altogether - which is especially detrimental to developing and maintaining grit.

References

Kahneman, Daniel and Amos Tversky. 1979. "Prospect theory: An analysis of decisions under risk". *Econometrica*, 47, 263-291.

Appendix D Disadvantage and the Objective Probability of Success

In this appendix, we study the implications of a world in which disadvantage negatively affects the objective probability of success but does not affect the returns to effort.

The objective probability of success, q , is now a function of disadvantage, $q(\omega)$, with the assumption that $\frac{\partial q}{\partial \omega} > 0$, that is, for the same level of effort, a higher level of disadvantage reduces the objective probability of success. As in the graph below, we do not impose any additional assumption on the shape of this relation. For simplicity, we will now assume that the returns to effort are independent of the level of disadvantage. This also implies that the threshold level for beliefs \hat{k} is independent of ω . Crucially, the decision rule has not changed, so an agent will exert effort (set $e = 1$) if $p_{t-1} \geq \hat{k} = \frac{c}{\bar{v}}$, where \bar{v} is the utility associated with success.

We can draw the counterpart of Figure 1 for this different version of the model.

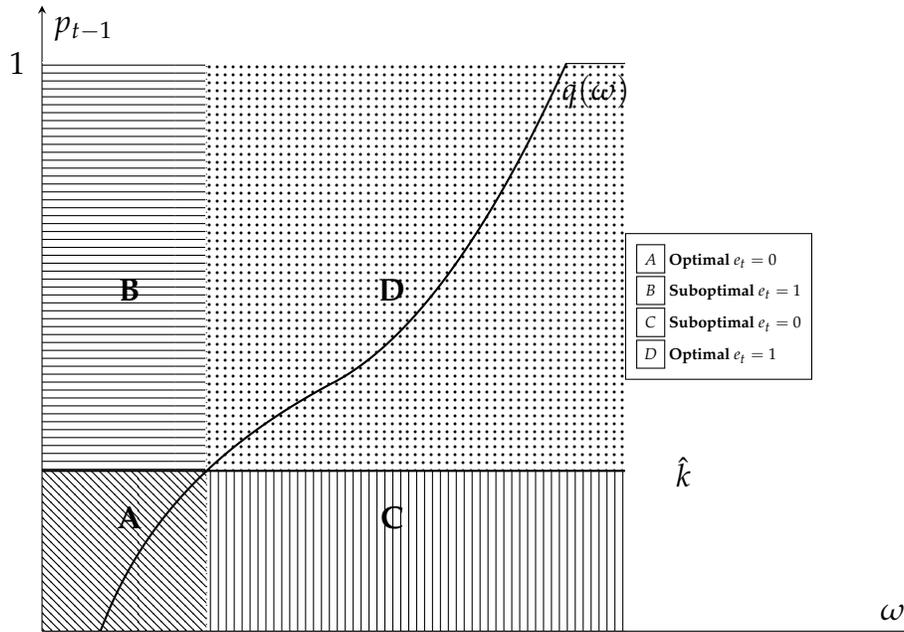


Figure 4: Short run decisions

In this case, the horizontal line will be the threshold \hat{k} as this does not depend on the disadvantage, while instead the curve is the objective probability of success, which is now a non-decreasing function of ω . The four areas identified in the graph mirror the ones in the main model, since the decision rule, which depends on the relative position between the threshold and the belief, and the optimality condition, which depends on the relative position of the threshold and the objective probability of success, have not changed. Therefore, regions A and C still characterize low effort exertion, as the belief is below the threshold, while in regions B and D an individual will be optimistic enough

to exert effort. The regions **A** and **B** define the areas in which the optimal effort level is 0, since q is below the threshold, while in the regions **C** and **D** the optimal effort is 1. Combining these two, we can deduce that in regions **A** and **D** the short term decision is the optimal one, while in regions **B** and **C** the short term decision is suboptimal.

Moving to long term outcomes, we can first see that there can exist thresholds $\underline{\omega}$ and $\bar{\omega}$, such that $q(\bar{\omega}) = 1$ and $q(\underline{\omega}) = 0$ and all agents with ω above $\bar{\omega}$ will have $q(\omega) = 1$ and all agents with ω below $\underline{\omega}$ will have $q(\omega) = 0$; because of the non decreasing assumption on the relation between q and ω , all agents with $\underline{\omega} < \omega < \bar{\omega}$, will have intermediate objective probability of success, i.e. $0 < q(\omega) < 1$. This means that we allow for extremely disadvantaged individuals who have no chance of succeeding and on the other extreme, extremely advantaged agents whose success is certain. We can immediately notice that for values of ω below $\underline{\omega}$ and above $\bar{\omega}$, the analysis from the previous model with $q = 0$ and $q = 1$ continues without changes. For extremely disadvantaged agents $\omega \leq \underline{\omega}$, it is never optimal to exert effort, and indeed they can only exert effort for a finite number of periods after which their belief will inevitably fall below the threshold. On the other hand, for an extremely advantaged agent $\omega \geq \bar{\omega}$, initial beliefs play a crucial role in their long term outcomes: if their initial belief is below the threshold, then they will be stuck in a trap; if instead they are moderately optimistic, then their beliefs will approach 1 and they will continue to exert effort. For intermediate levels of disadvantage, the same result applies: the only stable long term outcome is a trap. Therefore, we can restate Proposition 2 for this model.

Proposition 7. *Given ω , the long-run outcomes, their welfare ranking and stability is as follows:*

1. *Long-run outcomes: For all $T > 0$, there are two long-run outcomes, that is, (i) High effort ($e_T = 1$) and optimistic beliefs $p_T \geq \hat{k}$, and (ii) Low effort ($e_T = 0$) and pessimistic beliefs $p_T < \hat{k}$;*
2. *Welfare ranking: For all $T > 0$, if $q > \hat{k}$, the first long-term outcome with $e_T = 1, p_T > \hat{k}(\omega)$ welfare dominates the second long-term outcome with $e_T = 0, p_T < \hat{k}$, that is, $E(u|e_T = 1, q(\omega)) > E(u|e_T = 0, q(\omega))$.*
3. *Stability:*
 - *When $\omega \geq \bar{\omega}$ ($q(\omega) = 1$), if $p_0 \geq \hat{k}$, the sequence converges to the first long-term outcome (high effort, optimistic beliefs); if $p_0 < \hat{k}$, it converges to the second outcome (low effort, pessimistic beliefs).*
 - *When $\omega \leq \underline{\omega}$ ($q(\omega) = 0$), the sequence converges to the second long-term outcome (low effort, pessimistic beliefs).*
 - *When $\underline{\omega} < \omega < \bar{\omega}$ ($0 < q(\omega) < 1$), the second long-term outcome (low effort, pessimistic beliefs) is the unique absorbing state: $\lim_{T \rightarrow \infty} e_T = 0, \lim_{T \rightarrow \infty} p_T < \hat{k}$.*

Proposition 3 can also be restated as now more disadvantaged agents, by assumption, have a lower objective probability of succeeding, and therefore have a higher probability of their subjective beliefs falling below the threshold value $\hat{k} = \frac{c}{\bar{v}}$. Therefore, we can restate the previous Proposition 3 for this model.

Proposition 8. *The relationship between disadvantage and pessimism is as follows. Consider two agents i, j with intermediate levels of disadvantage where agent i is more disadvantaged than agent j , that is, $\underline{\omega} \leq \omega_i < \omega_j \leq \bar{\omega}$ and a common value of initial beliefs $p_{0,j} = p_{0,i}$. In a long-term outcome, for all $T > 0$, the likelihood that agent i is stuck in a pessimistic low-effort trap is greater than that of agent j .*

Once we established that the results of the main model are robust to this new specification, it is easier to see that the policy implications are also consistent. Given that the only long term outcome that is stable is the pessimistic low effort one, no policy alone, whether it is a reduction in the disadvantage or a reduction of the threshold, can revert this and induce agents to reach the optimistic high effort long term outcome; these interventions might mitigate the emergence of a pessimistic belief low-effort trap by delaying it and increase the expected periods in which agents are exerting effort. Therefore, also in this framework, there is a large scope for grit interventions to ensure that even disadvantaged agents continue to exert effort.

Moreover, in this framework it is clear how affirmative action can affect the relationship between disadvantage and the objective probability of success for more disadvantaged groups. Although we assumed a relationship between disadvantage and q , policies aimed at increasing the probability of success for highly disadvantaged agents can have a longer lasting impact, as long as the agents start optimistic enough. For example, diversity-based hiring policies might induce members of a disadvantage minority to maintain high effort as the chance of success increased, allowing the continuation of effort exertion by increasing the probability of keeping beliefs above the threshold.

Finally, note that all results would still hold even if we constructed a model in which disadvantage affected both the objective probability of success and the returns to effort, provided that the directions of these relationships are consistent with those in the main model and in this extension. In that case, more disadvantaged agents would be both more likely to fail and more likely, upon failure, to have their beliefs fall below the threshold.

Appendix E Prior over q

In this appendix, we assume that $q \in \{q_L, q_H\}$. Let P_0 denote the prior probability that $q = q_H$. We assume that this distribution is i.i.d. over time. For each $t > 0$, posterior beliefs are updated following the Bayes rule:

$$P_t(e_t) = \begin{cases} \frac{P_{t-1}q_H}{(1-P_{t-1})q_L + P_{t-1}q_H} & \text{if } e_t = 1 \text{ and } \theta_t = 2\omega \\ \frac{P_{t-1}(1-q_H)}{(1-P_{t-1})(1-q_L) + P_{t-1}(1-q_H)} & \text{if } e_t = 1 \text{ and } \theta_t = 0 \\ P_{t-1} & \text{if } e_t = 0 \end{cases}$$

As before, the agent has to exert effort to update their prior probability on q . We begin by proving the analog of Proposition 1, except that instead of convergence in expectations, conditional on $e_t = 1$ for all $t \geq 0$, we now obtain point-wise convergence to the true value of q :

Proposition 9. *If $0 < P_0 < 1$, conditional on $q = q_H$ (respectively, $q = q_L$) $\lim_{t \rightarrow \infty} E(P_t) = 1$ (respectively, $\lim_{t \rightarrow \infty} E(P_t) = 0$) if $e_t = 1$ for all $t \geq 1$.*

Proof of Proposition 8. Suppose without loss of generality $q = q_H$. There are two steps. First, consider the likelihood ratio $\frac{P(\theta_1, \dots, \theta_T | q_L)}{P(\theta_1, \dots, \theta_T | q_H)}$. We need to show that the likelihood ratio converges to zero in probability as T tends to infinity. Let $Z_T = \log\left(\frac{P(\theta_1, \dots, \theta_T | q_L)}{P(\theta_1, \dots, \theta_T | q_H)}\right) = \log\left(\prod_{t=1}^T \left(\frac{P(\theta_t | q_L)}{P(\theta_t | q_H)}\right)\right) = \sum_{t=1}^T (\log\left(\frac{P(\theta_t | q_L)}{P(\theta_t | q_H)}\right))$ (where we have used the independence property). Then, using Jensen's inequality, we have the following:

$$E\left(-\sum_{t=1}^T (\log\left(\frac{P(\theta_t | q_L)}{P(\theta_t | q_H)}\right))\right) = -\sum_{t=1}^T P(\theta_t | q_H) \log\left(\frac{P(\theta_t | q_L)}{P(\theta_t | q_H)}\right) \geq -\log\left(\sum_{t=1}^T P(\theta_t | q_H) \left(\frac{P(\theta_t | q_L)}{P(\theta_t | q_H)}\right)\right) = 0.$$

Hence, $E\left(-\sum_{t=1}^T (\log\left(\frac{P(\theta_t | q_L)}{P(\theta_t | q_H)}\right))\right) \geq 0$ with exact equality if $\frac{P(\theta_t | q_H)}{P(\theta_t | q_L)} = \frac{P(\theta_{t'} | q_H)}{P(\theta_{t'} | q_L)}$ for all $t, t' > 0$. It follows that $\lim_{T \rightarrow \infty} (P(-Z_T \geq c) = 1)$ for some $c > 0$, or equivalently, the likelihood ratio $\frac{P(\theta_1, \dots, \theta_T | q_L)}{P(\theta_1, \dots, \theta_T | q_H)}$ converges to zero in probability. So, we have verified that the likelihood function is a test function satisfying Schwartz's Theorem 3.5 (see Schwartz (1965)). Second, since q_H has a positive prior probability ($P_0 > 0$), both conditions of Schwartz's Theorem 3.5 are satisfied. It follows that P_t converges to 1. Symmetric results hold when $q = q_L$ and when $0 < P_0 < 1$. ■

The preferences remain unchanged. Therefore, following our short-term decision-protocol, an agent will choose to exert effort in the period t iff $E(u|e_t = 1, P_{t-1}) = P_{t-1}v(2\omega) - c \geq 0 = E(u|e_t = 0, P_{t-1})$. so that in a short-term decision problem, $e_t(P_{t-1}) = 1$ iff

$$P_{t-1} \geq \hat{k}(\omega) = \frac{c}{v(2\omega)}.$$

Suppose that $q_L < \hat{k} < q_H$ and $q = q_H$. Then, replacing using P_t instead of p_t throughout Propositions 2 and 3, it follows that an agent can get such in a pessimistic low effort trap

when $q = q_H$ and there will be multiple long-term outcomes ranked with welfare with the pessimistic low effort trap as the only absorbing state and the likelihood of getting such a trap for fine T is higher, the higher the degree of disadvantage of the agent. For completeness, we state the exact analog of the two Propositions. The proofs of the two Propositions are identical to those of Propositions 2 and 3 (simply replace p_0 and p_t with P_0 and P_t) and are omitted for brevity.

Proposition 10. *Suppose $q \in \{q_L, q_H\}$ with prior $0 < P_0 < 1$ and $q_L < \hat{k}(\omega) < q_H$. Given ω , and $q = q_H$, the long-term results, their welfare ranking, and stability are as follows. beginnenumerate element long-term outcomes: For all $T > 0$, there are two long-term outcomes, i.e., (i) High effort ($e_T = 1$) and optimistic beliefs $P_T \geq \hat{k}(\omega)$, and (ii) Low effort ($e_T = 0$) and pessimistic beliefs $P_T < \hat{k}(\omega)$; item welfare ranking: For all $T > 0$, the first long-term outcome with $e_T = 1, P_T > \hat{k}(\omega)$ welfare dominates the second long-term outcome with $e_T = 0, P_T < \hat{k}(\omega)$, that is, $E(u|e_T = 1, q) > E(u|e_T = 0, q)$.*

Stability:

- When $q_H = 1$, if $P_0 \geq \hat{k}(\omega)$, the sequence converges to the first long-term outcome (high effort, optimistic beliefs); if $P_0 < \hat{k}(\omega)$, it converges to the second outcome (low effort, pessimistic beliefs).
- When $0 < q_H < 1$, the second long-term outcome (low effort, pessimistic beliefs) is the unique absorbing state: $\lim_{T \rightarrow \infty} e_T = 0$, $\lim_{T \rightarrow \infty} P_T < \hat{k}(\omega)$.

Proposition 11. *The relationship between disadvantage and pessimism is as follows. Consider two agents i, j with intermediate levels of disadvantage where agent i is more disadvantaged than agent j , that is, $\omega_i < \omega_j$ and a common prior over $\{q_L, q_H\}$, $P_{0,i} = P_{0,j}$. In a long-term outcome, for all $T > 0$, the likelihood that agent i is stuck in a pessimistic low-effort trap is greater than that of agent j .*

It is perhaps useful to relate the above results to those obtained in one-armed bandit problems where decision-makers are Bayesian and far-sighted. The key result for our purposes is the one in which, once the known arm (in our case, low effort) is chosen, it is also chosen in all subsequent stages (Lemma 4.1 in Bradt 1956). This preserves the key feature of our short-term decision-protocol that once the decision-maker chooses low effort, then they do so in all subsequent periods of time, a key step in showing the existence of multiple welfare ranked long-term outcomes in Propositions 9 and 10 above.

However, with farsightedness, as before, there cannot be multiple, welfare ranked outcomes as in these two propositions, although the unique outcome could be welfare dominated because the agent is stuck with the wrong beliefs.

To see this, consider a T period model. Suppose $P_{T-1} < \hat{k}$. and Then, in the last period T , no effort will be exerted and $P_{T-1} = P_T$. It follows that either $P_{T-2} < \hat{k}$ or $\frac{P_{T-2}(1-q_H)}{(1-P_{T-2})(1-q_L)+P_{T-2}(1-q_H)} < \hat{k}$. Assume that $P_{T-2} < \hat{k}$. Then, again, given that $e_T = 0$, $e_{T-1} = 0$ and $P_{T-2} = P_{T-1} = P_T$. Proceeding iteratively backwards, it follows that if $P_0 < \hat{k}$ and $e_t = 0$ and $P_t = P_0$ for all $t = 1, \dots, T$. Hence, if $P_0 < \hat{k}$ or $\frac{P_0(1-q_H)}{(1-P_0)(1-q_L)+P_0(1-q_H)} < \hat{k}$, a fully farsighted agent can choose optimally not to exert effort all $0 < t < T$. In this

case, the agent will be stuck with wrong beliefs and the unique outcome will be welfare dominated if $q = q_H > \hat{k}$. Note that any intervention that changes P_0 above both threshold values will result in welfare improvements, and the earlier this is done, the higher the welfare gains.

References

Schwartz , Lorraine. 1965. "On Bayes Procedures." *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, Volume 4, pages 10–26, (1965)