

Control, Cost, and Confidence: Perseverance and Procrastination in the Face of Failure*

Inga Deimen

Julia Wirtz

University of Arizona

University of Bristol

February 2022

Abstract

We study effort provision and the development of the belief that effort matters over time: a student is uncertain whether she has control over success through her effort or whether success is determined by her innate ability, which she also does not know. In each period, what she can learn about her control and her ability depends on the level of effort she exerts. The student's optimal effort policy in this two-dimensional bandit problem takes the form of a linear belief cutoff rule and typically features repeated switching of the effort level. Moreover, we define perseverance and procrastination as indices for the student's behavior over time and analyze how they are affected by control, cost, and confidence. Finally, we relate our results to findings in educational psychology and discuss policies to foster perseverance and to lower procrastination.

JEL: D83, D81, D91, I21

Keywords: correlated bandit, optimal experimentation, perseverance, procrastination, locus of control, mindsets.

1 Introduction

The willingness to work hard even after experiencing setbacks is documented as a key predictor of success: According to Duckworth, Peterson, Matthews and Kelly (2007) grit is an

*Corresponding author: Inga Deimen: University of Arizona, Eller College of Management, Department of Economics, 1130 E. Helen St, Tucson, AZ 85721, United States, ideimen@arizona.edu; Julia Wirtz: University of Bristol, School of Economics, Priory Road Complex, Priory Road, Bristol, BS8 1TU, United Kingdom, j.wirtz@bristol.ac.uk. Declarations of interest: none.

essential determinant of high achievement, possibly even more important than talent.¹ Heckman, Stixrud and Urzua (2006) find that, for a given level of cognitive skills, non-cognitive skills such as perseverance strongly influence social and economic success.² Both in education and the workplace it is, therefore, essential to understand how students or workers can be motivated to exert high effort, even if they experience failures.

A central factor that seems to determine the willingness to work hard is the “control belief” – the belief that we have control over our own fate, or that we can achieve success if we work hard enough. The connection is simple: if we believe that we can achieve success eventually by exerting effort, we are less likely to give up after experiencing a failure.³ The control belief is at the core of the psychological concept of *mindset*: Dweck (2006) distinguishes subjects with a “fixed mindset” – who believe that success is based on innate ability – from subjects with a “growth mindset” – who believe that success comes from hard work. Consequently, when facing a failure fixed types stop exerting effort, whereas growth types increase their effort level. Closely related are the concepts of *locus of control*⁴ and *self-efficacy*⁵, which are also centered around the control belief. All three concepts consider the control belief to be a personality trait, that is, essentially, an exogenously given constant. Therefore, interventions aiming at increasing effort often try to directly enhance the students’ control beliefs.⁶

In this paper, we try to expand the set of possible policies to induce high effort by considering the control belief as a variable that evolves over time. In our analysis, we take the view of a student who decides whether to work hard or not. The key assumption is that the student is uncertain if high effort is decisive for success. As she experiences successes and failures, the student learns about how much her effort matters. We show how the development of the control belief over time depends on the student’s confidence in her own ability, as well as her effort cost, patience, and luck. Our analysis provides insights on possible policies to promote effort, both in the short and in the long run.

¹Duckworth et al. (2007) develop a survey-based measure for grit. They find that it is a significant predictor for success, even when controlling for factors like IQ.

²Heckman et al. (2006) report that early childhood programs such as the Perry Preschool Program or Headstart had positive impact on life outcomes through non-cognitive skills, such as perseverance, while they did not have an effect on IQ.

³For empirical evidence in a school context, see, for example, Coleman and DeLeire (2003). Caliendo, Cobb-Clark and Uhlendorff (2015) analyze effort in job search.

⁴According to Rotter (1966), subjects with an “external locus of control” see success and failure beyond their reach, whereas subjects with an “internal locus of control” attribute outcomes more to their own behavior. Thus, after failure, an internal locus of control leads to persistence of high effort, while an external locus results in despondence.

⁵Bandura (1997) holds that the extent to which people believe they have control over their lives determines their resilience in the face of difficulties.

⁶Dweck, in particular, has conducted several interventions in schools, to promote a high control belief in students (for a summary, see Dweck (2000)).

We design a simple theoretical model in which a student is uncertain about both the efficacy of effort and her own ability level: she does not know if success can be achieved through effort or if it is all down to innate ability. Moreover, if ability is decisive, she does not know if her ability is sufficient to succeed. This uncertainty translates to three possible states of the world: (α) success is down to ability and the student has high ability, (β) success is down to ability, but the student has low ability, or (ω) success can be achieved through high effort. There is always a chance of failure due to bad luck. The belief that the state is ω – i.e., that the student can affect the outcome through effort – is interpreted as her control belief. The student’s belief that ability matters and that she has high ability (α) rather than low ability (β), is interpreted as her confidence.

The student faces a two-dimensional correlated bandit problem in which she simultaneously learns about her ability and control. Each period, the student decides whether to exert costly high effort or not and, subsequently, observes a success or a failure. Different levels of effort, combined with success or failure, result in different posterior beliefs about the state. For example, a student who succeeds with low effort would be certain that ability matters and she is a high type (α); by contrast, a student who observes a failure with low effort does not know whether she failed because high effort is necessary (ω), because she is a low type (β), or because she is a high type and had bad luck (α). For a given control belief, the student’s confidence determines how she updates her control belief. This is particularly interesting in our setting, in which ability is irrelevant when effort matters and *vice versa*. Thus, two students with identical prior control beliefs and identical experienced histories can draw very different conclusions, depending on their confidence. A student with low confidence tends to attribute failure with high effort to her lack of ability, while a student with high confidence tends to blame bad luck.

Interestingly, we find that a simple stopping rule is not generally optimal: the student may change her effort level repeatedly, despite the fact that the state is fixed. Intuitively, after experiencing failure with low effort, the student becomes more convinced that high effort is necessary for success. Conversely, after experiencing a failure with high effort, the student becomes more convinced that effort was futile. Due to repeated switching of the effort level, the value function is a complex expression which depends on the precise sequence of high and low effort and which does not have a closed form solution. Thus, we are not able to solve the Bellman equation to find the value function and derive the optimal behavior. Instead, we derive a linear belief cutoff rule that characterizes the student’s effort choices in each period, i.e., the optimal Markov policy. The cutoff rule is derived from comparing two *pivotal plans* that prescribe switching the effort level immediately after observing any failure.

With uncertainty about the state, the optimal policy achieves the maximum expected

welfare. Thus, the student’s behavior is optimal, given her belief. For our welfare analysis, motivated by the educational psychology research cited above, we take the perspective of a teacher who knows that the state is ω and that the student can only be successful if she exerts high effort. From this perspective, the student suffers due to her uncertainty. Since what she can learn depends on her effort, learning may be delayed or she may never learn the true state. We characterize the student’s behavior over time using two indices that capture this loss of welfare due to uncertainty: procrastination and perseverance. We define *procrastination* as the tally of consecutive failures with low effort that it takes before the student is willing to exert high effort.⁷ Procrastination is a (negative) measure for short-run welfare. The longer it takes for the student to start exerting high effort, the longer she delays the possibility for success. The student may experiment with high and low effort for some time, but will eventually stop exerting high effort if she does not observe a success. We define *perseverance* as the overall tally of failures with high effort that it takes before the student gives up forever. Perseverance is thus a measure for long-run welfare. Intuitively, the longer the student tries to succeed with high effort, the more likely she eventually observes a success. Thus, to improve welfare we are interested in policies that decrease the student’s procrastination and that increase her perseverance.

In accordance with the psychological literature mentioned above, our model predicts that a higher control belief encourages effort. Our model complements and extends this literature since we are able to describe how the student’s beliefs and behavior change over time and how this is affected by the parameters of the model. We find that high effort costs lead to less perseverance and more procrastination. Equivalently, a student who assigns a lower payoff to success is less willing to work hard. We also predict that a student is more persevering if she is more patient. While these results are intuitive, the role of confidence is more surprising. The effect of confidence on effort crucially depends on the time-frame considered. In the short run, higher confidence can increase the student’s procrastination. In the long run, higher confidence can increase the student’s perseverance. These effects are due to the interdependence of learning about confidence and control.

From a theoretical perspective, the problem studied in our paper is a two-armed bandit problem: the student can choose between high and low effort which corresponds to choosing between two arms of a bandit. The probability of success from both high and low effort is uncertain. By choosing an effort level, the student receives both an immediate payoff and information about the payoff distribution which is valuable for her future actions. Consequently, the student faces the classical trade-off between exploitation of the option with the

⁷We use procrastination literally, signifying the delay of effort. While procrastination is often understood as a dysfunctional behavior or trait, we model it as a rational reaction to learning about the efficacy of effort. For a discussion, see Kim and Seo (2015).

highest expected immediate payoff and exploration of the other option.

Rothschild (1974) shows that an agent choosing between two arms with unknown (and possibly correlated) payoff distributions will not necessarily learn the optimal action. This central insight on bandit problems also applies to our setting. A student may stop exerting high effort for good – even though the state is ω . Gittins (1979) develops a simple index which characterizes the optimal policy for bandits with independent arms. The Gittins index is not applicable to our model, as the arms are correlated: the outcome with high effort also affects the student’s belief about the success probability with low effort and *vice versa*. For bandit problems with correlated arms, there is no general analytical solution for the optimal policy, and the existing theoretical work is limited.

Camargo (2007) is able to show the optimality of cutoff belief strategies for a class of correlated two-armed bandits that satisfy some mild continuity assumption. In particular, his characterization of optimal policies for two arms extends the classical result by Gittins (1979) to setups with finite horizon, non-geometric discounting, and correlated arms. By focusing on a specific information structure, we are able to give an explicit solution for the optimal cutoff belief strategy and thereby fully characterize the agent’s optimal policy. This allows us to derive a formal notion of procrastination and perseverance and to analyze how these are affected by changes in the different parameters of our model.

Klein and Rady (2011) consider the case of negatively correlated arms that are operated by two different agents. While we consider a single-agent decision problem, they analyze the strategic interaction between the agents.⁸ For perfectly negative correlation, they characterize all Markov perfect equilibrium strategies in closed form. The extension to imperfectly correlated arms is particularly related to our setting, even though the correlation between the arms is different. As in our case, Klein and Rady show existence of an equilibrium in cutoff strategies and that learning may remain incomplete. Pastorino (2019) provides an application of bandits with correlated arms. The paper analyzes a structural model of job and wage mobility in a firm and finds evidence for a novel mechanism through which learning about ability shapes wages.

Our paper connects to a literature featuring interdependent learning on two dimensions: the belief about one variable affects learning about another. Heidhues, Kőszegi and Strack (2018) consider misspecified beliefs: The agent’s outcome depends on her action, her ability, and some external factor. The agent is uncertain about the external factor, while she is certain, but incorrect about her ability. The agent’s overconfidence in her ability leads to incorrect inferences about the environment, and the agent’s adjusted actions lead her further away from the truth. In contrast to model misspecification, our paper features

⁸Starting with Bolton and Harris (1999) and Keller, Rady and Cripps (2005) there has been a rich economic literature on strategic experimentation in which arms are operated by multiple agents.

model uncertainty: agents learn about two variables at the same time. The agent is both uncertain about her ability and about the efficacy of her effort. The student’s beliefs are not misspecified as she attaches a positive probability to the true state. However, she might never learn the true production function, since what she can learn depends on her actions. Model uncertainty is also featured in Hestermann and Le Yaouanq (2021) and Piketty (1995).

In Hestermann and Le Yaouanq (2021), the agent may start with an incorrect prior belief about her ability. Similarly to our paper, the agent faces two-dimensional uncertainty and the ability belief influences learning in the other dimension. In a series of two-armed bandit problems, she can choose to take a new draw of the exogenous variable – i.e., sample a new environment – or stick to the current one. While an overconfident agent will eventually learn her true ability as she keeps sampling new environments, an underconfident individual may get stuck in a sub-optimal environment and never learn. This outcome corresponds to our finding that a student with low confidence in her ability may never learn that she can succeed with effort. There is no effort choice in Hestermann and Le Yaouanq (2021) and, thus, the agent’s problem is different from ours. We provide a characterization of the optimal strategy with forward-looking agents.

In Piketty (1995), as in our model, the agents are uncertain about the efficacy of effort, and the level of effort they exert determines the information they can acquire. Agents with low control beliefs never strive and never learn that exerting high effort would be valuable, while others – with high control beliefs – work unnecessarily hard. The main difference to our approach is that in Piketty (1995) agents only live one period and, thus, behave myopically in their effort choice. Instead, our paper solves a bandit problem: the agent takes the information rent and the expected future payoffs into account, when deciding about the effort level.

Lemoine (2020) considers agents who are uncertain about their ability and how their efforts translate into output. Differently from our model, effort and ability are complementary. Surprisingly, agents become either persistently overconfident or persistently underconfident on average, depending on how greater effort affects the variance of an agent’s payoffs.

Our paper also relates to a small literature in information acquisition and learning: Che and Mierendorff (2019), Nikandrova and Pancs (2018), and Mayskaya (2019) consider the choice between two actions that give different information about the state of the world. Differently from bandit problems, here, the agent does not receive any payoff until she takes a final decision.

We proceed as follows. We state our model in Section 2. Section 3 illustrates an example of two periods. In Section 4, we derive the student’s optimal policy. We then describe the sets of absorbing beliefs and introduce perseverance and procrastination as indices for the student’s behavior over time. At the end of the section, we describe how the parameters

of the model affect perseverance and procrastination and discuss possible implications for policy. A final section concludes. All proofs are in the appendix.

2 Model

We consider a single agent decision problem over time. Time is discrete, and we discount the output of future periods with discount factor $\delta \in (0, 1)$. In each period, the student chooses a level of effort $e \in \{H, L\}$. Exerting high effort, $e = H$, has cost $c \in (0, 1 - \varepsilon)$ for some $\varepsilon \in (0, 1)$, exerting low effort is costless.⁹ In each period, after choosing the level of effort, the student observes either a success S or a failure F . A success gives the student a payoff of 1, while a failure yields a payoff of 0.¹⁰

The state of the world takes a value in $\{\alpha, \beta, \omega\}$ (high ability, low ability, work) and is fixed for all periods. The student's choice of effort determines the probability of success for each state. In particular, if the state is ω (work), the student can only succeed with high effort, independent of her ability: the probability of success is $1 - \varepsilon$ if she exerts high effort, and 0 if she exerts low effort. On the other hand, if the state is α or β (high or low ability), the probability of success is independent of the level of effort. In state α , the probability of success is $1 - \varepsilon$, while in state β , the student cannot succeed. The parameter $\varepsilon \in (0, 1)$ denotes the probability of having bad luck – i.e., with probability ε , the student observes a failure, even though she had all the prerequisites for success. The student is uncertain about the realization of the state, and the triple $\mathbf{p} = (p_\alpha, p_\beta, p_\omega) \in [0, 1]^3$ with $p_\alpha + p_\beta + p_\omega = 1$ specifies the student's belief about the state. We interpret p_ω as the *control belief*, p_α as the *confidence* of the student in her own ability, and p_β as the belief that she has insufficient ability. The student's belief can be represented in a triangle as seen in Figure 1.

It proves convenient to specify the *outcome* of a period $h \in \{SH, SL, FH, FL\}$ as a combination of the choice of high effort H or low effort L and the observed success S or failure F in that period. The respective probabilities of these outcomes are given in Table 1. An outcome h , together with a belief $\mathbf{p} = (p_\alpha, p_\beta, p_\omega)$, induces a posterior belief $\mathbf{p}(h) = (p_\alpha(h), p_\beta(h), p_\omega(h))$. The student's updating of her belief follows Bayes rule. Her posterior beliefs after the different outcomes are given in Table 2.

In our setup, the student always updates her belief about one state relative to the other two states. In particular, with low effort, the student updates about state α relative to the alternatives β and ω . After observing a success with low effort SL , the student learns that she has high ability and succeeds because of that. Her belief jumps directly to $p_\alpha = 1$. After

⁹The assumption that c is bounded by $1 - \varepsilon$ ensures that exerting effort may be valuable for the student.

¹⁰We normalize the payoff to success to 1, but allow heterogeneity in effort cost c . Equivalently, we could normalize the cost and allow the payoff to vary, or both.




Figure 1: Belief triangle for $\mathbf{p} = (p_\alpha, p_\beta, p_\omega)$.

	α	β	ω
$Pr(SH)$	$1 - \varepsilon$	0	$1 - \varepsilon$
$Pr(SL)$	$1 - \varepsilon$	0	0
$Pr(FH)$	ε	1	ε
$Pr(FL)$	ε	1	1

Table 1: Probabilities of outcomes for the states.

observing a failure with low effort FL , the student cannot distinguish whether the reason was a lack of effort or a lack of ability. In terms of beliefs, p_α is reduced, while p_ω and p_β increase at the same rate. As the belief ratio of $\frac{p_\omega}{p_\beta}$ remains constant, the student's belief moves on a straight line between the point $p_\alpha = 1$ and the edge between $p_\omega = 1$ and $p_\beta = 1$. The student's updating after exerting low effort is illustrated by the solid (blue) arrows in Figure 2.

With high effort, the student updates about state β relative to the alternatives α and ω . As the belief ratio of $\frac{p_\omega}{p_\alpha}$ remains constant, the student's belief moves on a straight line between the point $p_\beta = 1$ and the edge between $p_\omega = 1$ and $p_\alpha = 1$. After observing a failure with high effort FH , the probability that her ability is insufficient to succeed increases. Her belief moves towards $p_\beta = 1$. After observing a success with high effort SH , the student cannot distinguish whether the reason was high effort or high ability. In terms of beliefs, p_ω and p_α increase at the same rate, p_β is reduced to zero, and the belief jumps to the edge between $p_\omega = 1$ and $p_\alpha = 1$. The student's updating after exerting high effort is illustrated by the dashed (red) arrows in Figure 2.

Note that a student's confidence affects how she updates her control belief for a given outcome. After any failure, a student with high confidence has a higher posterior control belief than a student with low confidence. For any fixed p_ω , the student's posterior be-

$\mathbf{p}(h)$	$p_\alpha(h)$	$p_\beta(h)$	$p_\omega(h)$
$\mathbf{p}(SH)$	$\frac{p_\alpha}{p_\alpha+p_\omega}$	0	$\frac{p_\omega}{p_\alpha+p_\omega}$
$\mathbf{p}(SL)$	1	0	0
$\mathbf{p}(FH)$	$\frac{p_\alpha\varepsilon}{p_\alpha\varepsilon+p_\beta+p_\omega\varepsilon}$	$\frac{p_\beta}{p_\alpha\varepsilon+p_\beta+p_\omega\varepsilon}$	$\frac{p_\omega\varepsilon}{p_\alpha\varepsilon+p_\beta+p_\omega\varepsilon}$
$\mathbf{p}(FL)$	$\frac{p_\alpha\varepsilon}{p_\alpha\varepsilon+p_\beta+p_\omega}$	$\frac{p_\beta}{p_\alpha\varepsilon+p_\beta+p_\omega}$	$\frac{p_\omega}{p_\alpha\varepsilon+p_\beta+p_\omega}$

Table 2: Posterior beliefs after observing outcomes.




Figure 2: Updating beliefs after high (dashed red) and low (solid blue) effort.

beliefs about control after failure (Table 2) can be rewritten as $p_\omega(FH) = \frac{p_\omega\varepsilon}{1-(1-\varepsilon)(p_\alpha+p_\omega)}$ and $p_\omega(FL) = \frac{p_\omega}{1-(1-\varepsilon)p_\alpha}$. Both are increasing in confidence p_α . The reason is that high ability reduces the unconditional probability of failure, which in turn increases the posterior control belief after failure.¹¹

Confidence also affects how much information each effort level provides. The effect is easiest to see for extreme beliefs. Assume a student is certain that, if ability matters, her ability is high – i.e., $p_\beta = 0$. In this case, the student learns nothing from high effort: her belief remains unchanged after observing SH or FH . Low effort, instead, moves her belief towards $p_\omega = 1$ for FL or makes it jump directly to $p_\alpha = 1$ for SL . Conversely, a student with $p_\alpha = 0$ learns nothing from exerting low effort.

Discussion of the assumptions. We assume binary effort and ability levels where ability does not matter if effort is decisive and *vice versa*. An alternative interpretation – which is equally consistent with the model – is that there are three different ability levels: at the lowest level (β), the student never succeeds, independently of effort; at the intermediate level of ability (ω), the student can succeed if and only if she exerts high effort; at the

¹¹The feature that, for a given prior control belief, the posterior control belief after failure increases in confidence is robust to any production function for which high ability decreases the probability of failure.

highest level (α), the student can succeed, independently of effort. In this interpretation, comparing across students, effort and ability can be seen as substitutes: some students may be in state α and can succeed on their ability alone, while others may be in state ω and can substitute effort for ability. As ability increases from β to ω , effort and ability act as complements: effort only pays off if the ability is high enough. As ability increases further from ω to α , effort becomes unnecessary, so effort and ability are no longer complements. In general, results will be similar for any model in which the return to effort depends on ability non-monotonically – i.e., if effort and ability are not always complementary.¹²

As is standard in the literature, we only model one type of noise – bad luck. As a result, the student learns immediately that the state is α upon observing SL and that the state cannot be β upon observing SH . We expect symmetrically similar result when replacing bad luck by good luck. In practice, it seems more reasonable to assume that a student may fail despite high effort and high ability, than that a student could succeed despite low effort and low ability.

Finally, we assume that the probability of bad luck ε is the same in the states ω and α . Different probabilities would allow for an additional channel of learning: the student could learn about the relative probabilities of the state by keeping her effort constant. In particular, after successfully exerting high effort for sufficiently long, the student could infer from different values of ε whether the state is ω or α just from observing the failure likelihood.

3 Illustrative example

To illustrate the basic trade-offs at work, we consider simplified one- and two-period versions of our model without discounting ($\delta = 1$).

One period: The choice of effort only impacts the immediate payoff but yields no informational gain. The student’s expected payoff from exerting high effort is $EU(H) = (p_\omega + p_\alpha)(1 - \varepsilon) - c$, while the expected payoff from low effort is $EU(L) = p_\alpha(1 - \varepsilon)$. Thus, the student prefers to exert high effort if and only if $p_\omega(1 - \varepsilon) \geq c$. The cost of high effort c is weighed against the expected benefit $p_\omega(1 - \varepsilon)$, that results from the probability of effort being necessary to get a success, p_ω , and not having bad luck, $1 - \varepsilon$.

Two periods: A plan $\pi = (\pi_1(\emptyset); \pi_2(h))$ specifies an action for both periods; the action in period two depends on the observed outcome h in period one. We solve the student’s decision

¹²In practice, it seems reasonable to assume that the return to effort is not always increasing in ability – especially, for any non-continuous payoff scheme, e.g., for pass/ fail exams or discrete grading schemes. When a student crosses a payoff threshold, a small increase in effort does not increase her payoff further.

problem by backward induction. In the last period, the student's problem is identical to the one-period case: she exerts high effort whenever her (posterior) belief $p_\omega(h)$ satisfies $p_\omega(h)(1 - \varepsilon) \geq c$. The choice of effort in the first period is more complex, since it affects not only the immediate payoff but also the information available in the second period.

Lemma 1 *For two periods, one of the following plans is optimal, depending on the student's prior \mathbf{p} .*

1. *Persevering: plan $(\pi_1(\emptyset) = H; \pi_2(SH) = H, \pi_2(FH) = H)$ is optimal for $p_\omega \geq \max\{c \frac{p_\alpha}{1-\varepsilon-c}, \frac{c}{c+\varepsilon} (\frac{1}{1-\varepsilon} - p_\alpha)\}$.*
2. *Despondent: plan $(\pi_1(\emptyset) = H; \pi_2(SH) = H, \pi_2(FH) = L)$ is optimal for $\frac{c}{c+\varepsilon} (\frac{1}{1-\varepsilon} - p_\alpha) > p_\omega \geq \max\{\frac{2cp_\alpha}{1-\varepsilon-c}, \frac{c}{2-\varepsilon-c} (\frac{1}{1-\varepsilon} + p_\alpha)\}$.*
3. *Responsive: plan $(\pi_1(\emptyset) = L; \pi_2(SL) = L, \pi_2(FL) = H)$ is optimal for $\min\{c (\frac{1}{1-\varepsilon} + p_\alpha), \frac{2cp_\alpha}{1-\varepsilon-c}\} > p_\omega \geq c (\frac{1}{1-\varepsilon} - p_\alpha)$.*
4. *Fatalist: plan $(\pi_1(\emptyset) = L; \pi_2(SL) = L, \pi_2(FL) = L)$ is optimal for $p_\omega < c (\frac{1}{1-\varepsilon} - p_\alpha)$.*

We find four optimal plans: For a high control belief p_ω , the student exerts high effort in both periods, regardless of the outcome of the first period (persevering). For a medium control belief p_ω and low confidence p_α , the student exerts high effort in $t = 1$. After observing a success, she continues with high effort in $t = 2$, while she switches to low effort after a failure (despondent). For a medium control belief p_ω and high confidence p_α , the student exerts low effort in $t = 1$. After observing a success, she continues with low effort in $t = 2$, while she switches to high effort after a failure (responsive). For a low control belief p_ω , the student exerts low effort in either period, regardless of the outcome of the first period (fatalist). We illustrate the different parameter regions described in the lemma in the left panel of Figure 3.

The right panel of Figure 3 shows the role of information for decision making in period 1 compared to period 2. The gray area indicates beliefs where the student exerts low effort in period 1, while the dotted area indicates low effort in period 2.¹³ Thus, in the white dotted area, the student exerts high effort in period one, but not in period 2. While the cost of high effort is higher than the immediate benefit, in this area, the information that can be gained makes high effort worthwhile: a success allows the student to exclude state β ; this information is particularly valuable if the confidence p_α is relatively low. Meanwhile, in the gray area without dots, the student chooses low effort in period one and high effort in period 2. In this area, the cost of high effort is lower than the immediate benefit, but experimenting

¹³Note that the beliefs are updated after observing the outcome of period 1. Thus, the beliefs change from period 1 to period 2 as illustrated in Figure 2.




Figure 3: Left: optimal plans depending on initial belief. Right: belief thresholds for effort in periods one and two. Both for $c = \frac{1}{4}$ and $\varepsilon = \frac{1}{4}$.

with low effort is worthwhile for the information gain: a success implies that the state is α ; this information is particularly valuable if the confidence p_α is relatively high.

To see the interplay between the student's confidence and control, note that the level of effort in period 1 can be lower for higher confidence p_α . In particular, increasing her confidence can move the student from *despondent* to *responsive*. The reason is that a student with high confidence and a medium control belief prefers to experiment with low effort in the hope of succeeding and learning that the state is α . In contrast, effort in period 2 can be higher for higher confidence p_α . Increasing her confidence can move the student from *despondent* to *persevering* or from *fatalist* to *responsive*. As discussed in the previous section, this is because a student with higher confidence has a higher posterior control belief after a failure and in period 2 only the control belief matters for the choice of effort.

The optimal plans for the two period example illustrate some general features of our model. Whenever the student observes a success, she sticks to her choice of effort for the remaining period. After observing SL , she is fully informed that low effort is optimal. After observing SH , the student continues to exert high effort. Assume she preferred to exert low effort after SH – i.e., $p_\omega(SH)(1 - \varepsilon) < c$. Then she would also exert low effort after FH , since $p_\omega(FH) < p_\omega(SH)$. However, if she does not condition her effort in period two on the outcome of period one, she does not gain valuable information from high effort in period one. This can only be the case when she purely cares about her immediate payoff. Hence she prefers to exert low effort in period one, since $p_\omega(1 - \varepsilon) < p_\omega(SH)(1 - \varepsilon) < c$, but this precludes observing SH .

After a failure, the student may optimally change her effort level. A student who starts with high (low) effort may switch to low (high) effort after a failure, since it is now less (more) likely that the state is ω . A student with low control belief p_ω never exerts high

effort, even after observing FL , and never learns if the state is actually ω and high effort would be worthwhile.

4 Infinite horizon

When the time horizon is infinite, the student's decision problem becomes stationary, and her effort choice is only determined by her current belief. The student's belief follows the rules outlined in Table 2, from any period to the next. A *history of outcomes* \mathbf{h} is a vector with entries $h \in \{SH, SL, FH, FL\}$. We denote by $n_h = n_h(\mathbf{h})$ the *tally of outcome* h in history \mathbf{h} . Notably, the student's posterior belief only depends on the tally of observed outcomes, but not on their order. For example, the student's posterior belief after observing FL and FH is identical to the student's posterior after observing FH and FL . Any history of outcomes \mathbf{h} can thus be expressed as the vector $(n_{SH}, n_{SL}, n_{FH}, n_{FL})$. The student's belief after observing history \mathbf{h} is denoted by $\mathbf{p}(n_{SH}, n_{SL}, n_{FH}, n_{FL})$.

A *Markov policy* $m(\mathbf{p})$ is a mapping $m : \{[0, 1]^3 : p_\omega + p_\alpha + p_\beta = 1\} \rightarrow \{H, L\}$ that specifies an effort level e for any belief \mathbf{p} . To characterize the optimal Markov policy, it proves useful to introduce the notion of a plan. A *plan* π is an infinite binary tree in which the nodes are the student's actions and the edges are the possible outcomes given the preceding action. The plan pre-specifies which action the student is going to take at each possible node. Given the student's action, nature chooses the outcome.¹⁴

A plan π is called *optimal at belief* \mathbf{p} if the student's expected payoff for π is higher than for any alternative plan π' at a given belief \mathbf{p} : $EU[\pi, \mathbf{p}] \geq EU[\pi', \mathbf{p}]$. Without loss of generality, we assume that the student chooses high effort $e = H$ whenever she is indifferent between exerting high and low effort. As a result, the optimal plan at belief \mathbf{p} is unique. An *optimal Markov policy* is consistent with the optimal plan at belief \mathbf{p} in the following sense: the action that an optimal Markov policy prescribes for belief \mathbf{p} is identical to the initial action in the plan that is optimal at this belief. Moreover, at any posterior belief the same must be true: the optimal policy must prescribe exactly the same action that the optimal plan dictates for any node that is reached. Since the optimal plan for each belief is unique, the optimal Markov policy is also unique.

4.1 Optimal policy

To find the optimal Markov policy, we have to compare infinitely many possible plans for each belief \mathbf{p} . The classical approach to this kind of problem involves analyzing the Bellman

¹⁴For example, a plan may specify high effort in the initial period and to continue forever after a success, but to switch to low effort forever after a failure.

equation to derive the value function and the optimal policy. We cannot solve the Bellman equation in our setting, because there is no closed form solution for the value function: Since the optimal policy does not invoke a simple stopping rule, the value function is an expression that depends on the exact sequence of high and low effort for all possible histories. In particular the student's optimal effort can switch back and forth between low and high effort, as already seen in the example. Moreover, the number of switches can be arbitrarily large depending on the parameter realizations. The value function would thus need to include the payoffs from an arbitrarily large number of possible histories.

We, therefore, develop an alternative method to obtain the optimal policy. The key insight is that it is sufficient to compare two *pivotal plans*. These plans pin down the belief threshold that determines the optimal effort choice and, thus, the optimal policy. We apply the principle of optimality and show that there are no profitable one-shot deviations. The use of pivotal plans to derive the optimal policy without recourse to the Bellman equation is, to the best of our knowledge, new.

The two pivotal plans π_H and π_L prescribe immediate switching of the effort level after experiencing a failure; after a success, the effort level is kept unchanged forever for both plans. After a failure with both high and low effort, both plans prescribe following an arbitrary fixed continuation plan $\bar{\pi}$. For an illustration, see Figure 4.

π_H : Choose H . After observing SH , play H forever. After observing FH , play L . After observing SL , play L forever. After observing FL , follow the plan $\bar{\pi}$.

π_L : Choose L . After observing SL , play L forever. After observing FL , play H . After observing SH , play H forever. After observing FH , follow the plan $\bar{\pi}$.




Figure 4: Pivotal plans π_H (left) and π_L (right).

Lemma 2 For any continuation plan $\bar{\pi}$, the student prefers plan π_H to π_L if and only if $p_\omega \geq \bar{p}_\omega(p_\alpha)$ with

$$\bar{p}_\omega(p_\alpha) := \frac{c}{(1 - \delta\varepsilon - \delta c)} \left[\frac{1 - \delta}{1 - \varepsilon} + p_\alpha \frac{\delta(2 - \delta - \delta\varepsilon)}{1 - \delta} \right]. \quad (1)$$

The derivation of $\bar{p}_\omega(p_\alpha)$ follows from equating the expected payoffs of the pivotal plans π_H and π_L and solving for p_ω . Note that the choice of the continuation plan $\bar{\pi}$ does not affect the comparison of the payoffs from π_H and π_L . In both plans, the student reaches the decision node of the continuation plan after observing FH and FL . The probability of this event and the student's posterior beliefs are independent of the order of observations. We show that the student optimally exerts high effort if and only if she prefers π_H to π_L , and, thus, if and only if her control belief p_ω is above the threshold \bar{p}_ω :

Proposition 1 *The optimal Markov policy has the form*

$$m^*(\mathbf{p}) = \begin{cases} H & \text{for } p_\omega \geq \bar{p}_\omega(p_\alpha) \\ L & \text{otherwise.} \end{cases}$$

In the appendix, we formally show that the threshold \bar{p}_ω indeed prescribes the optimal action for any possible belief \mathbf{p} . Hence, \bar{p}_ω characterizes the optimal policy. The student optimally exerts high effort H if her belief is weakly above the threshold, and low effort L otherwise. Note that \bar{p}_ω is increasing in p_α : increasing confidence can move the student's belief below the threshold for a given control belief p_ω , i.e., if she has higher confidence she is no longer willing to exert high effort in the current period. As in the two-period example, higher confidence gives an incentive to experiment with low effort in the hope of discovering that the state is α .

To prove optimality, we consider all possible one-step deviations. We show that either the deviation is not optimal or the deviation is coherent with the optimal policy, i.e., it induces the same immediate effort level. The proof has two parts:

Part 1 *In any optimal plan, after observing a success SH (SL), the student continues to exert high (low) effort forever.*

First, in both plan π_H and π_L , the student keeps the effort level unchanged after a success. To show that this must be the case for any optimal plan, we first consider a deviation to choosing H after observing SL . After SL the student knows with certainty that the state is α . In this case, there is no more learning and the student can save the effort cost by choosing L . Second, we consider a deviation to choosing L after observing SH . This is always dominated by either π_H or π_L and, thus, can never be optimal. Intuitively, if the student preferred to exert low effort after observing SH , she would have optimally chosen low effort already in the preceding period, precluding the observation of SH in the first place.¹⁵

¹⁵For an intuition, assume that the optimal plan prescribes high effort in period t and – independently of the outcome – low effort in period $t + 1$. That is, even after observing SH the student exerts low effort

Part 2 *When the optimal plan at belief \mathbf{p} begins with high (low) effort, the student prefers plan π_H to π_L (π_L to π_H).*

The plans π_H and π_L are the most responsive plans as they prescribe switching the effort level *immediately* after a failure. Therefore, they are decisive at the belief threshold \bar{p}_ω : π_H is the optimal plan just above this threshold and π_L just below. We show that any optimal plan prescribes the *same* effort level as the threshold \bar{p}_ω indicates. For beliefs further away from the threshold, other plans may be optimal – in particular, the student might want to try the same level of effort again despite a failure. However, the immediate effort level of these plans is coherent with that of the policy $m^*(\mathbf{p})$. To prove the second part, we need to address all plans in which – differently from π_H or π_L – the student does not switch the effort level immediately after a failure but only at a later period. We show that if a plan is optimal, it must induce the same immediate effort choice as indicated by $m^*(\mathbf{p})$. First, consider a plan that differs from π_L in that the student chooses L directly after FL . If such a plan is optimal, π_L is preferred to π_H after observing FL – i.e., at belief $\mathbf{p}(FL)$. We show that this implies that π_L is also preferred to π_H at the initial belief \mathbf{p} . Hence, we have $p_\omega < \bar{p}_\omega$, and this optimal plan is consistent with the optimal policy $m^*(\mathbf{p})$ which prescribes L . Second, consider a plan that differs from π_H in that the student chooses H directly after FH . If such a plan is optimal, π_H is preferred to π_L after observing FH – i.e., at belief $\mathbf{p}(FH)$. Again, we show that this implies that π_H is also preferred to π_L at the initial belief \mathbf{p} . Hence, we have $p_\omega \geq \bar{p}_\omega$, and this optimal plan is consistent with the optimal policy $m^*(\mathbf{p})$ which also prescribes H .

Since the analysis holds for any starting belief \mathbf{p} and the problem is stationary, the reasoning applies to any decision node and, therefore, covers all possible multi-step deviations. In conclusion, the threshold \bar{p}_ω gives a necessary and sufficient condition for the optimal action.

4.2 Absorbing beliefs

The optimal Markov policy results in a sequence of periods with high and low effort, which is not trivial to describe. We use two approaches to characterize the student’s behavior: In this subsection, we determine the sets of absorbing beliefs that the student will eventually

to have a chance to observe SL and discover that the state is α . Note that the likelihood of eventually observing SL in this plan is the same as in an alternative plan prescribing low effort in period t and high effort in $t + 1$, conditional on observing FL . However, if the student prefers low effort after SH , she strictly prefers the latter plan. This is because in the latter plan, the student has the possibility of saving the cost of high effort all together (if the state is in fact α) or delaying it (if the state is β). For a geometric intuition, note that the slope of updating after observing SH is steeper than that of the threshold \bar{p}_ω .

reach. In the next subsection, we compare how the student arrives at these beliefs using two indices for the dynamics of her behavior: perseverance and procrastination.

In each period, the student chooses between exerting high and low effort. In the short run, the student's behavior is solely determined by the belief threshold \bar{p}_ω . The student exerts high effort when her current belief \mathbf{p} is above the threshold and low effort when it is below. Importantly, after any failure, the student may change her effort level. In particular, after observing a failure with high effort, FH , her belief moves towards $p_\beta = 1$. If her posterior belief $\mathbf{p}(FH)$ crosses the threshold \bar{p}_ω , exerting low effort becomes optimal. For an illustration, see Figure 5.¹⁶ On the other hand, after observing a failure with low effort, FL , her confidence decreases and her belief moves towards the edge between $p_\omega = 1$ and $p_\beta = 1$. As a consequence, her posterior belief $\mathbf{p}(FL)$ moves towards the threshold \bar{p}_ω . If her posterior belief crosses this threshold, exerting high effort becomes optimal. Updating her beliefs in this fashion, after observing failures with low and high effort, the student may cross the threshold \bar{p}_ω repeatedly.




Figure 5: The thresholds \bar{p}_ω and \underline{p}_ω for $\delta = 0.7$, $\varepsilon = 0.1$, and $c = 0.4$.

The threshold \bar{p}_ω is also decisive for the student's behavior in the long run. For a sufficiently low control belief, the student's belief can never cross the threshold \bar{p}_ω again, even when the tally of failures with low effort, n_{FL} , goes to infinity. In this case, she continues to exert low effort forever. The next lemma characterizes the set of these beliefs.

Lemma 3 *If the student's belief satisfies $p_\omega < \underline{p}_\omega(p_\alpha)$ with*

$$\underline{p}_\omega(p_\alpha) := \frac{(1 - \delta)c}{(1 - \varepsilon)(1 - \delta\varepsilon - \delta c)} (1 - p_\alpha), \quad (2)$$

she exerts low effort forever.

¹⁶Geometrically, FH , decreases the beliefs in ω and α on a steeper slope than that of \bar{p}_ω .

We illustrate the threshold \underline{p}_ω in Figure 5. For beliefs below \underline{p}_ω , even observing failures with low effort repeatedly cannot increase the student's control belief sufficiently to cross the threshold \bar{p}_ω . The reason is that, with an increasing number of observations FL , her posterior belief moves on a straight line away from $p_\alpha = 1$ towards the set of absorbing beliefs $\overline{FL} := \{\mathbf{p} | p_\alpha = 0 \text{ and } p_\omega < \bar{p}_\omega\}$. Hence, once her belief drops below the threshold \underline{p}_ω , she exerts low effort forever.

In contrast to her behavior after a failure, the student never changes her effort after a success. In particular, after a success with low effort, SL , the student is fully informed that the state is α , and that low effort is optimal. Her belief remains in $\overline{SL} := \{\mathbf{p} | p_\alpha = 1\}$ and she continues to exert low effort forever. After a success with high effort, SH , the student knows that the state cannot be β , and her belief moves to $\mathbf{p}(SH) \in \overline{SH} := \{\mathbf{p} | p_\beta = 0 \text{ and } p_\omega \geq \bar{p}_\omega\}$ on the edge between $p_\omega = 1$ and $p_\alpha = 1$. She continues to exert high effort and, therefore, does not gain new information, so her belief remains fixed there.

To summarize the student's long term behavior, we find that asymptotically the student's belief remains fixed in one of three sets of absorbing beliefs.

Corollary 1 *The sets of absorbing beliefs are $\overline{SL} = \{\mathbf{p} | p_\alpha = 1\}$, $\overline{SH} = \{\mathbf{p} | p_\beta = 0 \text{ and } p_\omega \geq \bar{p}_\omega\}$, and $\overline{FL} = \{\mathbf{p} | p_\alpha = 0 \text{ and } p_\omega < \bar{p}_\omega\}$. With a belief in \overline{SL} and \overline{FL} the student exerts low effort forever, while she exerts high effort forever in \overline{SH} .*

In the set \overline{SH} , the student never learns whether the state is α or ω , while in the set \overline{FL} , she never learns whether it is β or ω . If the state is ω and effort could lead to success, in \overline{FL} the student exerts an inefficiently low level of effort forever.

4.3 Perseverance and procrastination in the face of failure

As we have seen in the previous subsection, the student never changes her effort level again after observing a success. In this subsection, we focus on the student's behavior if she only observes failures. We introduce perseverance as a long-term index for the student's willingness to work hard despite failures. Recall that the student may experiment with low and high effort for a while. However, there is a maximum number of failures with high effort she is willing to tolerate overall. We want to capture how persevering the student is in the long run, i.e., how often she exerts high effort overall without observing a success until she gives up forever.

Recall the dashed lines in Figure 5: Whenever the student observes one or multiple failures with high effort, FH , her belief moves on a straight line towards $p_\beta = 1$ and, thus, closer to the threshold \underline{p}_ω . Whenever the student observes one or multiple failures with low effort, FL , her belief moves on a straight line away from $p_\alpha = 1$. We know from Lemma 3 that the student never exerts high effort again once her belief falls below the threshold \underline{p}_ω .

The tally of failures with high effort n_{FH} until the belief crosses this threshold thus serves as a long-term index. Note that this index does not change in the tally of failures with low effort n_{FL} : If $n_{FH}, n_{FL} \in \mathbb{N}_0$ are such that $p_\omega(0, 0, n_{FH}, n_{FL}) < \underline{p}_\omega(p_\alpha(0, 0, n_{FH}, n_{FL}))$, then the inequality holds for all n_{FL} . In other words, once her belief crosses \underline{p}_ω , it will stay below this threshold, independently of the tally of failures with low effort. Moreover, since the student's posterior belief only depends on the tally of observed outcomes but not on their order, the tally of failures with low effort *before* her belief crosses \underline{p}_ω does not impact the index either.

Definition 1 (Perseverance) *For any tally of failures with low effort n_{FL} , the student's perseverance is given by*

$$\bar{n}_{FH}(\mathbf{p}) := \begin{cases} \min \left\{ n_{FH} \in \mathbb{N}_0 : p_\omega(0, 0, n_{FH}, n_{FL}) < \underline{p}_\omega(p_\alpha(0, 0, n_{FH}, n_{FL})) \right\}, & \text{if it exists} \\ \infty, & \text{if the set is empty,} \end{cases}$$

where $\mathbf{p}(0, 0, n_{FH}, n_{FL})$ is her belief after observing n_{FH} failures with high effort and n_{FL} failures with low effort.

We thus define the student's perseverance $\bar{n}_{FH}(\mathbf{p})$ as the overall tally of failures with high effort she is at most willing to observe, before her belief drops below \underline{p}_ω and she exerts low effort forever.¹⁷

For an illustration, see the left panel of Figure 6: A student with an initial belief below the threshold \underline{p}_ω (black area) exerts low effort forever without any additional observations of FH and thus $\bar{n}_{FH} = 0$. Instead, a student with a belief $p_\omega \geq \underline{p}_\omega$ (shaded area) is still willing to exert high effort and we have $\bar{n}_{FH} > 0$. Note that for beliefs that satisfy $p_\alpha + p_\omega = 1$, the student's beliefs are fixed in \overline{SH} and we have $\bar{n}_{FH} = \infty$, since such beliefs can not be changed by observing FH .

When the true state is ω , perseverance measures the probability that the student gives up exerting high effort due to a series of unlucky failures, given by $\varepsilon^{\bar{n}_{FH}}$, and never learns the true state.

In addition to the long-term index perseverance, we introduce procrastination as an index for the student's short-term behavior. It captures a temporary delay in high effort. We have seen that the student's control belief moves towards the threshold \bar{p}_ω when she observes a failure with low effort FL . Moreover, from Proposition 1 we know that the student exerts high effort when her belief rises above \bar{p}_ω . We therefore define the following.

¹⁷For a more explicit expression of perseverance in terms of the underlying parameters, see equation (A.6) in the appendix. In the derivation, the independence of perseverance on n_{FL} becomes obvious as the terms depending on it cancel out.




Figure 6: Perseverance (left) and procrastination (right) for $\delta = 0.8$, $c = 0.3$, and $\varepsilon = 0.4$.

Definition 2 (Procrastination) *The student's procrastination is given by*

$$\underline{n}_{FL}(\mathbf{p}) := \begin{cases} \min \{n_{FL} \in \mathbb{N}_0 : p_\omega(0, 0, 0, n_{FL}) \geq \bar{p}_\omega(p_\alpha(0, 0, 0, n_{FL}))\}, & \text{if it exists} \\ \infty, & \text{if the set is empty,} \end{cases}$$

where $\mathbf{p}(0, 0, 0, n_{FL})$ is her belief after observing n_{FL} failures with low effort.

In words, the student's procrastination is the tally of consecutive failures with low effort until she switches to high effort for the first time. Indeed, after observing a tally of $\underline{n}_{FL}(\mathbf{p})$ failures with low effort, a student with belief \mathbf{p} starts exerting high effort, as her belief is then above \bar{p}_ω .¹⁸

For an illustration, see the right panel of Figure 6: A student with belief $p_\omega \geq \bar{p}_\omega$ (black area) exerts high effort immediately, and we have $\underline{n}_{FL} = 0$. A student with belief $p_\omega < \bar{p}_\omega$ exerts low effort (shaded and white areas). If after a tally of $\underline{n}_{FL} \in \mathbb{N}$ failures her updated belief $p_\omega(0, 0, 0, \underline{n}_{FL})$ becomes sufficiently high, the student is willing to try exerting high effort (shaded area). However, a student with belief $p_\omega < \underline{p}_\omega$ (white area) is never willing to exert high effort for any tally of failures with low effort; thus we have $\underline{n}_{FL} = \infty$.

The higher the student's procrastination \underline{n}_{FL} , the more failures with low effort are required to motivate the student to exert high effort. Note that this is temporary, as she may again switch back and forth between low and high effort in the future. This is in contrast to perseverance, which measures the maximum number of periods the student will exert high effort overall before exerting low effort forever, unless she observes a success before.

The student's procrastination and perseverance depend on the student's beliefs as well as cost, patience, and noise. In the next subsection, we analyze how procrastination and perseverance change in these parameters.

¹⁸A more explicit expression of procrastination in terms of the underlying parameters is given in equation (A.7) in the appendix.

4.4 Welfare and policies

To analyze the optimal effort choices from the perspective of the student, we have assumed that she has an initial belief $\mathbf{p} \in (0, 1)^3$, where all states have strictly positive probabilities. The policy $m^*(\mathbf{p})$ derived in Proposition 1 achieves maximum welfare for any belief \mathbf{p} . Thus from the student's point of view her behavior is already optimal and there is no need for any intervention.

For this section, we take the perspective of a teacher who knows that the true state of the world is ω and that the student can only be successful if she exerts high effort. This is the perspective usually taken in the educational psychology literature, e.g., in the work of Dweck. From this perspective, the student suffers a loss of welfare due to her uncertainty about the state. Note that, while the student is uncertain about the state, her beliefs are not misspecified, as she attaches a positive probability to the true state. In principle, the student is able to learn the true state. However, since what she can learn depends on her effort, learning may be delayed or she may never learn the true state. The concepts of procrastination and perseverance allow to disentangle short- and long-run welfare effects of this uncertainty.

Procrastination \underline{n}_{FL} is a (negative) measure for short-run welfare. The longer it takes for the student to start exerting high effort, the longer she delays the possibility for success. Thus, policies that aim at *decreasing* the student's procrastination will improve her welfare. Perseverance \bar{n}_{FH} is a measure for long-run welfare. A student with higher perseverance has a lower probability of giving up high effort for good due to a series of unlucky failures. Intuitively, the longer the student tries to succeed with high effort, the more likely she eventually observes a success. Thus, policies that aim at *increasing* the student's perseverance will improve her welfare. Between the two measures, the long-term concept of perseverance clearly should be given greater weight. While procrastination only leads to a temporary delay, a lack of perseverance increases the chance that the student will never learn the true state. Once a student is trapped in the dark gray area below \underline{p}_ω (see Figure 5), she gives up forever and will never learn that she could have succeeded with high effort.

The next result shows how changes in the different beliefs and model parameters affect the student's perseverance and procrastination.

Proposition 2 *Perseverance \bar{n}_{FH} is decreasing in cost c and increasing in control p_ω , confidence p_α , and patience δ ; it is in- or decreasing in noise ε depending on the belief \mathbf{p} . Procrastination \underline{n}_{FL} is decreasing in control p_ω and increasing in confidence p_α , cost c and noise ε ; it is in- or decreasing in patience δ depending on the belief \mathbf{p} .*

Note that the comparative statics for the student's beliefs require a more precise statement of the usual *ceteris paribus* principle. Since the student's beliefs need to add up to

one, an increase in one belief needs to be compensated by a decrease in another belief. In particular, to derive the comparative statics for the control belief p_ω , we keep confidence p_α constant and set $p_\beta = 1 - p_\alpha - p_\omega$. This way, an increase in control translates to a decrease in the student’s belief p_β that she has insufficient ability to succeed. Similarly, to derive the comparative statics for confidence p_α , we keep control p_ω constant. Consequently, now an increase in confidence translates to a decrease in the student’s belief p_β that she has insufficient ability.¹⁹

Control A student with higher control belief thinks that it is more likely that high effort is necessary to succeed. Clearly, she is more likely to exert high effort in the short and in the long run. The student’s procrastination is decreasing and her perseverance is increasing in p_ω .

Directly increasing the control belief through feedback that emphasizes the importance of effort is the main focus of the interventions proposed in Dweck (2006). Similarly, Bandura (1997) and Duckworth et al. (2007) state that a high control belief is key to foster perseverance. Our analysis shows that this type of feedback has positive short- and long-term consequences, the student not only starts earlier to exert high effort but also stops later in the face of failures. Examples for feedback that stress effort would be “Good effort” or “Keep working hard”.

We see our paper as a complement to and extension of the approach usually taken in the educational psychology literature. Our model shares the feature that effort increases with the control belief. The value added lies in the description of the evolution of the control belief over time and the interaction with confidence. Moreover, we analyze the impact of confidence, cost, patience, and noise on perseverance and procrastination to learn how these affect the agent’s effort choices and learning over time.

Confidence Our model provides novel insights on the role of confidence. Mueller and Dweck (1998) find that praise for ability can decrease the students’ motivation. We find that the effect of the student’s ability belief on their effort level depends very much on the time frame considered; in the long term, a carefully designed message can, in fact, have a positive effect. For an illustration see Figure 7.

From a short-term perspective, the student’s procrastination \underline{n}_{FL} increases in confidence p_α . The student continues to exert low effort for longer before trying to succeed with high

¹⁹An alternative way to derive comparative statics would be to compensate an increase in one belief by a proportional decrease in both other beliefs. Note that this method would not change the results substantially, but it would complicate the exposition. Given our motivation, we find it more relevant to look at an increase in confidence for a given control belief and vice versa, at an increase in control for a given confidence level. This way, higher confidence results in a lower belief that ability is insufficient, and a higher control belief results in a lower belief that ability is so low that effort is useless.




Figure 7: Perseverance \bar{n}_{FH} (left) and procrastination \underline{n}_{FL} (right) as functions of confidence p_α , for $p_\omega = 0.2$, $\varepsilon = 0.4$, $\delta = 0.95$, and $c = 0.3$.

effort. This is coherent with the findings of Mueller and Dweck (1998) who consider a short time frame with a single instance of feedback. By contrast, from a long-term perspective, the student’s perseverance \bar{n}_{FH} increases in confidence p_α . Note that this effect is driven by the reduction in p_β , the student’s belief that she has insufficient ability. A student with high confidence is less easily discouraged as she tends to attribute failures to bad luck. This is coherent with the findings of Stinebrickner and Stinebrickner (2012). They elicit the ability beliefs of college students at the start of their course and find that, controlling for exam results, students with higher initial ability beliefs are less likely to drop out.

While, in the short run, high confidence makes it attractive to experiment with low effort for a chance to discover whether the state is α , this effect disappears in the long run. Perseverance measures the maximum attempts to succeed with high effort despite failure FH . Note that the observations of FH need not happen consecutively. Between periods with high effort the student may experiment with low effort arbitrarily often which has no impact on the measure of perseverance. Therefore, the motivation to experiment with low effort is not relevant for perseverance. In general, high confidence increases perseverance, as it takes more observations of failure with high effort FH to convince the student that the state is β . To summarize, while higher confidence may lead to less effort in the short run, it increases effort in the long run.

A key insight from our analysis is hence that interventions must be designed with care. Mueller and Dweck (1998) warn that praising ability can carry the message that ability – instead of effort – is decisive for success. This message would decrease the student’s control belief, which would both increase procrastination and reduce perseverance. Our model can give guidance on how to phrase messages about ability that are not detrimental to effort. Importantly, we find that a decrease in p_β is decisive for the effect of confidence on perseverance. Therefore a message that the student’s ability is not too low to achieve success (decreasing p_β) is most beneficial. An example for a helpful message would be: “You may have to work hard, but eventually you will make it.” Instead, an example for a potentially

harmful message would be: “You can do it if you are smart!”

Cost Unsurprisingly, an increase in the cost of effort c makes the student less willing to exert high effort, both in the short and in the long run. The student’s procrastination is increasing and her perseverance is decreasing in cost c . Hence, the agent suffers more from her uncertainty when costs are high, both in the long and in the short run.

This indicates that it is important to think about policies that decrease a student’s pecuniary and non-pecuniary costs of studying. Potential avenues could be to give individual support and to ensure a positive learning environment. This could be facilitated by increased investment in schools and teachers. Reducing the student’s costs will have a positive effect on perseverance and will decrease procrastination.

In our model, a high cost can equivalently be interpreted as a low payoff from success, i.e., low returns to education. There is evidence that students – especially students from disadvantaged backgrounds – underestimate returns to education. As a consequence, they underinvest in education. Better information about the returns to education or the exposure to relevant role models could encourage students to increase their investment in education; see, e.g., Jensen (2010), Wiswall and Zafar (2015), and Delavande, Bono and Holford (2020).

Patience In the short run, the impact of patience δ on the student’s procrastination n_{FL} depends on her confidence p_α . If her confidence is high, the student’s procrastination increases in her patience, if her confidence is low, it decreases. In other words, a confident student waits longer before trying to succeed with high effort, while a student with low confidence starts earlier. This corresponds to the *responsive* and the *despondent* plan of the student in the 2-period example. A more patient student has a higher gain from future payoffs and is, therefore, more willing to sacrifice current payoff in order to learn the true state. In the long term, higher patience increases perseverance, \bar{n}_{FH} increases in δ . Intuitively, future payoffs are relatively more valuable and the student is willing to tolerate a higher number of failures with high effort before giving up forever. Hence, in the long term, the agent suffers more from her uncertainty when she is less patient. Indeed, Golsteyn, Grönqvist and Lindahl (2014) find that a lower discount factor is associated with lower educational attainment.

According to Alan and Ertac (2018) it is possible to change a student’s patience: they provide evidence that patience can be enhanced through interventions teaching the imagination of future scenarios. While the short-run welfare implications are mixed, increasing patience will increase the student’s long-term welfare.

Noise From the perspective of the student, the connection between her effort and the outcome is noisy because of the possibility of bad luck. In the short run, procrastination

n_{FL} increases in the probability of bad luck ε . The effect is similar to the effect of a higher cost of effort. An increase in ε decreases the expected return to effort: effort is less likely to lead to success, even if the state is ω . This, clearly, decreases the incentives to exert high effort and the student's procrastination gets larger. The effect of an increase in noise on perseverance is ambivalent; \bar{n}_{FH} is non-monotonic in ε . Intuitively, on the one hand, the expected return to effort is reduced with more noise; on the other hand, a failure is less informative and, thus, less discouraging.

In a sports context, Foll, Rasche and Higgins (2006) find that, for a given control belief, subjects who attribute failure to unstable causes – such as luck – show greater perseverance than those who attribute failure to permanent causes. This indicates a long-term positive effect of some uncertainty. Indeed, with no uncertainty the student can rule out being a high type α after a failure. However, a large increase in noise substantially increases the likelihood of never exerting high effort despite failures, indicating that some reliability of the outcome is important to encourage students to try hard.

5 Conclusion

In this paper, we study how a student's control belief evolves over time. In particular, we are interested in the co-evolution of the control belief and confidence. Moreover, we aim at a better understanding of perseverance and procrastination in the face of failure. We develop a simple model of belief formation that addresses reasons for differing reactions to failure. In our model, a student who is uncertain about her own ability and the efficacy of effort tries to learn over time whether and how she can succeed. Depending on her control belief, cost of effort, patience, and confidence, she chooses whether to exert high or low effort in each period. We rationalize why some students react to failure by ceasing to exert effort, while others continue or increase their effort level. We characterize the optimal effort policy and show that it is a linear cutoff rule. Moreover, we show that the optimal effort choices cannot be described by a simple stopping rule: depending on her beliefs, the student may optimally repeatedly switch between high and low effort, despite the fact that the state does not change over time.

Our characterization result allows us to derive expressions for perseverance and procrastination in the face of failure. Under the assumption that effort is necessary, procrastination is detrimental for the student's welfare in the short run, while perseverance is beneficial in the long run. We show that a student's procrastination is decreasing in control but increasing in confidence, cost, and noise. Our model confirms that increasing the control belief is an effective method to promote perseverance in the face of failure. In addition, our analysis reveals that interventions that aim at increasing patience and decrease the cost of exerting

effort can also be used to foster perseverance. The effect of promoting confidence crucially depends on the time-frame considered: while higher confidence can decrease the effort level in the short run and lead to more procrastination, it increases perseverance in the long run.

Heckman and Rubinstein (2001) and Duckworth et al. (2007) explain success and failure by focusing on fixed personality traits (non-cognitive skills or grit). However, there has been criticism towards this approach: it can be tempting to blame a fault of character for failure while neglecting the inequalities in support and resources students face (Yeh; 2017). By endogenizing perseverance and procrastination we reconcile the two approaches. High effort can only prevail if the cost and the benefit are commensurate. If the reality is such that a student can not succeed by working hard, she reasonably stops trying eventually. Showing students goals they can aspire to and making sure that it is actually within their control to reach them are necessary preconditions for motivating them to work hard.

6 Acknowledgements

We thank two anonymous referees, Martin Cripps, Jeff Ely, Johannes Hörner, Nicolas Klein, Ilan Kremer, Christoph Kuzmics, Moritz Meyer-ter-Vehn, Claudio Mezzetti, Konrad Mierendorff, Sven Rady, Ernesto Rivera Mora, Larry Samuelson, seminar participants in Arizona, Bonn, Bristol, and Warwick, as well as the audiences at EEA (Geneva), GAMES (Maastricht), SCW (Lund), Vereinstagung (Augsburg), SFB-TR15 workshop (Bonn) for helpful comments.

This work was supported by the German Research Foundation (SFB-TR 15) and the Hausdorff Center of Mathematics.

A Appendix

Proof of Lemma 1.

In the second period, the student exerts high effort if and only if $p_\omega(h_1)(1 - \varepsilon) - c \geq 0$, for $h_1 \in \{SH, SL, FH, FL\}$. Using the posterior beliefs derived in Table 2, we get

$$\begin{aligned} p_\omega(SH)(1 - \varepsilon) - c \geq 0 &\Leftrightarrow p_\omega \geq p_\alpha \frac{c}{1 - \varepsilon - c} =: p_\omega^{SH}(p_\alpha); \\ p_\omega(FH)(1 - \varepsilon) - c \geq 0 &\Leftrightarrow p_\omega \geq \frac{c}{c + \varepsilon} \left(\frac{1}{1 - \varepsilon} - p_\alpha \right) =: p_\omega^{FH}(p_\alpha); \\ p_\omega(FL)(1 - \varepsilon) - c \geq 0 &\Leftrightarrow p_\omega \geq c \left(\frac{1}{1 - \varepsilon} - p_\alpha \right) =: p_\omega^{FL}(p_\alpha). \end{aligned}$$

The student never exerts effort after observing SL , since this implies $p_\alpha = 1$. Note that $p_\omega^{FH}(p_\alpha) \geq p_\omega^{FL}(p_\alpha)$. Moreover, $p_\omega^{FH}(p_\alpha) \geq p_\omega^{SH}(p_\alpha)$ inside the belief triangle. To see the latter, note that $p_\omega^{FH}(p_\alpha) \geq p_\omega^{SH}(p_\alpha)$ is equivalent to $p_\alpha \leq \frac{1 - c - \varepsilon}{1 - \varepsilon}$. Moreover, $p_\omega \geq p_\omega^{SH}(p_\alpha)$ is equivalent to $p_\alpha \leq p_\omega \frac{1 - c - \varepsilon}{c}$, which implies that $p_\alpha \leq (1 - p_\alpha) \frac{1 - c - \varepsilon}{c}$. Finally, $p_\alpha \leq (1 - p_\alpha) \frac{1 - c - \varepsilon}{c}$ is equivalent to $p_\alpha \leq \frac{1 - c - \varepsilon}{1 - \varepsilon}$. Thus $p_\omega \geq p_\omega^{SH}(p_\alpha)$ implies that $p_\omega^{FH}(p_\alpha) \geq p_\omega^{SH}(p_\alpha)$, or in words,

whenever the belief p_ω is above the threshold $p_\omega^{SH}(p_\alpha)$, it must also be above the threshold $p_\omega^{FH}(p_\alpha)$.

In the first period, the student anticipates her effort choice in the second period for each possible outcome of the first period. We already know that the student will always choose low effort after SL . Moreover, from the analysis above, we know that if $p_\omega > p_\omega^{FH}(p_\alpha)$ then $p_\omega > p_\omega^{FL}(p_\alpha)$ and $p_\omega > p_\omega^{SH}(p_\alpha)$. Therefore, there are five cases to consider, depending on the location of the initial belief p_ω relative to these thresholds.

1. For a prior \mathbf{p} such that $p_\omega \geq p_\omega^{FH}(p_\alpha)$, the student chooses high effort in period 2, after any outcome, except for SL . In period 1, the student exerts high effort if

$$\begin{aligned}
& EU((\pi_1(\emptyset) = H; \pi_2(SH) = H, \pi_2(FH) = H)) \\
& \geq EU((\pi_1(\emptyset) = L; \pi_2(SL) = L, \pi_2(FL) = H)) \\
\Leftrightarrow & \quad 2[(p_\alpha + p_\omega)(1 - \varepsilon) - c] \geq (2p_\alpha + p_\omega)(1 - \varepsilon) - c[1 - p_\alpha(1 - \varepsilon)] \\
\Leftrightarrow & \quad p_\omega \geq c \left(\frac{1}{1 - \varepsilon} + p_\alpha \right) =: p_\omega^1(p_\alpha).
\end{aligned}$$

Thus, plan $(\pi_1(\emptyset) = H; \pi_2(SH) = H, \pi_2(FH) = H)$ is optimal if $p_\omega \geq \max\{p_\omega^{FH}(p_\alpha), p_\omega^1(p_\alpha)\}$. Plan $(\pi_1(\emptyset) = L; \pi_2(SL) = L, \pi_2(FL) = H)$ is optimal if $p_\omega^1(p_\alpha) > p_\omega \geq p_\omega^{FH}(p_\alpha)$.

2. For a prior \mathbf{p} such that $p_\omega^{FH}(p_\alpha) > p_\omega \geq \max\{p_\omega^{SH}(p_\alpha), p_\omega^{FL}(p_\alpha)\}$, the student chooses high effort in period 2 after both FL and SH , but not after SL or FH . In period 1, the student exerts high effort if

$$\begin{aligned}
& EU((\pi_1(\emptyset) = H; \pi_2(SH) = H, \pi_2(FH) = L)) \\
& \geq EU((\pi_1(\emptyset) = L; \pi_2(SL) = L, \pi_2(FL) = H)) \\
\Leftrightarrow & \quad 2p_\alpha(1 - \varepsilon) + p_\omega(1 - \varepsilon)(2 - \varepsilon) - c[1 + (p_\alpha + p_\omega)(1 - \varepsilon)] \\
& \quad \geq (2p_\alpha + p_\omega)(1 - \varepsilon) - c[1 - p_\alpha(1 - \varepsilon)] \\
\Leftrightarrow & \quad p_\omega \geq \frac{2cp_\alpha}{1 - \varepsilon - c} =: p_\omega^2(p_\alpha).
\end{aligned}$$

Since $p_\omega^2(p_\alpha) > p_\omega^{SH}(p_\alpha)$, plan $(\pi_1(\emptyset) = H; \pi_2(SH) = H, \pi_2(FH) = L)$ is optimal if $p_\omega^{FH}(p_\alpha) > p_\omega \geq \max\{p_\omega^2(p_\alpha), p_\omega^{FL}(p_\alpha)\}$. Plan $(\pi_1(\emptyset) = L; \pi_2(SL) = L, \pi_2(FL) = H)$ is optimal if $\min\{p_\omega^{FH}(p_\alpha), p_\omega^2(p_\alpha)\} > p_\omega \geq \max\{p_\omega^{SH}(p_\alpha), p_\omega^{FL}(p_\alpha)\}$.

3. For a prior \mathbf{p} such that $p_\omega^{FL}(p_\alpha) > p_\omega \geq p_\omega^{SH}(p_\alpha)$, the student chooses high effort in period 2 after SH , but not after SL , FH , or FL . In period 1, the student exerts high

effort if

$$\begin{aligned}
& EU((\pi_1(\emptyset) = H; \pi_2(SH) = H, \pi_2(FH) = L)) \\
& \geq EU((\pi_1(\emptyset) = L; \pi_2(SL) = L, \pi_2(FL) = L)) \\
\Leftrightarrow & \quad 2p_\alpha(1 - \varepsilon) + p_\omega(1 - \varepsilon)(2 - \varepsilon) - c[1 + (p_\alpha + p_\omega)(1 - \varepsilon)] \geq 2p_\alpha(1 - \varepsilon) \\
\Leftrightarrow & \quad p_\omega \geq \frac{c}{2 - \varepsilon - c} \left(\frac{1}{1 - \varepsilon} + p_\alpha \right) =: p_\omega^3(p_\alpha).
\end{aligned}$$

Thus, plan $(\pi_1(\emptyset) = H; \pi_2(SH) = H, \pi_2(FH) = L)$ is optimal if $p_\omega^{FL}(p_\alpha) > p_\omega \geq p_\omega^3(p_\alpha)$. Plan $(\pi_1(\emptyset) = L; \pi_2(SL) = L, \pi_2(FL) = L)$ is optimal if $\min\{p_\omega^{FL}(p_\alpha), p_\omega^3(p_\alpha)\} > p_\omega \geq p_\omega^{SH}(p_\alpha)$

4. For a prior p such that $p_\omega^{SH}(p_\alpha) > p_\omega \geq p_\omega^{FL}(p_\alpha)$, the student chooses high effort in period 2 after FL , but not after SL , FH , or SH . In period 1, the student exerts high effort if

$$\begin{aligned}
& EU((\pi_1(\emptyset) = H; \pi_2(SH) = L, \pi_2(FH) = L)) \\
& \geq EU((\pi_1(\emptyset) = L; \pi_2(SL) = L, \pi_2(FL) = H)) \\
\Leftrightarrow & \quad (2p_\alpha + p_\omega)(1 - \varepsilon) - c \geq (2p_\alpha + p_\omega)(1 - \varepsilon) - c[1 - p_\alpha(1 - \varepsilon)] \\
\Leftrightarrow & \quad 0 \geq p_\alpha(1 - \varepsilon)c.
\end{aligned}$$

Since this can never hold in this interval, plan $(\pi_1(\emptyset) = L; \pi_2(SL) = L, \pi_2(FL) = H)$ is optimal if $p_\omega^{SH}(p_\alpha) > p_\omega \geq p_\omega^{FL}(p_\alpha)$.

5. For a prior p such that $\min\{p_\omega^{FL}(p_\alpha), p_\omega^{SH}(p_\alpha)\} > p_\omega$, the student always chooses low effort in period 2. In period 1, the student exerts high effort if

$$\begin{aligned}
& EU((\pi_1(\emptyset) = H; \pi_2(SH) = L, \pi_2(FH) = L)) \\
& \geq EU((\pi_1(\emptyset) = L; \pi_2(SL) = L, \pi_2(FL) = L)) \\
\Leftrightarrow & \quad (2p_\alpha + p_\omega)(1 - \varepsilon) - c \geq 2p_\alpha(1 - \varepsilon) \\
\Leftrightarrow & \quad p_\omega \geq \frac{c}{1 - \varepsilon}.
\end{aligned}$$

Since $\frac{c}{1 - \varepsilon} > p_\omega^{FL}(p_\alpha), p_\omega^{SH}(p_\alpha)$ plan $(\pi_1(\emptyset) = L; \pi_2(SL) = L, \pi_2(FL) = L)$ is optimal when $\min\{p_\omega^{FL}(p_\alpha), p_\omega^{SH}(p_\alpha)\} > p_\omega$.

□

Proof of Lemma 2.

To derive the threshold \bar{p}_ω , note first that the student's beliefs do not depend on the order of observed outcomes. Therefore, we can define $\mathbf{p}(FL, FH)$ as the belief resulting from \mathbf{p} after observing FH and FL in arbitrary order. Let $\bar{\pi}$ be some arbitrary fixed plan. We denote by $EU(\bar{\pi}, \mathbf{p}(FL, FH))$ the expected payoff that results from following an arbitrary plan $\bar{\pi}$ at belief $\mathbf{p}(FL, FH)$. Comparing the expected payoffs from plans π_H and π_L – defined in the main text – at belief \mathbf{p} yields

$$\begin{aligned}
& EU(\pi_H, \mathbf{p}) \geq EU(\pi_L, \mathbf{p}) \\
\Leftrightarrow & p_\omega(1 - \varepsilon) \left[1 + \frac{\delta(1 - \varepsilon)}{1 - \delta} \right] + p_\alpha \frac{1 - \varepsilon}{1 - \delta} \\
& - c \left[1 + (p_\omega + p_\alpha) \frac{\delta(1 - \varepsilon)}{1 - \delta} \right] + \delta^2 (p_\omega \varepsilon + p_\alpha \varepsilon^2 + p_\beta) EU[\bar{\pi}, \mathbf{p}(FL, FH)] \\
& \geq \delta p_\omega(1 - \varepsilon) \left[1 + \frac{\delta(1 - \varepsilon)}{1 - \delta} \right] + p_\alpha \frac{1 - \varepsilon}{1 - \delta} \\
& - \delta c \left[(p_\omega + p_\alpha \varepsilon + p_\beta) + (p_\omega + p_\alpha \varepsilon) \frac{\delta(1 - \varepsilon)}{1 - \delta} \right] + \delta^2 (p_\omega \varepsilon + p_\alpha \varepsilon^2 + p_\beta) EU[\bar{\pi}, \mathbf{p}(FL, FH)] \\
\Leftrightarrow & p_\omega \geq \frac{c}{(1 - \delta\varepsilon - \delta c)} \left[\frac{1 - \delta}{1 - \varepsilon} + p_\alpha \frac{\delta(2 - \delta - \delta\varepsilon)}{1 - \delta} \right]. \tag{A.1}
\end{aligned}$$

□

Proof of Proposition 1. For this proof, we apply the principle of optimality and show that there are no profitable one-shot deviations. Therefore, deviations after different histories can be considered separately. We proceed by a pairwise comparison of the pivotal plans to selected plans that differ only along one history. The proof consists of two parts.




Figure A.1: Plans π_H^{SL} (left) and π_L^{SH} (right).

Part 1: In any optimal plan, after a success, the student continues to exert the same effort level forever.

1.i) After a success with low effort (SL), it is always optimal to exert low effort forever. Consider a plan π_H^{SL} that prescribes H at any point after SL . See Figure A.1 (left), for an illustration. Note that after observing SL once, the student knows with certainty that the state is α and thus effort is futile. At any subsequent history, the immediate payoff from L is $1 - \varepsilon$ while the payoff from H is $1 - \varepsilon - c$. Moreover, no more information can be gained for the future. Therefore, an optimal plan can never prescribe high effort after observing SL .

1.ii) After a success with high effort (SH), it is always optimal to exert high effort forever. Assume for contradiction that a plan π_L^{SH} which prescribes L after SH is optimal. For an illustration, see Figure A.1 (right). If plan π_L^{SH} is optimal, it must give a higher payoff than π_H :

$$\begin{aligned}
& EU(\pi_L^{SH}, \mathbf{p}) > EU(\pi_H, \mathbf{p}) \\
\Leftrightarrow & p_\omega(1 - \varepsilon) \left[1 + \frac{\delta^2(1 - \varepsilon)}{1 - \delta} \right] + p_\alpha \frac{1 - \varepsilon}{1 - \delta} \\
& - c \left[1 + (p_\omega + p_\alpha \varepsilon) \frac{\delta^2(1 - \varepsilon)}{1 - \delta} \right] + \delta^2 (p_\omega \varepsilon + p_\alpha \varepsilon^2 + p_\beta) EU[\bar{\pi}, \mathbf{p}(FL, FH)] \\
& > p_\omega(1 - \varepsilon) \left[1 + \frac{\delta(1 - \varepsilon)}{1 - \delta} \right] + p_\alpha \frac{1 - \varepsilon}{1 - \delta} \\
& - c \left[1 + (p_\omega + p_\alpha) \frac{\delta(1 - \varepsilon)}{1 - \delta} \right] + \delta^2 (p_\omega \varepsilon + p_\alpha \varepsilon^2 + p_\beta) EU[\bar{\pi}, \mathbf{p}(FL, FH)] \\
\Leftrightarrow & \frac{p_\omega}{p_\omega + p_\alpha} < c \left[\frac{1}{1 - \varepsilon} + \frac{\delta}{1 - \delta} \frac{p_\alpha}{p_\omega + p_\alpha} \right] \\
& \Leftrightarrow p_\omega(SH) < c \left[\frac{1}{1 - \varepsilon} + \frac{\delta}{1 - \delta} p_\alpha(SH) \right] \\
& \Leftrightarrow p_\omega < c \left[\frac{p_\omega + p_\alpha}{1 - \varepsilon} + \frac{\delta}{1 - \delta} p_\alpha \right]. \tag{A.2}
\end{aligned}$$

If plan π_L^{SH} is optimal, it must also give a higher payoff than π_L . Assume instead that

π_L is preferred:

$$\begin{aligned}
& EU(\pi_L, \mathbf{p}) > EU(\pi_L^{SH}, \mathbf{p}) \\
\Leftrightarrow & \delta p_\omega(1-\varepsilon) \left[1 + \frac{\delta(1-\varepsilon)}{1-\delta} \right] + p_\alpha \frac{1-\varepsilon}{1-\delta} - \delta c \left[(p_\omega + p_\alpha \varepsilon + p_\beta) + (p_\omega + p_\alpha \varepsilon) \frac{\delta(1-\varepsilon)}{1-\delta} \right] \\
& + \delta^2 (p_\omega \varepsilon + p_\alpha \varepsilon^2 + p_\beta) EU[\bar{\pi}, \mathbf{p}(FL, FH)] \\
& > p_\omega(1-\varepsilon) \left[1 + \frac{\delta^2(1-\varepsilon)}{1-\delta} \right] + p_\alpha \frac{1-\varepsilon}{1-\delta} - c \left[1 + (p_\omega + p_\alpha \varepsilon) \frac{\delta^2(1-\varepsilon)}{1-\delta} \right] \\
& + \delta^2 (p_\omega \varepsilon + p_\alpha \varepsilon^2 + p_\beta) EU[\bar{\pi}, \mathbf{p}(FL, FH)] \\
\Leftrightarrow & p_\omega < c \left[\frac{1}{1-\varepsilon} + \frac{\delta}{1-\delta} p_\alpha \right]. \tag{A.3}
\end{aligned}$$

We see that it is equivalent to say that the student prefers plan π_L^{SH} to π_H at belief \mathbf{p} and that the student prefers plan π_L to π_L^{SH} at belief $\mathbf{p}(SH)$. However, inequality (A.3) is easier to fulfill than inequality (A.2). Therefore, whenever π_L^{SH} is preferred to π_H , it is inferior to π_L and whenever π_L^{SH} is preferred to π_L , it is inferior to π_H . Thus π_L^{SH} is either dominated by π_L , by π_H , or by both.

Next, consider a plan π_L^{SHi} which prescribes starting with H and playing L repeatedly after observing SH . Specifically, it prescribes playing L after observing SH and continuing with L after observing FL i times, with $i \geq 1$. If there is no success, it prescribes playing H forever, after observing FL for the $(i+1)$ -st time. If plan π_L^{SHi} is optimal, then it must yield a higher payoff than any other plan. We can construct an alternative plan π_H^{SHi} , which is identical to π_L^{SHi} except that it prescribes L after observing FL $(i-1)$ times and H forever after observing FL for the i -th time. We then compare these plans:

$$\begin{aligned}
& EU(\pi_L^{SHi}, \mathbf{p}) > EU(\pi_H^{SHi}, \mathbf{p}) \\
\Leftrightarrow & p_\omega(1-\varepsilon) \left[1 + \frac{\delta^{i+2}(1-\varepsilon)}{1-\delta} \right] + p_\alpha \frac{1-\varepsilon}{1-\delta} - c \left[1 + (p_\omega + p_\alpha \varepsilon^{i+1}) \frac{\delta^{i+2}(1-\varepsilon)}{1-\delta} \right] \\
& > p_\omega(1-\varepsilon) \left[1 + \frac{\delta^{i+1}(1-\varepsilon)}{1-\delta} \right] + p_\alpha \frac{1-\varepsilon}{1-\delta} - c \left[1 + (p_\omega + p_\alpha \varepsilon^i) \frac{\delta^{i+1}(1-\varepsilon)}{1-\delta} \right] \\
\Leftrightarrow & \frac{p_\omega}{p_\omega + p_\alpha \varepsilon^i} < c \left[\frac{1}{1-\varepsilon} + \frac{\delta}{1-\delta} \frac{p_\alpha \varepsilon^i}{p_\omega + p_\alpha \varepsilon^i} \right] \\
\Leftrightarrow & p_\omega(SH, FL^i) < c \left[\frac{1}{1-\varepsilon} + \frac{\delta}{1-\delta} p_\alpha(SH, FL^i) \right] \\
\Leftrightarrow & p_\omega < c \left[\frac{p_\omega + p_\alpha \varepsilon^i}{1-\varepsilon} + \frac{\delta}{1-\delta} p_\alpha \varepsilon^i \right]. \tag{A.4}
\end{aligned}$$

We see that inequality (A.4) is even harder to fulfill than inequality (A.2). Observing an additional FL makes low effort less attractive.

Next, consider a plan $\hat{\pi}_L^{SHi}$ which is identical to π_L^{SHi} , except that the first two periods are switched: $\hat{\pi}_L^{SHi}$ starts with L and, after observing FL , H is played in period 2. Then, it prescribes playing L after observing SH and continuing with L after observing FL for $(i - 1)$ times. If there is no success, it prescribes playing H forever, after observing FL for the $(i + 1)$ -st time overall. If plan π_L^{SHi} is optimal, it must also give a higher payoff than $\hat{\pi}_L^{SHi}$. Assume instead that $\hat{\pi}_L^{SHi}$ is preferred:

$$\begin{aligned}
& EU(\hat{\pi}_L^{SHi}, \mathbf{p}) > EU(\pi_L^{SHi}, \mathbf{p}) \\
\Leftrightarrow & \delta p_\omega(1 - \varepsilon) \left[1 + \frac{\delta^{i+1}(1 - \varepsilon)}{1 - \delta} \right] + p_\alpha \frac{1 - \varepsilon}{1 - \delta} - \delta c \left[(p_\omega + p_\beta + p_\alpha \varepsilon) + (p_\omega + p_\alpha \varepsilon^{i+1}) \frac{\delta^{i+1}(1 - \varepsilon)}{1 - \delta} \right] \\
& > p_\omega(1 - \varepsilon) \left[1 + \frac{\delta^{i+2}(1 - \varepsilon)}{1 - \delta} \right] + p_\alpha \frac{1 - \varepsilon}{1 - \delta} - c \left[1 + (p_\omega + p_\alpha \varepsilon^{i+1}) \frac{\delta^{i+2}(1 - \varepsilon)}{1 - \delta} \right] \\
\Leftrightarrow & p_\omega < c \left[\frac{1}{1 - \varepsilon} + \frac{\delta}{1 - \delta} p_\alpha \right]. \tag{A.5}
\end{aligned}$$

Comparing inequalities (A.4) and (A.5), we see that inserting the posterior belief $\mathbf{p}(SH, FL^i)$ into (A.5) results in an expression that is equivalent to (A.4). It is thus equivalent to say that the student prefers plan π_L^{SHi} to π_H^{SHi} at belief \mathbf{p} and that the student prefers plan $\hat{\pi}_L^{SHi}$ to π_L^{SHi} at belief $\mathbf{p}(SH, FL^i)$. Inequality (A.5) is easier to fulfill than inequality (A.4). Therefore, whenever π_L^{SHi} is preferred to π_H^{SHi} , it is inferior to $\hat{\pi}_L^{SHi}$ and whenever π_L^{SHi} is preferred to $\hat{\pi}_L^{SHi}$, it is inferior to π_H^{SHi} . Thus π_L^{SHi} is either dominated by $\hat{\pi}_L^{SHi}$, by π_H^{SHi} , or by both.

The argument extends to plans where L is played – one or several times – not directly after the first observation SH , but j periods later. The reason is that the student’s belief after observing SH adjusts to $p_\beta = 0$ and does not change as long as the student chooses H .

Part 2: When the optimal plan at belief \mathbf{p} begins with high (low) effort, the agent prefers plan π_H to π_L (π_L to π_H).

While any optimal plan must prescribe continuing with the same effort choice after a success, after a failure different effort choices may be optimal for different beliefs. Thus, π_H or π_L will not always be optimal. We now show that if a plan different from π_H or π_L is optimal at a given belief, then it induces exactly the same “next” effort choice as the preferred plan of π_H and π_L .

For this, it is useful to consider the effect of observing FL and FH on inequality (A.1) which signifies that plan π_H is weakly preferred to π_L . After observing FL , resulting in

posterior belief $\mathbf{p}(FL)$, $\boldsymbol{\pi}_H$ is weakly preferred to $\boldsymbol{\pi}_L$ if and only if

$$\begin{aligned} & EU[\boldsymbol{\pi}_H, \mathbf{p}(FL)] \geq EU[\boldsymbol{\pi}_L, \mathbf{p}(FL)] \\ \Leftrightarrow & \frac{p_\omega}{p_\omega + p_\beta + p_\alpha \varepsilon} \geq \frac{c}{(1 - \delta \varepsilon - \delta c)} \left[\frac{1 - \delta}{1 - \varepsilon} + \frac{p_\alpha \varepsilon}{p_\omega + p_\beta + p_\alpha \varepsilon} \frac{\delta(2 - \delta - \delta \varepsilon)}{1 - \delta} \right] \\ \Leftrightarrow & p_\omega \geq \frac{c}{(1 - \delta \varepsilon - \delta c)} \left[\frac{1 - \delta}{1 - \varepsilon} [1 - p_\alpha(1 - \varepsilon)] + p_\alpha \varepsilon \frac{\delta(2 - \delta - \delta \varepsilon)}{1 - \delta} \right], \end{aligned}$$

where

$$\frac{c}{(1 - \delta \varepsilon - \delta c)} \left[\frac{1 - \delta}{1 - \varepsilon} [1 - p_\alpha(1 - \varepsilon)] + p_\alpha \varepsilon \frac{\delta(2 - \delta - \delta \varepsilon)}{1 - \delta} \right] < \bar{p}_\omega(p_\alpha).$$

The minimum value of belief p_ω such that the student prefers plan $\boldsymbol{\pi}_H$ to $\boldsymbol{\pi}_L$ after observing FL is *below* the threshold \bar{p}_ω – for which the student prefers plan $\boldsymbol{\pi}_H$ to $\boldsymbol{\pi}_L$ with the prior belief \mathbf{p} . This means that observing FL makes it easier to fulfill inequality (A.1). Therefore, whenever a student prefers plan $\boldsymbol{\pi}_L$ to $\boldsymbol{\pi}_H$ at the posterior belief $\mathbf{p}(FL)$, she must also prefer it at the prior belief \mathbf{p} .

After observing FH , resulting in posterior belief $\mathbf{p}(FH)$, $\boldsymbol{\pi}_H$ is weakly preferred to $\boldsymbol{\pi}_L$ if and only if

$$\begin{aligned} & EU[\boldsymbol{\pi}_H, \mathbf{p}(FH)] \geq EU[\boldsymbol{\pi}_L, \mathbf{p}(FH)] \\ \Leftrightarrow & \frac{p_\omega \varepsilon}{p_\omega \varepsilon + p_\alpha \varepsilon + p_\beta} \geq \frac{c}{(1 - \delta \varepsilon - \delta c)} \left[\frac{1 - \delta}{1 - \varepsilon} + \frac{p_\alpha \varepsilon}{p_\omega \varepsilon + p_\alpha \varepsilon + p_\beta} \frac{\delta(2 - \delta - \delta \varepsilon)}{1 - \delta} \right] \\ \Leftrightarrow & p_\omega \geq \frac{c}{(1 - \delta \varepsilon - \delta c)} \left[\frac{(1 - \delta)}{1 - \varepsilon} \left[1 + \frac{1 - \varepsilon}{\varepsilon} (1 - p_\omega - p_\alpha) \right] + p_\alpha \frac{\delta(2 - \delta - \delta \varepsilon)}{1 - \delta} \right] \\ & > \frac{c}{(1 - \delta \varepsilon - \delta c)} \left[\frac{1 - \delta}{1 - \varepsilon} + p_\alpha \frac{\delta(2 - \delta - \delta \varepsilon)}{1 - \delta} \right] = \bar{p}_\omega(p_\alpha). \end{aligned}$$

The minimum value of belief p_ω such that the student prefers plan $\boldsymbol{\pi}_H$ to $\boldsymbol{\pi}_L$ after observing FH is *above* the threshold \bar{p}_ω – for which the student prefers plan $\boldsymbol{\pi}_H$ to $\boldsymbol{\pi}_L$ with the prior belief \mathbf{p} . This means that observing FH makes it harder to fulfill inequality (A.1). Therefore, whenever a student prefers plan $\boldsymbol{\pi}_H$ to $\boldsymbol{\pi}_L$ at the posterior belief $\mathbf{p}(FH)$, she must also prefer it at the prior belief \mathbf{p} .

2.i) If the optimal plan starts with L and prescribes to continue with L after failure (FL), then the policy $m^*(\mathbf{p})$ also prescribes L . We show that for such a plan to be optimal, we must have that the student prefers plan $\boldsymbol{\pi}_L$ to $\boldsymbol{\pi}_H$ at the posterior belief $\mathbf{p}(FL)$. As we have seen, she must then also prefer $\boldsymbol{\pi}_L$ at the prior belief \mathbf{p} .

Consider a plan $\boldsymbol{\pi}_L^{FL}$ which starts with L . After SL it prescribes L forever. After FL , it prescribes following plan $\boldsymbol{\pi}_L$ – i.e., play L once more and continue with L in case of success,




Figure A.2: Plans π_L^{FL} (left) and π_H^{FL} (right).

but switch to H in case of failure. If the plan π_L^{FL} is optimal, then it must yield a higher expected payoff than any other plan. We can construct an alternative plan π_H^{FL} , which is identical to π_L^{FL} , except that it prescribes following plan π_H after FL . For an illustration, see Figure A.2.

We show that it is equivalent to say that the student prefers plan π_L^{FL} to π_H^{FL} at belief \mathbf{p} and that the student prefers plan π_L to π_H at belief $\mathbf{p}(FL)$:

$$\begin{aligned}
& EU(\pi_L^{FL}, \mathbf{p}) > EU(\pi_H^{FL}, \mathbf{p}) \\
\Leftrightarrow & \delta^2 p_\omega (1 - \varepsilon) \left[1 + \frac{\delta(1 - \varepsilon)}{1 - \delta} \right] + p_\alpha \frac{1 - \varepsilon}{1 - \delta} \\
& - \delta^2 c \left[(p_\omega + p_\beta + p_\alpha \varepsilon^2) + (p_\omega + p_\alpha \varepsilon^2) \frac{\delta(1 - \varepsilon)}{1 - \delta} \right] \\
& + (p_\omega \varepsilon + p_\beta + p_\alpha \varepsilon^2) EU(\bar{\pi}, \mathbf{p}(FL, FL, FH)) \\
& > \delta p_\omega (1 - \varepsilon) \left[1 + \frac{\delta(1 - \varepsilon)}{1 - \delta} \right] + p_\alpha \frac{1 - \varepsilon}{1 - \delta} \\
& - \delta c \left[(p_\omega + p_\beta + p_\alpha \varepsilon) + (p_\omega + p_\alpha \varepsilon) \frac{\delta(1 - \varepsilon)}{1 - \delta} \right] \\
& + (p_\omega \varepsilon + p_\beta + p_\alpha \varepsilon^2) EU(\bar{\pi}, \mathbf{p}(FL, FL, FH)) \\
\Leftrightarrow & \frac{p_\omega}{p_\omega + p_\beta + p_\alpha \varepsilon} < \frac{c}{(1 - \delta \varepsilon - \delta c)} \left[\frac{1 - \delta}{1 - \varepsilon} + \frac{p_\alpha \varepsilon}{p_\omega + p_\beta + p_\alpha \varepsilon} \frac{\delta(2 - \delta - \delta \varepsilon)}{1 - \delta} \right] \\
\Leftrightarrow & EU[\pi_H, \mathbf{p}(FL)] < EU[\pi_L, \mathbf{p}(FL)].
\end{aligned}$$

The comparison of π_L^{FL} and π_H^{FL} at belief \mathbf{p} reduces to the comparison of π_L and π_H at the posterior belief $\mathbf{p}(FL)$. As the plans π_L^{FL} and π_H^{FL} are identical after observing SL , the respective payoffs cancel out. We have already shown that, whenever a student prefers plan π_L to π_H at the posterior belief $\mathbf{p}(FL)$, she must also prefer it at the prior belief \mathbf{p} . Therefore, if plan π_L^{FL} is optimal then it prescribes the same immediate action as the policy $m^*(\mathbf{p})$: L .




Figure A.3: Plans π_H^{FH} (left) and π_L^{FH} (right).

Since the analysis holds for any possible starting belief \mathbf{p} and the problem is stationary, the reasoning applies to any decision node. Thus, we can extend this analysis to plans prescribing playing L several times after observing FL : Consider a plan $\pi_L^{FL^i}$ which prescribes playing L i times after observing FL , with $i \geq 1$. Specifically, it prescribes to play L after observing FL ($i - 1$) times and, after observing FL for the i -th time, to follow plan π_L . We compare this plan to an alternative plan $\pi_H^{FL^i}$ which is identical except that it prescribes plan π_H after observing FL for the i -th time. The comparison of the expected payoffs is identical to the one above, but with new starting belief $\mathbf{p}(FL^i)$, which results from observing FL i times. Recall that, whenever a student prefers plan π_H to π_L at the posterior belief $\mathbf{p}(FL)$, she also prefers π_H at the prior belief \mathbf{p} . Since the problem is stationary, this implies that whenever a student prefers plan π_H to π_L at the posterior belief $\mathbf{p}(FL^i)$, she must also prefer π_H at the prior belief $\mathbf{p}(FL^{i-1})$, and thus, also at the prior belief \mathbf{p} . Thus, if a plan $\pi_L^{FL^i}$, which prescribes L i times after observing FL , is optimal, then the first action is coherent with $m^*(\mathbf{p}) = L$.

2.ii) If the optimal plan starts with high effort and prescribes to continue high effort even after failure (FH), then the policy $m^*(\mathbf{p})$ also prescribes H . We show that for such a plan to be optimal, we must have that the student prefers plan π_H to π_L at the posterior belief $\mathbf{p}(FH)$. As we have seen, she then also prefers π_H at the prior belief \mathbf{p} .

Consider a plan π_H^{FH} which starts with H . After SH it prescribes H forever. After FH it prescribes following plan π_H – i.e., play H once more and continue with H in case of success, but switch to L in case of failure. If the plan π_H^{FH} is optimal, then it must yield a higher expected payoff than any other plan. We can construct an alternative plan π_L^{FH} , which is identical to π_H^{FH} , except that it prescribes following plan π_L after FH . For an illustration, see Figure A.3.

We now show that it is equivalent to say that the student prefers plan π_H^{FH} to π_L^{FH} at

belief \mathbf{p} and that the student prefers plan π_H to π_L at belief $\mathbf{p}(FH)$:

$$\begin{aligned}
& EU(\pi_H^{FH}, \mathbf{p}) \geq EU(\pi_L^{FH}, \mathbf{p}) \\
\Leftrightarrow & p_\omega(1-\varepsilon) \left[1 + \delta + \frac{\delta^2(1-\varepsilon^2)}{1-\delta} \right] + p_\alpha \frac{1-\varepsilon}{1-\delta} \\
& - c \left[1 + \delta + (p_\omega + p_\alpha) \frac{\delta^2(1-\varepsilon^2)}{1-\delta} \right] + (p_\omega\varepsilon^2 + p_\alpha\varepsilon^2 + p_\beta) EU(\bar{\pi}, \mathbf{p}(FL, FH, FH)) \\
& \geq p_\omega(1-\varepsilon) \left[1 + \frac{\delta(1-\varepsilon)}{1-\delta} \right] + p_\alpha \frac{1-\varepsilon}{1-\delta} \\
& - c \left[1 + (p_\omega + p_\alpha) \frac{\delta(1-\varepsilon)}{1-\delta} \right] + (p_\omega\varepsilon^2 + p_\alpha\varepsilon^2 + p_\beta) EU(\bar{\pi}, \mathbf{p}(FL, FH, FH)) \\
\Leftrightarrow & \frac{p_\omega\varepsilon}{p_\omega\varepsilon + p_\alpha\varepsilon + p_\beta} \geq \frac{c}{(1-\delta\varepsilon - \delta c)} \left[\frac{1-\delta}{1-\varepsilon} + \frac{p_\alpha\varepsilon}{p_\omega\varepsilon + p_\alpha\varepsilon + p_\beta} \frac{\delta(2-\delta-\delta\varepsilon)}{1-\delta} \right] \\
\Leftrightarrow & EU[\pi_H, \mathbf{p}(FH)] \geq EU[\pi_L, \mathbf{p}(FH)].
\end{aligned}$$

The comparison of π_H^{FH} and π_L^{FH} at belief \mathbf{p} reduces to comparing the payoff from π_H and π_L at the posterior belief $\mathbf{p}(FH)$. As the plans π_H^{FH} and π_L^{FH} are identical for all other histories, the respective payoffs cancel out. We have already shown that, whenever a student prefers plan π_H to π_L at the posterior belief $\mathbf{p}(FH)$, she must also prefer it at the prior belief \mathbf{p} . Therefore, if plan π_H^{FH} is optimal then it prescribes the same immediate action as the policy $m^*(\mathbf{p})$: H . Analogously to the case above, this argument directly extends to plans which prescribe playing H several times after FH .

Since the plans after observing both FL and FH are unspecified, any possible action following these outcomes is covered in the analysis. In conclusion, the threshold \bar{p}_ω gives a necessary and sufficient condition for the optimal action. \square

Proof of Lemma 3. Consider a student with prior belief $\mathbf{p} = (p_\omega, p_\alpha, p_\beta)$. The posterior beliefs, after observing history $(0, 0, 0, n_{FL})$ are given by:

$$\begin{aligned}
p_\omega((0, 0, 0, n_{FL})) &= \frac{p_\omega}{p_\omega + p_\alpha\varepsilon^{n_{FL}} + 1 - p_\omega - p_\alpha}, \\
p_\alpha((0, 0, 0, n_{FL})) &= \frac{p_\alpha\varepsilon^{n_{FL}}}{p_\omega + p_\alpha\varepsilon^{n_{FL}} + 1 - p_\omega - p_\alpha}.
\end{aligned}$$

Taking the limits, we obtain

$$\begin{aligned}
\lim_{n_{FL} \rightarrow \infty} p_\omega((0, 0, 0, n_{FL})) &= \frac{p_\omega}{1 - p_\alpha}, \\
\lim_{n_{FL} \rightarrow \infty} p_\alpha((0, 0, 0, n_{FL})) &= 0.
\end{aligned}$$

From Proposition 1, we know that the student exerts high effort if and only if $p_\omega \geq \bar{p}_\omega$. This inequality becomes easier to fulfil as n_{FL} increases. In order to find beliefs for which the threshold \bar{p}_ω cannot be reached, no matter how often the student observes the outcome FL , we consider the posterior beliefs as $n_{FL} \rightarrow \infty$. Using the posterior beliefs derived above, we obtain the following condition for low effort forever:

$$\begin{aligned} \lim_{n_{FL} \rightarrow \infty} p_\omega((0, 0, 0, n_{FL})) &< \lim_{n_{FL} \rightarrow \infty} \bar{p}_\omega(p_\alpha((0, 0, 0, n_{FL}))) \\ \Leftrightarrow \quad \frac{p_\omega}{1 - p_\alpha} &< \frac{c}{(1 - \delta\varepsilon - \delta c)} \left[\frac{1 - \delta}{1 - \varepsilon} + 0 \cdot \frac{\delta(2 - \delta - \delta\varepsilon)}{1 - \delta} \right] \\ \Leftrightarrow \quad p_\omega &< \frac{c(1 - \delta)}{(1 - \varepsilon)(1 - \delta\varepsilon - \delta c)} (1 - p_\alpha) =: \underline{p}_\omega(p_\alpha). \end{aligned}$$

The threshold \underline{p}_ω defines all beliefs such that inequality (A.1) is violated in the limit for $n_{FL} \rightarrow \infty$. \square

Proof of Proposition 2. Recall, that we do comparative statics by setting $p_\beta = 1 - p_\alpha - p_\omega$, such that changes with respect to p_ω are captured by p_β while p_α remains constant and similarly, changes with respect to p_α are captured by p_β while p_ω remains constant.

First, we rewrite the definition of perseverance and calculate the comparative statics. Perseverance is derived from the condition in Definition 1 that the belief p_ω falls below the threshold $\underline{p}_\omega(p_\alpha)$. For $p_\omega = 0$, the student exerts low effort forever, and $\bar{n}_{FH} = 0$. Thus, we assume $p_\omega > 0$ for this part of the proof. We rewrite the condition in Definition 1 as follows:

$$\begin{aligned} p_\omega(0, 0, n_{FH}, n_{FL}) &< \underline{p}_\omega(p_\alpha(0, 0, n_{FH}, n_{FL})) \\ \Leftrightarrow \quad \frac{p_\omega \varepsilon^{n_{FH}}}{p_\omega \varepsilon^{n_{FH}} + p_\alpha \varepsilon^{(n_{FH} + n_{FL})} + 1 - p_\omega - p_\alpha} &< \frac{c(1 - \delta)}{(1 - \varepsilon)(1 - \delta\varepsilon - \delta c)} \left(1 - \frac{p_\alpha \varepsilon^{(n_{FH} + n_{FL})}}{p_\omega \varepsilon^{n_{FH}} + p_\alpha \varepsilon^{(n_{FH} + n_{FL})} + 1 - p_\omega - p_\alpha} \right) \\ \Leftrightarrow \quad \varepsilon^{n_{FH}} &< \frac{1 - p_\alpha - p_\omega}{p_\omega} \frac{c(1 - \delta)}{(1 - \varepsilon - c)(1 - \delta\varepsilon)} \\ \Leftrightarrow \quad n_{FH} &> \frac{1}{\ln(\varepsilon)} \ln \left(\frac{1 - p_\alpha - p_\omega}{p_\omega} \frac{c(1 - \delta)}{(1 - \varepsilon - c)(1 - \delta\varepsilon)} \right) =: \tilde{n}_{FH}. \end{aligned} \tag{A.6}$$

Note that \bar{n}_{FH} is the minimal step function satisfying condition (A.6). The relevant range is $\mathbf{p} \geq \underline{p}_\omega$, since outside of this range \bar{n}_{FH} is constant. We derive monotonicity properties of \bar{n}_{FH} by taking the derivatives of \tilde{n}_{FH} , the right side of inequality (A.6). Consequently, if the derivative of \tilde{n}_{FH} with respect to a variable is positive (negative) in the relevant range,

\bar{n}_{FH} weakly increases (decreases) in this variable. We obtain

$$\begin{aligned}
\frac{\partial}{\partial p_\omega} \tilde{n}_{FH} &= -\frac{1-p_\alpha}{p_\omega(1-p_\alpha-p_\omega)\ln(\varepsilon)} > 0, \\
\frac{\partial}{\partial p_\alpha} \tilde{n}_{FH} &= -\frac{1}{(1-p_\alpha-p_\omega)\ln(\varepsilon)} > 0, \\
\frac{\partial}{\partial c} \tilde{n}_{FH} &= \frac{1-\varepsilon}{c(1-c-\varepsilon)\ln(\varepsilon)} < 0, \\
\frac{\partial}{\partial \delta} \tilde{n}_{FH} &= -\frac{1-\varepsilon}{(1-\delta)(1-\delta\varepsilon)\ln(\varepsilon)} > 0, \\
\frac{\partial}{\partial \varepsilon} \tilde{n}_{FH} &= \frac{1+\delta(1-c-2\varepsilon)}{(1-c-\varepsilon)(1-\delta\varepsilon)\ln(\varepsilon)} - \frac{1}{\varepsilon\ln(\varepsilon)^2} \ln\left(\frac{c(1-p_\alpha-p_\omega)(1-\delta)}{p_\omega(1-c-\varepsilon)(1-\delta\varepsilon)}\right).
\end{aligned}$$

Second, we rewrite the definition of procrastination and calculate the comparative statics. procrastination is derived from the condition in Definition 2 that the belief p_ω rises above the threshold $\bar{p}_\omega(p_\alpha)$. For $p_\alpha = 0$, \underline{n}_{FL} is constant: For $p_\alpha = 0$ and $p_\omega < \bar{p}_\omega$, we have $\underline{n}_{FL} = \infty$. For $p_\alpha = 0$ and $p_\omega \geq \bar{p}_\omega$, we have $\underline{n}_{FL} = 0$. Thus, we assume $p_\alpha > 0$ for this part of the proof. We rewrite the condition in Definition 2 as follows:

$$\begin{aligned}
p_\omega((0, 0, 0, n_{FL})) &\geq \bar{p}_\omega(p_\alpha(0, 0, 0, n_{FL})) \\
\Leftrightarrow \frac{p_\omega}{p_\omega + p_\alpha \varepsilon^{n_{FL}} + 1 - p_\omega - p_\alpha} & \\
&\geq \frac{c}{(1-\delta\varepsilon-\delta c)} \left[\frac{1-\delta}{1-\varepsilon} + \frac{p_\alpha \varepsilon^{n_{FL}}}{p_\omega + p_\alpha \varepsilon^{n_{FL}} + 1 - p_\omega - p_\alpha} \frac{\delta(2-\delta-\delta\varepsilon)}{1-\delta} \right] \\
\Leftrightarrow \varepsilon^{n_{FL}} &\leq \frac{1-\delta}{p_\alpha(1-\delta\varepsilon)} \left[p_\omega(1-\varepsilon) \left(\frac{1}{c} - \frac{\delta}{1-\delta\varepsilon} \right) - (1-p_\alpha) \frac{1-\delta}{1-\delta\varepsilon} \right] \\
\Leftrightarrow n_{FL} &\geq \frac{1}{\ln(\varepsilon)} \ln\left(\frac{(1-\delta)[p_\omega(1-\varepsilon)(1-\delta\varepsilon-\delta c) - c(1-\delta)(1-p_\alpha)]}{cp_\alpha(1-\delta\varepsilon)^2} \right) := \tilde{n}_{FL}. \quad (\text{A.7})
\end{aligned}$$

Note that \underline{n}_{FL} is the minimal step function satisfying condition (A.7). The relevant range is $\underline{p}_\omega \leq \mathbf{p} < \bar{p}_\omega$, since outside of this range the function is constant. We derive monotonicity properties of \underline{n}_{FL} , by taking the derivatives of \tilde{n}_{FL} , the right side of inequality (A.7). Consequently, if the derivative of \tilde{n}_{FL} with respect to a variable is positive (negative) in the relevant range, \bar{n}_{FL} weakly increases (decreases) in this variable. The derivatives are

given by

$$\begin{aligned}
\frac{\partial}{\partial p_\omega} \tilde{n}_{FL} &= \frac{(1-\varepsilon)(1-\delta c - \delta\varepsilon)}{\ln(\varepsilon) [p_\omega(1-\varepsilon)(1-\delta\varepsilon - \delta c) - c(1-\delta)(1-p_\alpha)]} \leq 0, \\
\frac{\partial}{\partial p_\alpha} \tilde{n}_{FL} &= -\frac{p_\omega(1-\varepsilon)(1-\delta\varepsilon - \delta c) - c(1-\delta)}{p_\alpha \ln(\varepsilon) [p_\omega(1-\varepsilon)(1-\delta\varepsilon - \delta c) - c(1-p_\alpha)(1-\delta)]} \geq 0, \\
\frac{\partial}{\partial c} \tilde{n}_{FL} &= -\frac{p_\omega(1-\varepsilon)(1-\delta\varepsilon)}{c \ln(\varepsilon) [p_\omega(1-\varepsilon)(1-\delta\varepsilon - \delta c) - c(1-\delta)(1-p_\alpha)]} \geq 0, \\
\frac{\partial}{\partial \delta} \tilde{n}_{FL} &= -\frac{(1-\varepsilon) [p_\omega(1-\varepsilon)(1-\delta\varepsilon - \delta c) + p_\omega c(1-\delta) - 2c(1-\delta)(1-p_\alpha)]}{(1-\delta)(1-\delta\varepsilon) \ln(\varepsilon) [p_\omega(1-\varepsilon)(1-\delta\varepsilon - \delta c) - c(1-p_\alpha)(1-\delta)]}, \\
\frac{\partial}{\partial \varepsilon} \tilde{n}_{FL} &= \frac{p_\omega \ln(\varepsilon) [1 + \delta(1-c-2\varepsilon)]}{p_\omega(1-\varepsilon)(1-\delta\varepsilon - \delta c) - c(1-p_\alpha)(1-\delta)} - \frac{2\delta \ln(\varepsilon)}{1-\delta\varepsilon} \\
&\quad - \frac{1}{\varepsilon} \ln \left[\frac{(1-\delta) [p_\omega(1-\varepsilon)(1-\delta\varepsilon - \delta c) - c(1-p_\alpha)(1-\delta)]}{c p_\alpha (1-\delta\varepsilon)^2} \right] \geq 0.
\end{aligned}$$

For the derivative with respect to δ , note that the sign is characterized by

$$p_\omega \geq \frac{2c(1-\delta)}{(1-\varepsilon)(1-\delta\varepsilon) + c(1-2\delta + \delta\varepsilon)}(1-p_\alpha),$$

which is a line through $\{p_\alpha = 1\}$ in the belief triangle, with a steeper negative slope than \underline{p}_ω has. The derivative is negative between \underline{p}_ω and this line, and positive above this line. Thus, procrastination \underline{n}_{FL} decreases in patience δ for low values of confidence p_α but increases for high values. The cross partial with respect to δ and p_α is positive:

$$\frac{\partial}{\partial \delta \partial p_\alpha} \tilde{n}_{FL} = -\frac{c p_\omega (1-\varepsilon)(1-c-\varepsilon)}{[c(1-\delta)(1-p_\alpha) - p_\omega(1-\varepsilon)(1-\delta\varepsilon - \delta c)]^2 \ln(\varepsilon)} \geq 0.$$

□

References

- Alan, S. and Ertac, S. (2018). Fostering patience in the classroom: Results from randomized educational intervention, *Journal of Political Economy* **126**(5): 1865–1911.
- Bandura, A. (1997). *Self-efficacy: The exercise of control*, Macmillan.
- Bolton, P. and Harris, C. (1999). Strategic experimentation, *Econometrica* **67**(2): 349–374.

- Caliendo, M., Cobb-Clark, D. A. and Uhlendorff, A. (2015). Locus of control and job search strategies, *Review of Economics and Statistics* **97**(1): 88–103.
- Camargo, B. (2007). Good news and bad news in two-armed bandits, *Journal of Economic Theory* **135**(1): 558–566.
- Che, Y.-K. and Mierendorff, K. (2019). Optimal dynamic allocation of attention, *American Economic Review* **109**(8): 2993–3029.
- Coleman, M. and DeLeire, T. (2003). An economic model of locus of control and the human capital investment decision, *Journal of Human Resources* **38**(3): 701–721.
- Delavande, A., Bono, E. D. and Holford, A. (2020). Academic and non-academic investments at university: The role of expectations, preferences and constraints, *Journal of Econometrics* .
- Duckworth, A. L., Peterson, C., Matthews, M. D. and Kelly, D. R. (2007). Grit: perseverance and passion for long-term goals, *Journal of personality and social psychology* **92**(6): 1087–1101.
- Dweck, C. (2000). *Self-theories: Their role in motivation, personality, and development*, Psychology Press.
- Dweck, C. (2006). *Mindset: The new psychology of success*, Random House.
- Foll, D. L., Rasclé, O. and Higgins, N. (2006). Persistence in a putting task during perceived failure: Influence of state-attributions and attributional style, *Applied Psychology* **55**(4): 586–605.
- Gittins, J. C. (1979). Bandit processes and dynamic allocation indices, *Journal of the Royal Statistical Society: Series B (Methodological)* **41**(2): 148–164.
- Golsteyn, B. H., Grönqvist, H. and Lindahl, L. (2014). Adolescent time preferences predict lifetime outcomes, *The Economic Journal* **124**(580): 739–761.
- Heckman, J. J. and Rubinstein, Y. (2001). The importance of noncognitive skills: Lessons from the GED testing program, *American Economic Review* **91**(2): 145–149.
- Heckman, J. J., Stixrud, J. and Urzua, S. (2006). The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior, *Journal of Labor economics* **24**(3): 411–482.
- Heidhues, P., Kőszegi, B. and Strack, P. (2018). Unrealistic expectations and misguided learning, *Econometrica* **86**(4): 1159–1214.

- Hestermann, N. and Le Yaouanq, Y. (2021). Experimentation with self-serving attribution biases, *American Economic Journal: Microeconomics* **13**(3): 198–237.
- Jensen, R. (2010). The (perceived) returns to education and the demand for schooling, *The Quarterly Journal of Economics* **125**(2): 515–548.
- Keller, G., Rady, S. and Cripps, M. (2005). Strategic experimentation with exponential bandits, *Econometrica* **73**(1): 39–68.
- Kim, K. R. and Seo, E. H. (2015). The relationship between procrastination and academic performance: A meta-analysis, *Personality and Individual Differences* **82**: 26–33.
- Klein, N. and Rady, S. (2011). Negatively correlated bandits, *The Review of Economic Studies* **78**(2): 693–732.
- Lemoine, D. (2020). Rationally misplaced confidence. mimeo.
- Mayskaya, T. (2019). Dynamic choice of information sources. mimeo.
- Mueller, C. M. and Dweck, C. S. (1998). Praise for intelligence can undermine children’s motivation and performance, *Journal of personality and social psychology* **75**(1): 33.
- Nikandrova, A. and Panco, R. (2018). Dynamic project selection, *Theoretical Economics* **13**(1): 115–143.
- Pastorino, E. (2019). Careers in firms: the role of learning and human capital. mimeo.
- Piketty, T. (1995). Social mobility and redistributive politics, *Quarterly Journal of Economics* **110**(3): 551–584.
- Rothschild, M. (1974). A two-armed bandit theory of market pricing, *Journal of Economic Theory* **9**(2): 185–202.
- Rotter, J. B. (1966). Generalized expectancies for internal versus external control of reinforcement, *Psychological monographs: General and applied* **80**(1): 1–28.
- Stinebrickner, T. and Stinebrickner, R. (2012). Learning about academic ability and the college dropout decision, *Journal of Labor Economics* **30**(4): 707–748.
- Wiswall, M. and Zafar, B. (2015). How do college students respond to public information about earnings?, *Journal of Human Capital* **9**(2): 117–169.
- Yeh, C. (2017). Forget grit: Focus on inequality, *Education Week* **36**(28).