

# Pessimism Traps and Algorithmic Interventions

Avrim Blum  


Toyota Technological Institute at Chicago, IL, USA

Emily Diana 

Carnegie Mellon University, Tepper School of Business, Pittsburgh, PA, USA

Kavya Ravichandran  

Toyota Technological Institute at Chicago, IL, USA

Alexander Tolbert 

Emory University, Department of Quantitative Theory and Methods, Atlanta, GA, USA

---

## Abstract

In this paper, we relate the philosophical literature on pessimism traps to information cascades, a formal model derived from the economics and mathematics literature. A pessimism trap is a social pattern in which individuals in a community, in situations of uncertainty, copy the sub-optimal actions of others, despite their individual beliefs. This maps nicely onto the concept of an information cascade, which involves a sequence of agents making a decision between two alternatives, with a private signal of the superior alternative and a public history of others' actions. Key results from the economics literature show that information cascades occur with probability one in many contexts, and depending on the strength of the signal, populations can fall into the incorrect cascade very easily and quickly. Once formed, in the absence of external perturbation, a cascade cannot be broken – therefore, we derive an intervention that can be used to nudge a population from an incorrect to a correct cascade and, importantly, maintain the cascade once the subsidy is discontinued. We extend this to the case of multiple communities, each of which might have a different optimal action, and a government providing subsidies that cannot discriminate between communities and does not know which action is optimal for each. We study this both theoretically and empirically.

**2012 ACM Subject Classification** Applied computing → Economics

**Keywords and phrases** Pessimism trap, opinion dynamics, algorithmic interventions, subsidy, decision-making

**Digital Object Identifier** 10.4230/LIPIcs.FORC.2025.5

**Related Version** *Full Version:* <https://arxiv.org/abs/2406.04462>

**Funding** This work was supported in part by the National Science Foundation under grants CCF-2212968 and ECCS-2216899, by the Simons Foundation under the Simons Collaboration on the Theory of Algorithmic Fairness, and by the Office of Naval Research MURI Grant N000142412742.

## 1 Introduction

Studying decision-making under uncertainty has interested scholars in several disciplines, such as psychology, economics, and philosophy. In economics, specifically in statistical discrimination, scholars such as [7] have studied how individuals make decisions, particularly in marginalized communities where societal structures and perceived barriers can impact their choices. This often leads to self-censorship or a withdrawal from available opportunities due to a pessimistic outlook. The concept of pessimism has also been explored in philosophy, specifically in recent work by [11]. [11]'s research examines the dilemmas faced by individuals belonging to marginalized groups when presented with evidence that diminishes their chances of success. As a result, they may rationally choose to redirect their efforts elsewhere to pursue other less valuable alternatives. However, this decision only reinforces the pessimistic



© Avrim Blum, Emily Diana, Kavya Ravichandran, and Alexander Tolbert;  
licensed under Creative Commons License CC-BY 4.0

6th Symposium on Foundations of Responsible Computing (FORC 2025).

Editor: Mark Bun; Article No. 5; pp. 5:1–5:19



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

mindset that influenced it, perpetuating a cycle of limited aspirations and self-doubt. [11] criticizes the simplistic solution of turning towards optimism and proposes a more critical and situational application of optimistic beliefs. [11] characterizes *pessimism traps* as follows:

A **pessimism trap** . . . is meant to capture how negative beliefs about one’s likelihood of success can play a role in agents pursuing less risky, modest ends instead of ambitious ones, thereby further entrenching the negative evidence that the agent herself and other agents in a similar position face. The central idea, however, has broader applicability. . . . for this paper, I will focus on those cases in which poverty, prejudice, and discrimination play a role in providing agents with the sort of evidence that would *prima facie* make it **rational** for them to arrive at the pessimistic beliefs that play a role in thwarting their ambition.

Our paper extends Morton’s qualitative description of pessimism traps (summarized in Appendix A) with mathematical formalism and studies algorithmic decision-making within this formalism. We explore the theoretical model of *information cascades* (also referred to as “herding” in the opinion dynamics literature [12, 13]) to shed light on decision-making under uncertainty and pessimism in such contexts. Using this model, we propose interventions to break pessimism traps and redirect the population into favorable states.

## 1.1 Our Contributions

In this paper, we study pessimism traps as conceptualized by Morton and offer interventions to shift communities out of these traps. Our first major contribution is to link the concept of a pessimism trap to the opinion dynamics literature. We do so by interpreting a single-dimensional information cascade model as a decision sequence in a community. We extend this in Section 4 to consider  $k$  parallel information cascades representing the decision processes of  $k$  independent communities.

Our next major contribution is an intervention for the single community setting that sustainably shifts communities out of pessimism. We consider a subsidy enacted by the government to incentivize certain behavior in an agent. Two natural ideas fail: first, suppose we subsidize the ambitious action by a large fixed amount. While this intervention incentivizes agents to act ambitiously in the short term, once the subsidy program ends, subsequent agents will ascribe prior agents’ ambition entirely to the subsidy they received and fall back into pessimism traps. Secondly, subsidizing by too small an amount fails to incentivize the ambitious action over the moderate one at all. Thus, we derive the precise size of the subsidy that will incentivize those who were already leaning toward ambition while not moving those who were anyway not considering the ambitious choice. Herein lies our first surprising insight – the non-monotonicity associated with the effect of the size of subsidy.

Next, we study  $k$  communities, each behaving as above. The government still intervenes but must act impartially with respect to community membership. Each of these groups may have a different optimal choice among the two options, and the government (a) does not know which choice is optimal for each group and (b) must be blind to community membership in providing subsidies. We show that in this case, we can construct a distribution  $\mathcal{D}$  such that if the government draws the subsidy randomly from that distribution, then eventually each community will shift toward what is the optimal end for them. This intervention relies crucially on the previous insight: if the provided subsidy is in the “just right” range discussed above for an agent, the agent’s action will shift their community toward optimal action. However, if it is too large or too small (outside of that range), the community will remain in the same state. Thus, provided the distribution assigns at least a minimum probability

to the “just right” range for a fixed community, that group will eventually settle into their optimal action. Therefore, our second main insight is that since we can guarantee (1) the existence of a range of subsidy values that shift the community toward the correct action for that community and (2) a single community faces no *negative* effect when provided a subsidy outside of that range, the government can use randomization to shift *all* groups into optimal behaviors, even without knowing what is optimal for any given community.

Finally, we verify the effectiveness of our proposed intervention in the one-group case via experiments. Our experiments show that the intervention we develop is indeed successful in shifting sequential decisions into optimism. Further, we confirm that the required budget for the intervention is tractable. Our experiments provide additional confirmation of the benefits of our intervention in the simple model that we study.

We view our work as a step forward in providing a theoretical model in which to study pessimism traps and, more broadly, in developing empowering interventions for marginalized communities. While our model is necessarily stylized, our major insights could help guide interventions in more complex settings.

In summary, our contributions are:

1. Linking the opinion dynamics and pessimism traps literatures by studying pessimism trap formation in the information cascade model.
2. Identifying an intervention that sustainably shifts a community of agents away from pessimism traps.
3. Extending this result into a setting where there are  $k$  different groups and where the government knows neither which the correct action is for a group, nor how close they are to escaping the trap and showing the power of randomization in this setting.
4. Corroborating the success of our theoretical interventions experimentally.

The remainder of our paper is organized as follows. First, we briefly survey work from a diverse array of fields related to our work. Next, we present the formal mathematical model in which we study pessimism traps along with some preliminary facts about them, and we reflect on the strengths and weaknesses of this modeling approach. In Section 3, we introduce and analyze our intervention to shift a single group toward optimism. We then extend this in Section 4 to the setting with  $k$  groups each of whom has an unknown optimal action. Finally, we provide experiments that show the success of our proposed interventions, substantiating our theoretical results.

## 1.2 Related Work

Here, we provide a brief survey of related works that contextualize our approach to studying pessimism traps and interventions aimed at promoting optimism.

As quoted above, [11] formalizes the notion of pessimism traps, building upon several empirical characterizations from scholars studying the phenomenon in the field of education [8]. Morton highlights the importance of *belief*, which is why we focus our modeling and interventions on shifting beliefs in a sustainable manner.

Much work in the literature studies the connection between individual behavior and community-wide effects. This includes the literature on opinion dynamics and social learning, which are surveyed in [14] and [6]. They discuss both simple and more complex models for how people’s beliefs vary in relation to the beliefs of those with whom they interact. [1] studies both Bayesian and non-Bayesian models for how agents update their beliefs, investigating consensus and asymptotic learning of state (i.e., what are the true beliefs in the world). On the behavioral economics side, [10] study the psychological reasons behind why individuals might act in a way that disregard their private information.

Since individual behavior can impact community-wide outcomes, as studied in the works above, it is reasonable to consider “nudging” or intervening at the individual level. Works that have studied this include [15] in behavioral economics, showing how minor policy adjustments can realign individual decision-making with optimal outcomes. Similarly, work in theoretical computer science and game theory, such as [3], study “nudging” specifically in the context of equilibria, where the goal is to redirect a population from a less desirable to a more desirable equilibrium.

Additionally, empirical studies have supported the application of nudges in various domains. For example, [9] extends the discussion on the efficacy of nudges in real-world settings. Integrating nudging into public policy, particularly in health and environmental strategies, as discussed by [4], demonstrates how these concepts have evolved beyond theoretical discussions to practical implementations.

Therefore, we utilize information cascade models to develop “nudging”-style interventions aimed at combating the pessimism traps that frequently arise in marginalized communities. [5] initially framed the classic information cascade model, demonstrating how individuals, despite possessing private information, often conform to the erroneous actions of predecessors due to the strong influence of prior actions. This model has served as a baseline for exploring various dimensions of information processing within groups. Building on this foundational work, [2] conducted laboratory experiments to observe cascade behavior in controlled settings, adding empirical evidence to the theoretical predictions. We mainly focus on the former.

## 2 Preliminaries and Basic Model

We study two settings in this paper: first, we look at a single sequence of  $T$  agents deciding whether to take action  $A$  or  $B$ ; second, we will expand this to  $k$  parallel sequences of agents, where the best action (between  $A$  and  $B$ ) differs group-to-group. Here, we define the signal and sequence dynamics for a single sequence; in Section 4, we flesh out the setting with  $k$  parallel groups. We take  $T$  to be finite, but our analysis extends to infinite  $T$ .

### 2.1 Single Sequence Model

First, for the case of a single community (sequence) we assume that  $A$  is uniformly the better option (that is, it is best for all agents), though this fact is unknown to the agents. When we consider  $k$  parallel sequences, then each sequence will have its own better option. We let  $\mathbf{E}_A$  represent the event that action  $A$  is correct and  $\mathbf{E}_B$  represent the event that action  $B$  is correct. When referring to the set of these possible events, we define  $\mathcal{E} \triangleq \{\mathbf{E}_A, \mathbf{E}_B\}$  to indicate the set of possibilities. In particular, we assume that *a priori*, agents have no bias toward either action, which we model by saying that they have a common, uniform prior over  $\mathcal{E}$ . Formally,  $\mathbb{P}[\mathbf{E}_A] = \frac{1}{2}$  and  $\mathbb{P}[\mathbf{E}_B] = \frac{1}{2}$ . Further, one action has associated with it reward 1 and the other reward 0. In addition to the public prior, each agent  $t$  receives a private signal  $s_t$  indicating that one of  $A$  or  $B$  is the correct action. This signal is correct with probability  $p > \frac{1}{2}$  and incorrect with probability  $1 - p$ . Therefore, we have  $\mathbb{P}[s_t | \mathbf{E}_A] = p^{\delta_{s_t, A}} (1 - p)^{\delta_{s_t, B}}$  and  $\mathbb{P}[s_t | \mathbf{E}_B] = p^{\delta_{s_t, B}} (1 - p)^{\delta_{s_t, A}}$ , where  $\delta$  is the Kronecker delta function.

Given this signal and the actions of preceding agents, the agent decides whether to take action  $A$  or  $B$ . Let  $h_t$  indicate the action of agent  $t$ , where an agent is identified by the time they act and  $H_{t-1}$  be the history of actions for the first  $t - 1$  agents. Agent  $t$  will take an action if its expected reward exceeds the expected reward of the other action. For concreteness, we adopt the tie-breaking convention that if an agent is indifferent between the two options based on the calculated posteriors, they follow their private signal. Now, we define an information cascade.

► **Definition 1** (Information Cascade). *An information cascade occurs when an agent's action does not depend on their private signal. This means that one of the following occurs:*

1.  $\mathbb{P}[\mathbf{E}_A | s_t = 1, H_{t-1}] > \frac{1}{2}$  and  $\mathbb{P}[\mathbf{E}_A | s_t = 0, H_{t-1}] > \frac{1}{2}$
2.  $\mathbb{P}[\mathbf{E}_A | s_t = 1, H_{t-1}] < \frac{1}{2}$  and  $\mathbb{P}[\mathbf{E}_A | s_t = 0, H_{t-1}] < \frac{1}{2}$

► **Remark 1.** *A cascade begins when the observed history  $H_{t-1}$  becomes sufficiently skewed towards either adoption or rejection. Therefore, enough previous agents have taken the same action to make it rational for subsequent agents to follow the trend, based on a simple Bayesian updating procedure, even if the history contradicts their private signal.*

To formally capture the concept of a pessimism trap for a sequence setting, we define the conditions under which agents, influenced by the actions of their predecessors, consistently choose the inferior option  $B$ . This phenomenon occurs when the aggregated evidence from prior decisions leads to a bias that overrides the agents' private signals. We represent this situation mathematically as follows:

► **Definition 2** (Pessimism Trap). *A pessimism trap occurs when an information cascade leads agents to consistently choose the inferior action  $B$  due to an overwhelming influence of prior agents' incorrect actions. Formally, a pessimism trap is defined by the following condition:*

$$\mathbb{P}[\mathbf{E}_B | s_t = 1, H_{t-1}] > \frac{1}{2} \text{ and } \mathbb{P}[\mathbf{E}_B | s_t = 0, H_{t-1}] > \frac{1}{2}$$

*This indicates that despite the private signal  $s_t$  suggesting action  $A$  is better, history  $H_{t-1}$  has led agent  $t$  to believe that  $B$  is better. Subsequently, agents' decisions are based purely on observed history rather than private signals.*

► **Remark 2.** *When computing posteriors, we assume that agents only consider actions of those up until a cascade began, since if agents are rational, they realize that no further inferential information can be obtained from individuals who do not use their signals.*

As a result, we will often be interested in  $\bar{H}_{t-1}$ , the history of actions taken by agents who chose an action if and only if it was their signal. We express mathematically the Bayesian updating procedure through which agents come to their posterior belief about the best action based on the history of actions before a cascade begins and then formalize the kinds of governmental / central interventions we study:

$$\mathbb{P}[\mathbf{E}_A | s_t, H_{t-1}] = \mathbb{P}[\mathbf{E}_A | s_t, \bar{H}_{t-1}] = \frac{\mathbb{P}[s_t | \mathbf{E}_A] \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_A]}{\sum_{w \in \{A, B\}} \mathbb{P}[s_t | \mathbf{E}_w] \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_w]} \quad (1)$$

► **Definition 3.** *A subsidy of size  $r$  toward action  $a$  is a benefit provided to an agent taking said action independent of the true world. Thus, in the world where  $a$  is the correct action, instead of receiving just  $R$  reward for taking action  $a$ , the agent receives  $R + r$  reward, and in a world in which it is the incorrect action, the agent still receives  $r$  reward.*

**Formulation as Random Walk.** The main technical tools used for analysis of the cascade and subsidy come from framing the process as a random walk. We isolate the probability with which action  $A$  is taken and the probability with which action  $B$  is taken, and we analyze the information cascade as taking steps on a random walk. This walk in general terminates when either an up or a down cascade has been reached. Once we introduce a subsidy, the stopping point of interest will be the up cascade. In Appendix C, we show how we can use the random walk formalism to derive the length of time for which the subsidy must be in place and the required budget.

## 2.2 On Our Modeling Choices

A significant contribution of our work is that we propose using the information cascade model to study pessimism traps and associated interventions. In this section, we justify our use of this model and discuss the trade-offs made.

### Justification

A natural question is: why choose this particular level of abstraction, i.e., why model the pessimism trap as hinging in this way on decisions made by previous agents? Several works have empirically characterized this phenomenon in the context of education, building on which Morton gives an epistemic characterization. At its core, as conceptualized by Morton, a pessimism trap occurs as a result of people’s beliefs (about the world, about the rationality of others, etc), realized as herding behavior. Thus, in order to model this faithfully, we require a model in which we may quantitatively update the *beliefs*. A Bayesian formulation is natural for this. Further, though sequentialness is not inherent to Morton’s characterization of pessimism traps, there is a natural sense in which agents consider the context of those who went *before* them when making decisions. Accordingly, a sequential model is a good choice for organizing the manner in which agents are influenced.

### Strengths

In some sense, the Bayesian posterior update represents the “optimal” rational decision. By modeling pessimism traps this way, we are showing that even with perfect rationality, a trap can form, corroborating Morton’s point that pessimism traps do not occur due to lack of rationality. Further, as summarized in the appendix, two key features Morton identifies regarding pessimism traps are: (1) there is evidence of similar people not succeeding at the ambitious end and (2) not pursuing the ambitious end will not change the agent’s view of its value. Both of these are well-represented in the information cascade model. Particularly once we extend to multiple groups, agents are looking to the history of actions taken by people like them (i.e., in the same group), and witnessing several people taking a certain action would indicate to them that taking the opposite action does not tend to be beneficial for people like them. Likewise, in the information cascade model, there is no feedback for whether a different action would have been correct, and therefore there is no reason for an agent to change their pessimistic view about it.

### What this model misses

On the other hand, since the reward from the taken action is received in one step, this model does not reflect the fact that the ambitious choice requires investment and is contrasted with a choice that has a reasonable payoff throughout. Similarly, the model does not capture risk associated with the ambitious choice. These are important modeling considerations for future work.

## 3 Time-Varying Subsidy

Without external intervention, once a cascade begins, it persists by definition. However, is it possible to derive an intervention from an external entity, such as the government, that can lift a population out of an incorrect cascade and redirect it toward a correct one? Importantly, can this subsidy be designed so that the correct cascade remains stable once the subsidy is removed? Note that the naïve strategy of subsidising the “correct” action fails

the sustainability criterion, because upon removal, agents have no reason to believe that that action was correct. They would simply believe that agents who took it during the subsidized period did so due to said subsidy. In this section, we focus on the question of how to design a good intervention for a single group / sequence that has *sustainable* effects.

Since a cascade occurs when an agent makes the same decision regardless of their signal, breaking a cascade requires influencing at least some agents to act according to their respective private signals. Indeed, this is our approach to designing a subsidy. We consider subsidies for the “correct” action, which, in this section, we assume to be action  $A$ . The net reward for choosing the correct action is  $R$ .

Let the subsidy the government provides at time  $t$  be  $r_t$ . Recall the two possible states of the world: world  $A$ , where  $A$  is the correct action, and world  $B$ , where  $B$  is the correct action. Let  $|A|$  indicate the number of choices for  $A$  outside of a cascade state, and similarly for  $|B|$ . Finally, we assume that the entity providing this subsidy is not trusted by the agents, and so the agents cannot infer from the direction of the subsidy which action is correct.

Algorithm 1 implements this subsidy scheme. We assume that this algorithm is applied after an incorrect cascade has already begun, and the purpose is to strategically provide a subsidy  $r_t$  to the agent acting at time  $t$  to break the community out of the pessimism trap. When agents are acting according to their own signals, or when the correct cascade has been reached, the subsidy need no longer be applied. The guarantees and derivations for Algorithm 1 are provided in Theorems 1 and 2, with complete proofs in the appendix. The main idea is that once agents act in accordance with their signals, after we see enough agents reveal their signal, a simple majority vote will, in expectation, reveal the correct action. At this point, even if the subsidy is removed, rational agents will act optimistically.

► **Theorem 1.** *The subsidy scheme used in Algorithm 1 causes all agents to act according to their signals until the population falls into the correct cascade. The subsidy value is  $r_t = R \left( \frac{\gamma_t - 1}{1 + \frac{1-p}{p}\gamma_t} \right)$ , where  $\gamma_t = \left( \frac{1-p}{p} \right)^{2|A| - t} \frac{1-p}{p}$ .*

**Proof Sketch.** We want the subsidy to incentivize taking action  $A$  only if it is already aligned with the agent’s signal. We seek to determine  $r_t$  such that it ensures action  $A$  is at least as preferable when  $S_t = A$  and less preferable when  $S_t = B$ . That is, we want  $r_t$  to satisfy the following conditions, so we simply solve:

$$\begin{aligned} (1) \quad & r_t + R \cdot \mathbb{P}[\mathbf{E}_A \mid H_{t-1}, s_t = A] \geq R \cdot \mathbb{P}[\mathbf{E}_B \mid H_{t-1}, s_t = A] \\ (2) \quad & r_t + R \cdot \mathbb{P}[\mathbf{E}_A \mid H_{t-1}, s_t = B] \leq R \cdot \mathbb{P}[\mathbf{E}_B \mid H_{t-1}, s_t = B] \end{aligned}$$

Due to the fact that the subsidy induces signal-revealing, in order to compute the posterior, we can view the history as a series of revealed signals and consider the likelihood of seeing that stream in each of the worlds. ◀

► **Theorem 2.** *In expectation, the subsidy is provided over fewer than  $\frac{4}{2p-1}$  rounds and totals no more than  $R \frac{4}{2p-1}$ .*

**Proof Sketch.** To show this, we appeal to the random walk formulation discussed above. If the subsidy value is too small, the agent will decide based on the history, if the subsidy value is too large, the agent will choose the subsidized action, and if the subsidy is “just right,” the agent will act according to their signal. The first two cases correspond to not taking any steps on the random walk, and the last case corresponds to taking a step in the direction of the signal on the random walk. Thus, we can analyze the number of total steps that need to be taken in order to net sufficiently many steps to the correct directions. We apply Wald’s equation to do this. A detailed computation can be found in Appendix D.2. ◀

■ **Algorithm 1** Redirecting Pessimism Traps.

---

**Require:** Start of incorrect cascade  $t'$ , history  $H_{t''}$  for  $t'' \geq t'$ ,  $(T - t'') \gg \frac{4}{2^p - 1}$ , correct action  $A$ , incorrect action  $B$ , private signals  $\{s\}_{t=1}^T$ , signal strength  $p$ , reward  $R$  for correct action

- 1:  $|A| \leftarrow$  choices for  $A$  in  $\bar{H}_{t'-1}$
- 2:  $|B| \leftarrow$  choices for  $B$  in  $\bar{H}_{t'-1}$
- 3:  $t \leftarrow t''$
- 4: **while**  $t < T$  **do**
- 5:   **if**  $|A| - |B| \leq -2$  (In incorrect cascade) **then**
- 6:      $\gamma_t \leftarrow \left(\frac{1-p}{p}\right)^{2|A|-t} \frac{1-p}{p}$ ;  $r_t \leftarrow R \left(\frac{\gamma_t - 1}{1 + \frac{1-p}{p} \gamma_t}\right)$
- 7:   **else**
- 8:      $r_t \leftarrow 0$
- 9:   **end if**
- 10:   **if**  $r_t + R \cdot \mathbb{P}[\mathbf{E}_A \mid H_{t-1}, s_t] > R \cdot \mathbb{P}[\mathbf{E}_B \mid H_{t-1}, s_t]$  **then**
- 11:     Agent chooses action  $A$
- 12:     **if**  $|A| - |B| \leq 1$  (Not in cascade) **then**
- 13:        $|A| = |A| + 1$
- 14:     **end if**
- 15:   **else**
- 16:     Agent chooses action  $B$
- 17:     **if**  $|A| - |B| \leq 1$  (Not in cascade) **then**
- 18:        $|B| = |B| + 1$
- 19:     **end if**
- 20:   **end if**
- 21:   Update history and increment  $t$
- 22: **end while**

---

■ **4 Extension to multiple groups**

In this section, we remove the strong assumptions of the previous section, namely that (1) the government knows the correct action and (2) everyone has the same correct action. Perhaps a more natural assumption would be that there are  $k$  groups, where each group has its own correct action (between the two, A and B). In this section, we show how to develop a subsidy scheme in which *all groups* end up in their respective correct cascade, even if the government is not privy to the respective correct actions.

In this setting, an agent in a group only sees other agents from their group. For example, consider a small town where the decisions of a student are primarily, if not completely, affected by those in the same town. Suppose the government knows the strength of signals provided to agents, but it does not know the agent's town nor which action they take. At each time  $t$ , the government may give a subsidy  $r_t$  with the goal that, eventually, all groups will reach a stable up cascade on the correct action for that group.

In this section, we show how the government can construct a distribution  $\mathcal{D}$  without knowledge of anything more than described above, such that if at each time  $r_t \sim \mathcal{D}$ , after sufficiently many steps, with high probability, all groups will have converged to their respective correct cascades. Importantly, the key idea here is the same as in Algorithm 1: one kind of subsidy that will help each group make decisions that are right for them in the long run is one that encourages an agent to reveal their signal. Further, recall that we derived exactly

a signal-revealing subsidy in the previous section. Thus, we know that for each value of  $|A| - |B|$ , i.e., for each location along the random walk, there exists a value of the subsidy that incentivizes the agent to reveal their signal. Provided this subsidy value is chosen with at least some fixed minimum probability, we can use a similar random walk analysis to before, only this time slowed down by that probability factor. In this section, we flesh this argument out. Proofs can be found in the Appendix in the full version.

## 4.1 Formally Defining The Setting

We define the setting formally below.

► **Definition 4.** *Suppose that in the world, there are two potential actions  $A$  and  $B$ . Each agent in the world has an index in  $[k]$  associated with them which we call their group. For each group in  $[k]$ , one of the actions is the “right” one and the other is the “wrong” one. Notably, which action is correct differs between groups. Formally, we may assume the correct action for a group is  $A$  with probability  $1/2$  and  $B$  with probability  $1/2$ .*

*At time  $t$ , the “universe” draws an index  $j \sim \mathcal{G}$ , where  $\mathcal{G}$  is a distribution supported on  $[k]$ , with  $j$  signifying the group index, and the minimum value of the probability mass function is  $g_{\min}$ . It is a new agent’s turn to make a decision, and they are a member of group  $j$ . This agent sees the subsidy history for all agents and need rationally only consider the action history of those in group  $j$  who have gone before them. They receive a private signal  $s_t$ , which is the correct action for the group to which they belong with probability  $p_j > 1/2$  and incorrect with probability  $1 - p_j$ . They consider their group history and private signal to update their posterior belief as shown in Equation 1, and then incorporate the present subsidy to make their final decision.*

Next, we define the subsidy scheme followed by the government. In fact, it suffices to define a uniform distribution over the possible values the subsidy would need to take, i.e., all values that would encourage signal-revealing in an agent before that agent’s group hits a cascade.

► **Definition 5.** *Let  $v_{x,p}$  be the size of subsidy that causes signal revelation by the current agent when there are  $x$  more  $A$  actions than  $B$  ones. Define distribution  $\mathcal{D}$  as supported on  $\mathcal{V}$ , all possible values of  $v_{x,p}$ , and having probability mass associated with any  $v_{x,p}$  as  $1/|\mathcal{V}|$ . The government draws the subsidy they provide from this distribution  $\mathcal{D}$ .*

## 4.2 Main Result

In this section, we show that the random process described in Definition 5 applied to the setting described in Definition 4 for a reasonable amount of time shifts groups into correct cascades with high probability.

► **Theorem 3.** *Suppose the government provides the subsidy detailed in Definition 5 for the game in Definition 4. That is, at each time step  $t > 0$ , the government (without knowledge of the group of the current agent or the history) draws a subsidy at random  $r_t \sim \mathcal{D}$ , where the probability of any element in the support of  $\mathcal{D}$  is at least  $p_{\min}$ . Then, for all  $\delta > 0$ , after  $\frac{2k}{g_{\min}} \left( \frac{2 \log(3k/\delta-1)}{\log(p/(1-p))} + \frac{2 \log(3k/\delta)}{p_{\min}^2(2p-1)^2} + \log(3k/\delta) \right)$  steps, with probability at least  $1 - \delta$ , all  $k$  groups will end up in what is for them an up-cascade.*

**Proof Sketch.** At a high level, the proof proceeds as follows: (1) We fix a group and show that the subsidy behaves, as before, as a random walk that group takes on the number line. (2) We show how long it takes to achieve with high probability a sufficient condition for the walk to finish in a cascade on the correct action. (3) We union bound over the failure probability and appropriately scale the time required to achieve the stated result.

(1) Let us first describe a random walk  $\mathcal{R}_{\text{govt}}$  that models the setting in Definition 4 with the government subsidy described in Definition 5. See the appendix for a proof. Next, we define a related but simpler random walk,  $\mathcal{R}$ . In both cases, the reverse walk describes the same for a group for whom the correct direction is the left.

► **Lemma 6.** *The following random walk, which we shall call  $\mathcal{R}_{\text{govt}}$ , describes the walk taken by a group for whom the correct action is to the right when the government provides the subsidy described in Definition 5:  $\mathbb{P}[i \text{ to } i-1] = \alpha_i \cdot (1-p)$ ;  $\mathbb{P}[i \text{ to } i] = 1 - \alpha_i$ ;  $\mathbb{P}[i \text{ to } i+1] = \alpha_i \cdot p$ , where  $\alpha_i := \mathbb{P}_{\text{govt choice}}[\text{signal is revealed in this state}]$ .*

► **Definition 7.** *Let us call the random walk with the following transition probabilities  $\mathcal{R}$ :  $\mathbb{P}[i \text{ to } i-1] = p_{\text{min}} \cdot (1-p)$ ;  $\mathbb{P}[i \text{ to } i] = 1 - p_{\text{min}}$ ;  $\mathbb{P}[i \text{ to } i+1] = p_{\text{min}} \cdot p$ .*

We can show that we can analyze  $\mathcal{R}$  instead of  $\mathcal{R}_{\text{govt}}$  (details in full version).

(2) Next, we consider three modes of failure: first, if there are not sufficiently many signals aligned to the correct direction, the majority vote will not align with the correct action for the group; second, since the government only encourages signal-revealing but does not behave differently depending on which action is correct for a group, we may accidentally hit a bad cascade before hitting a good one just due to a bad ordering in the sequence of signals; third, we may not see enough people from this group to take sufficiently many steps on the random walk. We reason formally about each of these, computing first a number of steps after which with probability  $1 - \delta/(3k)$  the walk is sufficiently far to one side and then arguing that the side to which the walk is shifted is the *correct* one for that group with probability  $1 - \delta/(3k)$ . Finally, we apply the Hoeffding bound with failure probability  $\delta/(3k)$  to ensure we see sufficiently many people. With that, the analysis is complete for a single fixed group.

(3) Finally, we union bound over the failure probability so that the result holds and report after how long of the government providing such a subsidy, with high probability *all* groups stabilize to optimism cascades. Detailed proofs for this theorem and its constituent lemmas are in the appendix of the full version. ◀

► **Remark 3.** *For sake of generality, we present the result in Theorem 3 in terms of  $p_{\text{min}}$ , and  $g_{\text{min}}$ . However, let us plug these in and discuss the scaling for intuition. First, the support of the distribution  $\mathcal{D}$  has size at most  $2L$ , so  $p_{\text{min}} \leq 1/(2L) \leq 1/(2 \frac{\log(1/\delta_0-1)}{\log(p/(1-p))})$ . Plugging this in, upper bounding, and ignoring constants, we get that the number of total required steps scales like  $\frac{kL^2}{g_{\text{min}}} \cdot \frac{\log(3k/\delta)}{(p-\frac{1}{2})^2}$ , where  $L = \frac{\log(3k/\delta-1)}{\log(p/(1-p))}$ . We can see that the number of steps scales like  $1/(p-1/2)^2$ , implying that the closer the signal strength to  $1/2$ , the longer it takes to drift far enough in the walk. This is standard for problems where we must distinguish whether a “coin flip” has bias  $1/2 + \epsilon$  or  $1/2 - \epsilon$ . Next, we see the standard dependence on the failure probability,  $\log(1/\delta)$ , and union bound,  $\log 3k$ . Finally, we see the inverse dependence on  $g_{\text{min}}$ , meaning that the lower the minimum probability of seeing a group, the longer this subsidy needs to be in place. Now, if the probability of an agent*

belonging to a group is uniform across the  $k$  groups, then  $g_{\min} = 1/k$ , and so the scaling is like  $L^2 k^2 \cdot \frac{\log(3k/\delta)}{(p-\frac{1}{2})^2}$ . From this, we can see that the dominant dependence of this bound on the number of groups is through the prevalence of the lowest-prevalence group and making sure each group takes sufficiently many steps.

## 5 Experiments

Finally, we conducted simulations to assess the effectiveness of financial supplements in overcoming pessimism traps in our theoretical model. We study the first setting, with a single group and correct action known to government. The primary objectives of our experiments were to evaluate the impact of financial supplements on the proportion of correct cascades, analyze the magnitude of the subsidies, and understand the scalability of interventions across population sizes.

### 5.1 Data Generation and Procedure


The data for our experiments were generated through simulations that model the sequential decision-making process of agents. Each agent must choose between two actions,  $A$  and  $B$ , where  $A$  is uniformly the better option, though this is unknown to the agents *a priori*. The agents receive private signals indicating the correctness of their choice, with a probability  $p$  of being correct. Each simulation was conducted as follows: We initialized a population of  $N \in \{10, 100, 1000\}$  agents to explore the effects of population size on eventual cascade behavior. Each agent received a private signal with strength  $p \in [0.51, 0.99]$ . For each pair of values for  $N$  and  $p$ , we repeat the experiment 100 times and average the results. For the proportion of correct cascades, we repeat each such experiment 10 times and report one standard deviation/ $\sqrt{10}$  as the error bars. For the subsidy size, we report the standard deviation over the 100 trials /  $\sqrt{100}$  as the error bars. The agents sequentially made their decisions after observing the actions of all preceding agents. Financial supplements as derived above were introduced to influence the agents' decisions.

## 5.2 Results


### 5.2.1 Impact of Financial Supplement on Correct Cascades

Figure 1 shows the proportion of correct cascades when a financial supplement is provided (orange line) to 100 agents compared to when it is not (blue line). The results indicate a significant improvement in the proportion of correct cascades, especially at lower signal strengths. For example, with a signal strength of 0.6, the proportion of correct cascades increases markedly (including across different population sizes (10, 100, and 1000 agents), see appendix for plots) when the supplement is used. This demonstrates the benefit of the proposed financial incentive in decision-making.

In contrast, Figure 2 shows the proportion of correct cascades without the supplement for varying numbers of agents. The performance is notably low, particularly for weaker signals. This highlights the promise of financial interventions in overcoming pessimism traps and steering agents toward optimal decisions. Also observe that the number of agents in the sequence has very little effect – once a cascade forms, there is no new information to switch out of it even when there are new agents.



■ **Figure 1** Impact of Financial Supplement on Finding Correct Cascade with Supplement, 100 agents.




■ **Figure 2** Probability of Finding Correct Cascade without Supplement.

### 5.2.2 Average Subsidy Progression

To understand the dynamics of the financial supplement, we analyzed the progression of subsidies over time for different signal strengths. We present results for 100-agent cascades here, and further experiments can be found in the appendix in the full version.

Figure 3 presents the average subsidy progression for a population of 100 agents. The required subsidy stabilizes after a few agents have made their decisions, indicating that once the initial agents are guided correctly, the need for subsequent subsidies diminishes, a benefit for allocation of resources. The line for the weakest signal levels off – based on our theoretical results, we know that it takes a long time to stabilize when the signal margin is narrow.



■ **Figure 3** Average Subsidy Progression for 100 Agents.

Interestingly, the required supplement was smallest among the first few agents in the setting with the weakest private signal. We observe a spike in the supplement needed at the beginning, yet the necessary supplement typically returns to zero quite quickly. All three plots, however, show that in the case where  $p = 0.51$ , the subsidy stays consistent across the population. This is likely because the signal is weak enough that it takes too long on average for the best action to be detected compared to the number of agents. It is also interesting to observe the average starting times for the subsidies, which align with the beginning of the pessimism trap and almost always begin within the first 10 rounds.

### 5.3 Discussion and Implications

The experiments validate our theoretical predictions regarding the impact of financial supplements on decision-making under uncertainty. The simulations provide several key insights. First, the introduction of financial supplements significantly improves the proportion of correct cascades, particularly in scenarios with lower signal strength and smaller populations. This underscores the importance of targeted interventions in guiding agents towards optimal decisions. Second, the progression of subsidies indicates that early interventions are critical. Once the initial agents are influenced correctly, the necessity for continued subsidies diminishes, suggesting an efficient use of resources. Lastly, the results suggest that the scalability of financial interventions is feasible. Larger populations benefit from reduced per-capita subsidies, making such policies more practical for broader applications.

---

#### References

- 1 Daron Acemoglu and Asuman Ozdaglar. Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1(1):3–49, 2011. doi:10.1007/s13235-010-0004-1.

- 2 Lisa R. Anderson and Charles A. Holt. Information cascades in the laboratory. *American Economic Review*, 87(5):847–862, 1997.
- 3 Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Circumventing the price of anarchy: Leading dynamics to good behavior. *SIAM Journal on Computing*, 42(1):230–264, 2013. doi:10.1137/110821317.
- 4 Shlomo Benartzi, John Beshears, Katherine L. Milkman, Cass R. Sunstein, Richard H. Thaler, Maya Shankar, Will Tucker-Ray, William J. Congdon, and Steven Galing. Should governments invest more in nudging? *Psychological Science*, 28(8):1041–1055, 2017. PMID: 28581899. doi:10.1177/0956797617702501.
- 5 Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992. doi:10.1086/261849.
- 6 Christophe P. Chamley. *Rational Herds: Economic Models of Social Learning*. Cambridge University Press, Cambridge, 2003. Online publication January 2010. doi:10.1017/CB09780511616372.
- 7 Stephen Coate and Glenn C. Loury. Will affirmative-action policies eliminate negative stereotypes? *The American Economic Review*, 83(5):1220–1240, 1993. URL: <http://www.jstor.org/stable/2117558>.
- 8 Geoffrey L. Cohen, Julio Garcia, Nancy Apfel, and Allison Master. Reducing the racial achievement gap: A social-psychological intervention. *Science*, 313(5791):1307–1310, 2006. doi:10.1126/science.1128317.
- 9 Eric J. Johnson, Suzanne B. Shu, Benedict G. C. Dellaert, Craig Fox, Daniel G. Goldstein, Gerald Häubl, Richard P. Larrick, John W. Payne, Ellen Peters, David Schkade, Brian Wansink, and Elke U. Weber. Beyond nudges: Tools of a choice architecture. *Marketing Letters*, 23(2):487–504, 2012. URL: <http://www.jstor.org/stable/23259227>.
- 10 Daniel Kahneman and Amos Tversky. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.
- 11 Jennifer M. Morton. Resisting pessimism traps: The limits of believing in oneself. *Philosophy and Phenomenological Research*, 104(3):728–746, 2022. doi:10.1111/phpr.12809.
- 12 Elchanan Mossel, Manuel Mueller-Frank, Allan Sly, and Omer Tamuz. Social learning equilibria. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, page 639, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3219166.3219207.
- 13 Lones Smith and Peter Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000. URL: <http://www.jstor.org/stable/2999431>.
- 14 Alina Sirbu, Vittorio Loreto, Vito D. P. Servedio, and Francesca Tria. Opinion dynamics: Models, extensions and external effects. In Vittorio Loreto, Muki Haklay, Andreas Hotho, Vito D.P. Servedio, Gerd Stumme, Jan Theunis, and Francesca Tria, editors, *Participatory Sensing, Opinions and Collective Awareness*, pages 363–401. Springer International Publishing, 2017. doi:10.1007/978-3-319-25658-0\_17.
- 15 Richard H. Thaler and Cass R. Sunstein. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press, 2008.

## **A** Key features of pessimism traps

Morton’s epistemic characterization of pessimism traps builds upon several empirical characterizations from scholars studying the phenomenon in the field of education. Morton elucidates pessimism traps via the following key distinguishing features [pgs. 732-3 of [11], paraphrased here]:

1. Making the ambitious choice is an investment, involving a long duration before payoff and hard work.
2. There is a feasible alternative that still has a reasonable payoff.

3. There is risk involved in the ambitious choice, enough that it would affect the choice of a risk-averse agent.
4. There is strong evidence that succeeding in the ambitious end when coming from the group of interest is of low likelihood.
5. Not pursuing the ambitious end will not change the pessimistic view they hold about it.

Importantly, the choices in question aren't even necessarily the "right" and "wrong" ones, but rather they are "ambitious" and "moderate."

## B Proof for Posterior Update (Eqn. 1)

**Proof.** First, applying Bayes' theorem and given the conditional independence of  $s_t$  and  $\bar{H}_{t-1}$  when conditioned on the best action, we have:

$$\mathbb{P}[\mathbf{E}_A | s_t, \bar{H}_{t-1}] = \frac{\mathbb{P}[s_t | \mathbf{E}_A] \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_A] \mathbb{P}[\mathbf{E}_A]}{\mathbb{P}[s_t, \bar{H}_{t-1}]}$$

Expanding the denominator using the law of total probability:

$$\begin{aligned} \mathbb{P}[s_t, \bar{H}_{t-1}] &= \mathbb{P}[\mathbf{E}_A] \mathbb{P}[s_t, \bar{H}_{t-1} | \mathbf{E}_A] + \mathbb{P}[\mathbf{E}_B] \mathbb{P}[s_t, \bar{H}_{t-1} | \mathbf{E}_B] \\ \text{cond. indep. } \Rightarrow &= \mathbb{P}[\mathbf{E}_A] \mathbb{P}[s_t | \mathbf{E}_A] \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_A] + \mathbb{P}[\mathbf{E}_B] \mathbb{P}[s_t | \mathbf{E}_B] \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_B] \end{aligned}$$

Finally, because we assume a uniform prior over which action is correct, we have  $\mathbb{P}[\mathbf{E}_A] = \mathbb{P}[\mathbf{E}_B]$  and can thus cancel out all of those terms. ◀

## C Formulation as a Random Walk

Calculating the posterior can be difficult depending on the tie-breaking rule, as agents have to reverse engineer previous individuals' thought processes. However, with the tie-breaking rule we employ and our eventual subsidy, we can reformulate this process as a one-dimensional random walk on the integers. We begin with a toy example for intuition.

► **Lemma 8** (Asymmetric Simple Random Walk in -1 to 1). *While the number of choices for action A and action B in the history  $\bar{H}_{t-1}$  differ by no more than 1, agent  $t$  will follow their signal. Thus, the process is an asymmetric simple random walk in this interval.*

**Proof.** We will prove this through induction. Our inductive hypothesis is that if the previous  $t - 1$  agents acted according to their signal, and if the number of choices of action A and B differ by at most one, then agent  $t$  will follow their signal. To begin, we establish the base case. Consider the first agent. Note that the hypothesis that the previous  $t - 1$  agents acted according to their signal is vacuously true. Now, let us look at the two signal options for the first agent to verify that they indeed follow their signal in each case. If the first agent receives signal  $s_1 = A$ , they will update their posterior, according to Eqn. 1 as:

$$\mathbb{P}[\mathbf{E}_A | s_1 = A, H_0 = \{\}] = \frac{\mathbb{P}[s_1 = A | \mathbf{E}_A] \mathbb{P}[H_0 = \{ } | \mathbf{E}_A]}{\sum_{w \in \{A, B\}} \mathbb{P}[s_1 = A | \mathbf{E}_w] \mathbb{P}[H_0 = \{ } | \mathbf{E}_w]}$$

Now, because  $H_0$  is always  $\{\}$ , we can set  $\mathbb{P}[H_0 = \{ } | \mathbf{E}_A] = \mathbb{P}[H_0 = \{ } | \mathbf{E}_B] = 1$  and simplify the above expression to:

$$\mathbb{P}[\mathbf{E}_A | s_1 = A, H_0] = \frac{\mathbb{P}[s_1 = A | \mathbf{E}_A]}{\mathbb{P}[s_1 = A | \mathbf{E}_A] + \mathbb{P}[s_1 = A | \mathbf{E}_B]} = \frac{p}{p + (1 - p)} = p$$

## 5:16 Pessimism Traps and Algorithmic Interventions

Because  $p > \frac{1}{2}$ , the first agent will choose action  $A$  in this case, thus following their signal. If, instead the agent receives signal  $B$ , they will update their posterior as:

$$\mathbb{P}[\mathbf{E}_A | s_1 = B, H_0] = \frac{\mathbb{P}[s_1 = B | \mathbf{E}_A]}{\mathbb{P}[s_1 = B | \mathbf{E}_A] + \mathbb{P}[s_1 = B | \mathbf{E}_B]} = \frac{1-p}{1-p+p} = 1-p$$

In this case, they will choose action  $B$ , again following their signal. From this, we can say that the first agent's action is identical to the action of their signal and so is distributed as a Bernoulli random variable with success parameter  $p$ .

This satisfies the base case. Now, we proceed by induction. Suppose that the previous  $t-1$  agents acted according to their signal, i.e.,  $H_{t-1} = \bar{H}_{t-1}$ . Let  $|A|$  indicate the number of choices for  $A$  observed among the first  $t-1$  agents, and let  $|B|$  be the number of choices for  $B$  among those agents. Then, if the  $t^{\text{th}}$  agent gets signal  $A$ , their posterior is:

$$\mathbb{P}[\mathbf{E}_A | s_t = A, H_{t-1}] = \frac{p^{|A|+1}(1-p)^{|B|}}{p^{|A|+1}(1-p)^{|B|} + p^{|B|}(1-p)^{|A|+1}}$$

If  $|A| = |B|$ , then the posterior is equal to  $p$ , and the agent follows their signal. If  $|A| = |B| - 1$ , the posterior is equal to  $\frac{1}{2}$ , and the agent breaks the tie by following their signal. Finally, if  $|A| = |B| + 1$ , the posterior is  $\frac{p^2}{p^2 + (1-p)^2}$ , which is greater than  $\frac{1}{2}$  for  $p > \frac{1}{2}$ .

Equivalently, if they get the signal  $B$ , the posterior is:

$$\mathbb{P}[\mathbf{E}_A | s_t = B, H_{t-1}] = \frac{p^{|A|}(1-p)^{|B|+1}}{p^{|A|}(1-p)^{|B|+1} + p^{|B|+1}(1-p)^{|A|}} \quad \blacktriangleleft$$

► **Lemma 9** (Stopping Points at -2 or 2). *A cascade will begin if the number of choices for action  $A$  and action  $B$  in the history  $H_{t-1}$  differ by 2.*

**Proof.** To see this, assume that at time  $t$ ,  $|A| - |B| \geq 2$  or  $|A| - |B| \leq -2$ . Consider the first case. If agent  $t$  gets signal  $A$ , their posterior is

$$\frac{p^{|A|-|B|+1}}{1-p} \geq \frac{p^3}{1-p} \geq 1$$

and so they take action  $A$ . If they get signal  $B$ , their posterior is

$$\frac{p^{|A|-|B|-1}}{1-p} \geq \frac{p}{1-p} \geq 1$$

So, regardless of their signal, they choose action  $A$ . Thus, a cascade begins, because this condition will be maintained for each subsequent agent. An equivalent analysis applies to the second case. ◀

► **Lemma 10** (Probability of Wrong Cascade). *The probability of an incorrect cascade for  $p > \frac{1}{2}$  is*

$$\frac{\left(\frac{1-p}{p}\right)^2 + \left(\frac{1-p}{p}\right)^3}{1 + \left(\frac{1-p}{p}\right) + \left(\frac{1-p}{p}\right)^2 + \left(\frac{1-p}{p}\right)^3}.$$

**Proof.** We can compute this probability via a recurrence. Consider a random walk on the number line. Define  $X_i$  as the probability a walk starting at  $i$  reaches 2 before it reaches -2. Then, we have that:

$$X_i = pX_{i+1} + (1-p)X_{i-1} \quad ; \quad X_2 = 1 \quad ; \quad X_{-2} = 0$$

We can solve this system as follows:

$$\begin{aligned}
Y_i &:= X_i - X_{i-1} \\
pX_i + (1-p)X_i &= pX_{i+1} + (1-p)X_{i-1} \Leftrightarrow (1-p)Y_i = pY_{i+1} \\
Y_i &= Y_0 \left(\frac{1-p}{p}\right)^i \\
\sum_{i=-1}^2 Y_i = X_2 - X_{-2} = 1 &\Rightarrow Y_0 = \frac{1}{\sum_{i=-1}^2 \left(\frac{1-p}{p}\right)^i} \\
X_i = X_{-2} + \sum_{j=-1}^i Y_j &= \frac{1}{\sum_{j=-1}^2 \left(\frac{1-p}{p}\right)^j} \sum_{k=-1}^i \left(\frac{1-p}{p}\right)^k.
\end{aligned}$$

Simplifying this somewhat, we have that starting at 0:

$$\begin{aligned}
\mathbb{P}[\text{up cascade}] &= \frac{1 + \left(\frac{1-p}{p}\right)}{1 + \left(\frac{1-p}{p}\right) + \left(\frac{1-p}{p}\right)^2 + \left(\frac{1-p}{p}\right)^3} \\
\mathbb{P}[\text{down cascade}] &= \frac{\left(\frac{1-p}{p}\right)^2 + \left(\frac{1-p}{p}\right)^3}{1 + \left(\frac{1-p}{p}\right) + \left(\frac{1-p}{p}\right)^2 + \left(\frac{1-p}{p}\right)^3} \quad \blacktriangleleft
\end{aligned}$$

Therefore, there is a substantial probability of an incorrect cascade, increasing as  $p$  gets closer to  $\frac{1}{2}$ .

► **Lemma 11** (Expected Starting Time). *In expectation, it will take  $\frac{2}{1-2p+2p^3}$  agents for a cascade to begin.*

**Proof.** To compute this, let us consider the same random walk defined in Lemma 10. With respect to this random walk, let us now define  $Z_i = \mathbb{E}[\text{time to } \pm 2 \text{ starting from } i]$ . That is, we are interested in the expected time it takes to converge to a cascade. We can write the following equations and initial conditions ( $Z_2 = 0; Z_{-2} = 0$ ) and solve manually:

$$\begin{aligned}
Z_i &= p(Z_{i+1} + 1) + (1-p)(Z_{i-1} + 1) = 1 + p(Z_{i+1}) + (1-p)(Z_{i-1}) \\
Z_0 &= 1 + p(Z_1) + (1-p)(Z_{-1}) \quad ; \quad Z_1 = 1 + p(Z_2) + (1-p)Z_0 = 1 + (1-p)Z_0 \\
Z_{-1} &= 1 + pZ_0 + (1-p)Z_{-2} = 1 + pZ_0 \\
Z_0 &= 1 + p(1 + (1-p)Z_0) + (1-p)(1 + pZ_0) = 2 + 2p(1-p)Z_0 = \frac{2}{1-2p(1-p)}. \quad \blacktriangleleft
\end{aligned}$$

## D Time-Varying Subsidy Proofs

### D.1 Proof of Theorem 1

**Proof.** Consider the scenario where the net expected reward for choosing action  $A$  includes both the base reward  $R$  and a potential subsidy  $r_t$ :

$$(R + r_t) \mathbb{P}[\mathbf{E}_A | \bar{H}_{t-1}, s_t] + r_t \cdot \mathbb{P}[\mathbf{E}_B | \bar{H}_{t-1}, s_t] = r_t + R \cdot \mathbb{P}[\mathbf{E}_A | H_{t-1}, S_t].$$

Similarly, the reward for choosing action  $B$  is:

$$R \cdot \mathbb{P}[\mathbf{E}_B | \bar{H}_{t-1}, s_t]$$

## 5:18 Pessimism Traps and Algorithmic Interventions

We seek to determine  $r_t$  such that it ensures action  $A$  is at least as preferable when  $S_t = A$  and less preferable when  $S_t = B$ . This leads to the following conditions:

- (1)  $r_t + R \cdot \mathbb{P}[\mathbf{E}_A | \bar{H}_{t-1}, s_t = A] \geq R \cdot \mathbb{P}[\mathbf{E}_B | \bar{H}_{t-1}, s_t = A]$
- (2)  $r_t + R \cdot \mathbb{P}[\mathbf{E}_A | \bar{H}_{t-1}, s_t = B] \leq R \cdot \mathbb{P}[\mathbf{E}_B | \bar{H}_{t-1}, s_t = B]$

Simplifying these inequalities, we find:

$$\begin{aligned} r_t &\geq R \left( \mathbb{P}[\mathbf{E}_B | \bar{H}_{t-1}, s_t = A] - \mathbb{P}[\mathbf{E}_A | \bar{H}_{t-1}, s_t = A] \right), \\ r_t &\leq R \left( \mathbb{P}[\mathbf{E}_B | \bar{H}_{t-1}, s_t = B] - \mathbb{P}[\mathbf{E}_A | \bar{H}_{t-1}, s_t = B] \right). \end{aligned}$$

Now, let us rewrite the probabilities by expanding the posterior in terms of the evidence (and recalling that the prior over best actions is uniform)

$$\begin{aligned} \mathbb{P}[\mathbf{E}_A | \bar{H}_{t-1}, s_t = A] &= \frac{\mathbb{P}[\bar{H}_{t-1}, s_t = A | \mathbf{E}_A]}{\mathbb{P}[\bar{H}_{t-1}, s_t = A | \mathbf{E}_A] + \mathbb{P}[\bar{H}_{t-1}, s_t = A | \mathbf{E}_B]} \\ &= \frac{p \cdot \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_A]}{p \cdot \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_A] + (1-p) \cdot \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_B]} \end{aligned}$$

$$\begin{aligned} \mathbb{P}[\mathbf{E}_A | \bar{H}_{t-1}, s_t = B] &= \frac{\mathbb{P}[\bar{H}_{t-1}, s_t = B | \mathbf{E}_A]}{\mathbb{P}[\bar{H}_{t-1}, s_t = B | \mathbf{E}_A] + \mathbb{P}[\bar{H}_{t-1}, s_t = B | \mathbf{E}_B]} \\ &= \frac{(1-p) \cdot \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_A]}{p \cdot \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_A] + (1-p) \cdot \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_B]} \end{aligned}$$

$$\begin{aligned} \mathbb{P}[\mathbf{E}_B | \bar{H}_{t-1}, s_t = A] &= \frac{\mathbb{P}[\bar{H}_{t-1}, s_t = A | \mathbf{E}_B]}{\mathbb{P}[\bar{H}_{t-1}, s_t = A | \mathbf{E}_A] + \mathbb{P}[\bar{H}_{t-1}, s_t = A | \mathbf{E}_B]} \\ &= \frac{p \cdot \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_B]}{p \cdot \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_A] + (1-p) \cdot \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_B]} \end{aligned}$$

$$\begin{aligned} \mathbb{P}[\mathbf{E}_B | \bar{H}_{t-1}, s_t = B] &= \frac{\mathbb{P}[\bar{H}_{t-1}, s_t = B | \mathbf{E}_B]}{\mathbb{P}[\bar{H}_{t-1}, s_t = B | \mathbf{E}_B] + \mathbb{P}[\bar{H}_{t-1}, s_t = B | \mathbf{E}_A]} \\ &= \frac{(1-p) \cdot \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_B]}{p \cdot \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_A] + (1-p) \cdot \mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_B]} \end{aligned}$$

Suppose  $|A|$  people have decided on  $A$  in the history  $\bar{H}_{t-1}$  (excluding any decisions made during a cascade). We define:

$$\gamma_t := \frac{\mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_B]}{\mathbb{P}[\bar{H}_{t-1} | \mathbf{E}_A]} = \frac{(1-p)^{|A|} p^{t-1-|A|}}{p^{|A|} (1-p)^{t-1-|A|}} = \left( \frac{1-p}{p} \right)^{2|A|-t} \frac{1-p}{p}.$$

Then, we can re-write the above expressions as

$$\mathbb{P}[\mathbf{E}_A | \bar{H}_{t-1}, s_t = A] = \frac{1}{1 + \frac{1-p}{p} \gamma_t} \quad ; \quad \mathbb{P}[\mathbf{E}_B | \bar{H}_{t-1}, s_t = A] = \frac{\gamma_t}{1 + \frac{1-p}{p} \gamma_t}$$

Similarly, we can re-write the expressions in our second condition as

$$\mathbb{P}[\mathbf{E}_A | \bar{H}_{t-1}, s_t = B] = \frac{1}{\frac{p}{1-p} + \gamma_t} \quad ; \quad \mathbb{P}[\mathbf{E}_B | \bar{H}_{t-1}, s_t = B] = \frac{\gamma_t}{\frac{p}{1-p} + \gamma_t}$$

$$\text{This implies that } R \left( \frac{\gamma_t - 1}{1 + \frac{1-p}{p} \gamma_t} \right) \leq r_t \leq R \left( \frac{\gamma_t - 1}{\frac{p}{1-p} + \gamma_t} \right)$$

Finally, we show that there is always a non-negative value for the subsidy in this interval. Although this is not strictly required, we want to avoid a situation where an external entity is actually taking money from the agents. To do this, it is sufficient to see that  $\gamma_t \geq 1$  in the region where the subsidy is needed to cause agents to act according to their signal and when the correct cascade has not yet been reached. By Lemmas 8 and 9, this is the region where  $|B| - |A| \geq 2$ . Then,

$$\gamma_t = \frac{(1-p)^{|A|} p^{t-1-|A|}}{p^{|A|} (1-p)^{t-1-|A|}} = \frac{(1-p)^{|A|} p^{|B|}}{p^{|A|} (1-p)^{|B|}} = \frac{p^{|B|-|A|}}{1-p} \geq 1 \quad \blacktriangleleft$$

## D.2 Proof of Theorem 2

**Proof.** We apply Wald's first equation to upper bound the amount of time for which the subsidy must be released. This time, we imagine that we must move from a state where there are two more choices for  $B$  than  $A$  to a state where there are two more choices for  $A$  than  $B$ . This reduces to the problem of a random walk with probability  $p$  of moving right hitting 2 when starting from -2 (with no left stopping point). Let  $N$  be the stopping time at which this event occurs. Using Wald's equation, we have  $\mathbb{E}[S_N] = 4$  and  $\mathbb{E}[X_1] = p \cdot 1 + (1-p) \cdot -1 = 2p-1$ . Then,  $\mathbb{E}[N] = \frac{\mathbb{E}[S_N]}{\mathbb{E}[X_1]} = \frac{4}{2p-1}$ . The corresponding subsidy calculation considers the worst-case scenario where each round necessitates the maximum subsidy  $R$ , leading to:

$$\mathbb{E}[\text{Total subsidy}] = R \times E(N) = R \times \frac{4}{2p-1}. \quad \blacktriangleleft$$