PHBS WORKING PAPER SERIES

Information Design for Social Learning with Patient Agents

Chen Lyu
Peking University

Lu Liao
Shandong University School of Economics

November 2025

Working Paper 20251104

Abstract

Sequentially coming agents may adopt a new technology. Early adoption can generate information about its value, which is either high or low, and an intermediary decides how such information will be dynamically published. Because individuals tend to wait and free-ride on information generated by others, efficient social learning is hard to achieve. Facing this challenge, we study how the intermediary can improve social welfare by designing its information publishing policy. To incentivize early adoption, we show it is optimal to restrain future information flow via inducing individually sub-optimal adoption but not via excessive waiting. The optimal design features a simple threshold stopping structure: in every period, recommend adoption if the intermediary's current belief is more optimistic than a threshold; otherwise, recommend waiting forever. While the first-best design uses a constant threshold, the optimal design features time-varying thresholds that typically cross the first-best one. We also examine special cases where learning is via conclusive news. In good-news environments, the optimal design needs to involve a middle phase in which exploration is randomly terminated; in bad-news environments, adoption may be continued even if bad news has arrived. These serve to mitigate the individuals' incentive problem efficiently.

Keywords: information design, Bayesian persuasion, collective experimentation, long-lived agent, bandit problem
JEL Classification: D83, D82, C73

Peking University HSBC Business School
University Town, Nanshan District
Shenzhen 518055, China

Information Design for Social Learning with Patient Agents*

Chen Lyu† Lu Liao‡

First version: August 15, 2025

This version: October 25, 2025

(click here for the latest version)

Abstract

Sequentially coming agents may adopt a new technology. Early adoption can generate information about its value, which is either high or low, and an intermediary decides how such information will be dynamically published. Because individuals tend to wait and free-ride on information generated by others, efficient social learning is hard to achieve. Facing this challenge, we study how the intermediary can improve social welfare by designing its information publishing policy. To incentivize early adoption, we show it is optimal to restrain future information flow via inducing individually sub-optimal adoption but not via excessive waiting. The optimal design features a simple threshold stopping structure: in every period, recommend adoption if the intermediary's current belief is more optimistic than a threshold; otherwise, recommend waiting forever. While the first-best design uses a constant threshold, the optimal design features time-varying thresholds that typically cross the first-best one. We also examine special cases where learning is via conclusive news. In good-news environments, the optimal design needs to involve a middle phase in which exploration is randomly terminated; in bad-news environments, adoption may be continued even if bad news has arrived. These serve to mitigate the individuals' incentive problem efficiently.

Key words: information design, Bayesian persuasion, collective experimentation, long-lived agent, bandit problem

*We are grateful to helpful discussions with Andriy Zapechelnyuk, Yeon-Koo Che, Hülya Eraslan, Jan Knoepfle, Yingni Guo, Junjie Zhou and Ilwoo Hwang. All mistakes are our own.

†Peking University HSBC Business School, clyu@phbs.pku.edu.cn

‡Shandong University School of Economics, lu.liao@sdu.edu.cn

1 Introduction

Dynamic social (collective) learning often plays a key role in people's learning about the value of new things, be it a new technology, a new medical treatment, or a new product. For examples, manufacturers adopting a new green technology early may generate lessons for others about whether they should also upgrade to that technology in the next facility renovation; patients accepting a new therapy in its early stage help to yield crucial information for evaluating the therapy, which guides later patients; early readers of a book may provide feedback about its quality that later potential readers can refer to. In these practices, adoption by early agents can generate useful information for their later peers. Concerning social efficiency, a positive externality problem then arises, especially when the agents are not in urgent need. Everyone will have a tendency to wait for and free-ride on information generated by others.

Fortunately, in many applications, there is an intermediary who can control or influence how early information is communicated to later agents. For new technologies or medical treatments, it may be an industry association or government agency, who collects information and provides practical guidance; for products like books, it may be an online platform (e.g., Goodreads) who elicits feedback and publishes recommendations. Such an intermediary may help to achieve more efficient information generation and usage by properly designing its dynamic information provision policy. What does the optimal design look like? How should it cope with the agents' incentive problem? How is it different from the first-best design ignoring agents' incentive? These are our central questions.

In the model, a new technology's value may be either high or low. In every period, a unit mass of agents arrive, who may adopt the technology right away or wait. Those who adopt will leave with expected utility equaling to the technology's value; those who wait will face the same choices next period, with future utility being discounted. When adoption happens, signals about the technology's value will be generated and observed by an intermediary. In every period, based on previous signals, the intermediary can publish a guiding message, which will remain publicly available for all later cohorts. We assume the intermediary can commit to an arbitrary information publishing policy ex-ante, and all agents will act Bayes-rationally. Our design problem is to find such a policy to maximize a discounted sum of all cohorts' welfare.

Because all agents active in a period face the same decision problem, by the revelation principle, it suffices for us to consider *recommendation* policies, i.e., to recommend either adopting or waiting in every period, subject to the incentive compatibility (IC) constraints that agents will always follow. Without the IC constraints, the design problem can be treated as a standard two-arm bandit problem. The constraints, however, significantly complicate it because each constraint involves the agents' expectation about the intermediary's current belief and the entire future information flow. To overcome the challenge, we first show it is optimal to use *no-pause policies*, which never induces temporary stops in adoption.[1] Under such policies, every cohort should either be incentivized to adopt right after arrival or never adopt. This allows us to simplify the IC constraints and transform the problem into a simpler constrained Markov decision process. A duality approach then helps to characterize the optimal design and enables numerical solution.

How should the optimal design incentivize early agents to adopt rather than wait? A natural idea is that we must restrain the future information flow to reduce the value of waiting. Intuitively there are two ways to do so:

A. Under-recommending: when the past signals are positive enough to make adoption individually optimal, sometimes recommend waiting (may be permanent) instead.
B. Over-recommending: when the past signals are negative enough to make it individually optimal to wait, sometimes recommend adoption instead.

We show it is optimal to only rely on over-recommending. In particular, any optimal no-pause policy should recommend adoption as long as it is myopically optimal. It thus never under-recommends. Intuitively, although both approaches can restrain the value of public information flow, over-recommending generates more information and maintains more opportunity of continuing adoption in the future. It is thus dynamically preferred. Moreover, in our application, over-recommending alone suffices to limit the dynamic information flow enough. The optimal design can hence solely rely on it.

The optimal policy is shown to feature a simple threshold structure. In every period, adoption should be continued when the current intermediary's belief about high value is above a certain threshold. Notice that the first-best policy neglecting IC constraints also has a similar structure. However, the first-best policy features a constant threshold, while the optimal policy generally needs to use time-varying thresholds in order to solve the incentive problem. The optimal threshold typically starts above the first-best one, and will later reach somewhere below it before finally turning back to it. This dynamic pattern of recommendation standard helps to boost up early cohorts' expected utility from recommended adoption and restrict the speed of information flow. The fact that the standard can sometimes be even lower than the first-best one is a novel feature of our optimal design, compared to that in related studies focusing on short-lived agents (Kremer et al., 2014; Che & Hörner, 2018; Lyu, 2023). We explain this in details at the end of Section 4.

The optimal design can be solved (semi-)explicitly when the intermediary's learning is via conclusive news. Such learning environments are often considered in the literature for its tractability and ease to interpret. We hence also provide detailed analyses for them.

In the good-news environment, agents' adoption in every period generates either no news or a conclusive piece of news that confirms the technology's value being high. Ab-

[1] Under mild conditions, one can further show that all other policies are strictly suboptimal.

sence of news then makes the intermediary increasingly pessimistic over time. In this environment, the first-best design is well-known to have two phases. The first phase features full exploration, i.e., adoption always happens even without good news; the second phase features full exploitation, i.e., adoption continues only if good news has arrived. In contrast, we show that the optimal incentive-compatible design in general needs to involve another *partial-exploration* phase in between, where adoption randomizes between stopping and continuing absent good news. This makes individually sub-optimal exploration terminate in a soft and gradual manner, which avoids the abrupt information release between the two phases in the first-best design. Comparative statics regarding this middle phase can also be studied. In particular, we show that when the agents become more patient, the partial-exploration phase should end (weakly) later.

In the bad-news environment, agents' adoption generates either no news or a piece of bad news confirming the technology's value being low. In non-trivial scenarios, the first-best design will keep adopting without news and stop adoption right away after bad news. In sharp contrast, the optimal policy may keep recommending adoption even after bad news has arrived. This helps to limit the speed of information flow and thereby dis-incentivize waiting. Our result suggests randomly delaying the release of bad news can be an efficient way to improve social learning.

Our analyses can also be directly applied to certain settings beyond our main model. In particular, one may consider what if the agents also have another option to adopt, or what if the designer is not sure about how well agents can observe past recommendations. We discuss how our results can remain meaningful in these settings in Section 7.

Related literature – This paper is closely related to Kremer et al. (2014), Che & Hörner (2018) and Lyu (2023), who also study information design for improving social learning of sequentially coming agents, where early adoption generates informational externality. In particular, Lyu (2023) also handles general non-conclusive signals. The main difference is that the current paper studies long-lived agents who can wait, while those earlier papers study short-lived agents who either adopt or quit right after arrival. This makes the central dynamic incentive problem in our setting absent in those earlier papers. Moreover, Kremer et al. (2014) and Lyu (2023) consider private recommendations, while this paper considers public communication.

Also closely related is Chen et al. (2024), who also study information design for improving social learning among long-lived agents. The main difference is that we consider multiple cohorts of agents coming in sequence with the same preference, while they consider agents all present since the beginning with heterogeneous values. This leads to very different optimal designs. For example, with conclusive news learning, their optimal design fully reveals past information whenever adoption is continued, which easily contrasts with our design mentioned earlier. Moreover, Chen et al. (2024) focuses on learning with conclusive news, while we also consider general non-conclusive signals. Also related is Knoepfle & Salmi (2024). They study optimal evidence disclosure instead of information design a la Bayesian persuasion — their designer must either disclose nothing or fully disclose the evidence it has observed. Moreover, like Chen et al. (2024), they focus on conclusive-news learning and agents present since the beginning (with heterogeneous discounting instead of values).

Generally, our paper belongs to the broad literature on information design (Kamenica & Gentzkow, 2011; Rayo & Segal, 2010), and especially to studies on dynamic designs (e.g., Ely, 2017; Renault et al., 2017; Smolin, 2021; Ely & Szydlowski, 2020; Ball, 2019; Orlov et al., 2020; Lorecchio, 2021). One difference between our paper and many studies in this literature is that we consider a designer whose private information flow is controlled by the receivers' decisions, rather than being exogenous.

The paper also relates to the broad literature about social experimentation. In particular, Frick & Ishii (2024) studies a similar learning environment. They focus on characterizing the equilibrium outcome, but do not consider optimal design for information transmission.

The paper is arranged as follows: Section 2 presents the model; Section 3 provides the main characterization for optimal design; Section 4 reveals major properties of the optimal design; Sections 5 and 6 consider learning with conclusive news; Section 7 discusses additional issues and extensions. Proofs are provided in Appendix A.

2 The Model

A new technology is available for adoption. Its true value $\theta \in \{L, H\}$ is unknown, where $L < 0 < H$ and the society shares common prior p_0 about $\theta = H$. In every period $t = 1, 2, ...$, a new cohort of agents with unit mass arrive. They can decide whether to adopt the technology or wait. Once an agent adopts, she leaves with expected utility equaling to θ; an agent who has not adopted can always wait towards next period, and the flow utility during a waiting period is normalized to zero.[2] Each cohort of agents discount their future utility with factor $\delta_A \in (0, 1)$. We assume that whenever being indifferent, an agent will choose to adopt rather than wait.[3]

Adoption in every period will generate information about θ that can be observed by an intermediary, which we call the designer. We assume each agent is informationally small in the sense that her individual decision has zero probability to affect the information being

[2]If there is a flow waiting utility b, one can normalize it to zero and redefine θ to be $\theta - \frac{b}{1-\delta_A}$.

[3]This guarantees that the designer's problem introduced below has a solution. Alternatively, one can assume that the designer can choose what an agent will do when being indifferent. In this case, it is without loss of optimality for the designer to ask an indifferent agent to adopt, since adoption does not hurt this agent and helps to generate information faster.

generated. Together with the agents' tie-breaking rule mentioned before, this implies that all agents in the same cohort will always move in the same way.[4] It then suffices for us to specify what information will be generated when an entire cohort of agents adopt. Let s_i denote the signal generated from cohort i's adoption. We assume $(s_i)_{i=1}^{\infty}$ are conditionally i.i.d. given θ. We can also allow the designer to have an initial piece of signal s_0 reflecting its internal research or data, which is independent from $(s_i)_{i=1}^{\infty}$ conditional on θ. Notice all these signals are observable to the designer, but not directly to the agents.

In every period t, let p_t denote the designer's initial belief about $\theta = H$. Let μ_1 denote the distribution of $p_1 := \mathbb{P}(\theta = H|s_0)$. For any $k \in \mathbb{N}$, let $G^k(\cdot|\cdot)$ denote the transition probability of the belief reflecting Bayesian updating when k cohorts of agents adopt. For $k = 1$, we will simply write G^1 as G. The process of $(p_t)_{t=1}^{\infty}$ then follows transition rule:

$$p_1 \sim \mu_1 \tag{2.1}$$

$$p_{t+1}|p_t, K_t \sim G^{K_t}(\cdot|p_t) \tag{2.2}$$

where K_t denotes the number of cohorts adopting in period t. Let $u(p_t) := Hp_t + L(1-p_t)$, i.e., the expected value of θ given p_t. We impose the following assumption throughout:

Assumption 1. (i) $\mathbb{P}\big(u(p_1) > 0\big) > 0$; (ii) $G(\cdot|p)$ is non-degenerate for any $p \notin \{0, 1\}$.

Condition (i) guarantees that for at least some p_1 it is myopically strictly optimal for the first cohort to adopt. If it fails, we can never induce adoption with strictly positive surplus, which trivializes the designer's problem.[5] Condition (ii) guarantees $(s_i)_{i=1}^{\infty}$ are indeed informative.

At the beginning of every period, the designer can send a public message based on previous signals, which will remain observable in all later periods. We assume the designer can commit to such an information transmission policy ex-ante (before the realization of s_0), and the agents will react Bayes-rationally. Since all agents remaining in a period face the same binary-choice problem with the same tie-breaking rule, the revelation principle (Myerson, 1986) implies that it suffices for the designer to send binary recommendations – "adopt" or "wait" – subject to incentive compatibility requirement that all present agents want to follow. We use $a_t \in \{0, 1\}$ to denote the message sent in period t, where $a_t = 1$ means "adopt". In general, for any vector x, we will use $x_{<t}$ to denote $(x_{t'} : t' < t)$. Similarly is $x_{\leq t}$ defined.

We assume the designer wants to maximize the total discounted expected utility of all cohorts. Specifically, let U_t denote the expected utility of cohort t. That is, if χ_t denotes

[4] For informationally small agents, they want to move differently only when they are indifferent, but this implies that they all choose to adopt according to the tie-breaking rule.

[5] If the condition fails, then the first adopting agents must have non-positive surplus. For them not to deviate to waiting, all later cohorts must also have non-positive surplus.

the adoption time of cohort t, then $U_t = \mathbb{E}[\delta_A^{\chi_t - t}\theta]$. Then the designer wants to maximize:

$$\sum_{t=1}^{\infty} \delta_D^t U_t$$

where $\delta_D \in (0,1)$ represents the designer's cross-cohort discounting. Since the belief process summarizes past signals, as usual it suffices to consider *belief-based* policies. Such a policy is represented by a sequence of measurable mappings $\phi := (\phi_t)_{t=1}^{\infty}$, where $\phi_t : [0,1]^t \times \{0,1\}^{t-1} \to [0,1]$. Given any $p_{\le t}$ and recommendation history $a_{<t}$, $\phi_t(p_{\le t}, a_{<t})$ decides the probability of $a_t = 1$. In the main text below, we will always focus on these belief-based policies. In some proofs, however, it will be more convenient to work on policies directly conditioning on past signals. We provide explicit measure-theoretic definition of those general policies in Appendix A.1.1.

3 Characterization for the Optimal Design

3.1 Formulation of the designer's problem

Let Φ denote the set of all policies. Provided that the agents follow the recommendations, any $\phi \in \Phi$ will induce a probability measure over events about $\big(\theta, (a_t)_{t=1}^{\infty}, (p_t)_{t=1}^{\infty}, (K_t)_{t=1}^{\infty}\big)$. We use $\mathbb{P}_\phi$ and $\mathbb{E}_\phi$ to denote the corresponding probability measure and expectation operator.

Given the process of $(a_t)_{t=1}^{\infty}$, define $\nu(t) := \min\{m \ge t : a_m = 1\}$. Then $\nu(t)$ denotes the first time at which adoption is recommended since period t. If the agents follow recommendations, we then have $U_t = \mathbb{E}_\phi[\delta_A^{\nu(t)-t}\theta]$. The designer's objective function can thus be written as:

$$\sum_{t=1}^{\infty} \delta_D^t \mathbb{E}_\phi[\delta_A^{\nu(t)-t}\theta]$$

The agents' incentive compatibility requires that for all t and $x_{<t} \in \{0,1\}^{t-1}$:

$$\mathbb{P}_\phi(a_t = 1, a_{<t} = x_{<t}) > 0 \Rightarrow \quad \mathbb{E}_\phi[\theta | a_t = 1, a_{<t} = x_{<t}] \ge \delta_A \mathbb{E}_\phi[\delta_A^{\nu(t+1)-(t+1)}\theta | a_t = 1, a_{<t} = x_{<t}] \tag{3.1}$$

$$\mathbb{P}_\phi(a_t = 0, a_{<t} = x_{<t}) > 0 \Rightarrow \quad \mathbb{E}_\phi[\theta | a_t = 0, a_{<t} = x_{<t}] < \delta_A \mathbb{E}_\phi[\delta_A^{\nu(t+1)-(t+1)}\theta | a_t = 0, a_{<t} = x_{<t}] \tag{3.2}$$

Intuitively, in constraint (3.1), the LHS of the inequality is the agent's expected utility from adoption after observing $a_t = 1$ and $a_{<t} = x_{<t}$, while the RHS is her discounted expected utility if she deviates to waiting now but subsequently follows the recommendations. The constraint guarantees the one-shot deviation is not profitable. Similarly,

constraint (3.2) guarantees that the agents will not want to one-shot deviate to adoption after observing $a_t = 0$ and $a_{<t} = x_{<t}$. By the one-shot deviation principle, these constraints together guarantee global incentive compatibility.

Since agents do not internalize their positive informational externality, at any moment they should be less willing to adopt than the designer. This makes it reasonable to conjecture that constraint (3.2) will not be restrictive. We hence omit it for now, and will check it later. The designer's optimization can then be summarized as follows:

$$\max_{\phi \in \Phi} \sum_{t=1}^{\infty} \delta_D^t \mathbb{E}_\phi[\delta_A^{\nu(t)-t}\theta] \tag{3.3}$$

$$\text{s.t. IC constraint (3.1)} \tag{3.4}$$

3.2 The optimality of no-pause policies

Our first result shows that it is without loss of optimality to use policies featuring *no pause in adoption*. That is, once adoption is not recommended, it stops forever.

Definition 1. $\phi \in \Phi$ is a no-pause policy if for all t: $a_t = 0 \overset{\mathbb{P}_\phi\text{-}a.s.}{\Longrightarrow} a_{t'} = 0$ for all $t' > t$.[6]

Proposition 1. *For any $\phi \in \Phi$ satisfying (3.1), there exists a no-pause policy ϕ' satisfying (3.1) such that all cohorts are weakly better off under ϕ' than under ϕ.*

This result is good to know because practically no-pause policies are much simpler to implement and interpret than more general policies. To see the intuition behind, imagine the designer uses an incentive-compatible policy ϕ with pauses. Then after some history leading to a period t, she will recommend waiting while knowing that adoption will be re-recommended some time later. At such a moment, what if she switches to recommend adoption right away? Agents active in period t will be induced to adopt earlier, which benefits them since the adoption's value must be non-negative in expectation for the original policy to be incentive-compatible, and there is no need to discount anymore. For other agents, this cannot harm either since the designer can always keep them at least as happy as before with faster generated information. The only concern is earlier agents' incentive compatibility. The fact that being active in period t becomes more advantageous now makes waiting before period t more appealing. However, the designer can solve this problem by further over-recommending adoption in period t to make being active in period t as good as before. This restores incentive compatibility without eventually harming any cohort. In the proof, we show in details how ϕ' in the proposition can be constructed from ϕ using the intuition above.

Proposition 1 implies it is optimal to focus on no-pause policies. Under mild conditions, one can further show that any other policies must be sub-optimal. To state the

[6]By saying $A \Rightarrow B$ almost surely, we mean the probability $\mathbb{P}(A \setminus B) = 0$.

result, we will say signals $(s_i)_{i=0}^{\infty}$ are *negatively-inconclusive* if they never conclusively reveal $\theta = L$.[7] Let $\bar{p}$ be the myopically indifferent belief, i.e., $u(\bar{p}) = 0$. We also define:

Definition 2. Signals $(s_i)_{i=1}^{\infty}$ are *upwardly unbounded* if $G((\bar{p}, 1]|p) > 0$ for any $p > 0$.

Intuitively, the signals are upwardly unbounded if no matter how low p_t is, as long as it is non-zero, a new piece of signal may be strong enough to turn it above $\bar{p}$. We consider this as a mild condition because the chance for such a strong signal to realize can be arbitrarily small and shrink to zero arbitrarily fast when $p_t \downarrow 0$.

We have the following result:

Proposition 2. *Assume $(s_i)_{i=0}^{\infty}$ are negatively-inconclusive and $(s_i)_{i=1}^{\infty}$ are upwardly unbounded. Then any optimal policy must be a no-pause policy.*

The intuition is that under the proposition's assumptions, more or faster adoption will always have strictly positive informational value for the designer. This makes the modifications over ϕ we discussed right below Proposition 1 strictly beneficial.

3.3 The constrained Markov decision process

Due to Propositions 1 and 2, from now on we will focus on no-pause policies. Let $\Phi^{\dagger}$ denote the set of all no-pause policies. The designer's problem can be simplified to:

$$\max_{\phi \in \Phi^{\dagger}} \sum_{t=1}^{\infty} \delta_D^t \mathbb{E}_{\phi}[\theta a_t] \tag{3.5}$$

$$\text{s.t. } \mathbb{E}_{\phi}[\theta a_t] \geq \delta_A \mathbb{E}_{\phi}[\theta a_{t+1}], \forall t \tag{3.6}$$

To see the objective is correct, notice that under a no-pause policy, cohort t either adopts in period t following $a_t = 1$ or never adopts. Thus $U_t = \mathbb{E}_{\phi}[\theta a_t]$. To see the IC constraint is correct, notice under a no-pause policy, whenever $a_t = 1$, there is only one possible public history $a_1 = ... = a_{t-1} = 1$. Thus constraint (3.1) can be simplified to $\mathbb{E}_{\phi}[\theta | a_t = 1]\mathbb{P}_{\phi}(a_t = 1) \geq \delta_A \mathbb{E}_{\phi}[\delta_A^{\nu(t+1)-(t+1)}\theta | a_t = 1]\mathbb{P}_{\phi}(a_t = 1)$, which is equivalent to

$$\mathbb{E}_{\phi}[\theta a_t] \geq \delta_A \mathbb{E}_{\phi}[\delta_A^{\nu(t+1)-(t+1)}\theta a_t]$$

Moreover, with a no-pause policy, the next adoption time $\nu(t+1)$ equals to either $t+1$ (when $a_{t+1} = 1$) or ∞ (when $a_{t+1} = 0$). Thus $\delta_A \mathbb{E}_{\phi}[\delta_A^{\nu(t+1)-(t+1)}\theta a_t] = \delta_A \mathbb{E}_{\phi}[\theta a_t a_{t+1}]$, which further equals to $\delta_A \mathbb{E}_{\phi}[\theta a_{t+1}]$ since $a_{t+1} = 1 \Rightarrow a_t = 1$ a.s.

For later analyses, we also need to replace θ in the optimization with expressions about $(p_t)_{t=1}^{\infty}$. Recall that $u(p_t) := Hp_t + L(1-p_t)$. Since p_t summarizes all information about θ

[7] Formally, this means $\mu_1(\{0\}) = 0$ and $G(\{0\}|p) = 0$ for any $p > 0$.

available at time t, we have $\mathbb{E}_\phi[\theta|a_t, p_t] = \mathbb{E}_\phi[\theta|p_t] = u(p_t)$, and thus $\mathbb{E}_\phi[\theta a_t] = \mathbb{E}_\phi[a_t u(p_t)]$. Optimization (3.5) – (3.6) can then be written as:

$$\max_{\phi \in \Phi^\dagger} \sum_{t=1}^{\infty} \delta_D^t \mathbb{E}_\phi[a_t u(p_t)] \tag{3.7}$$

$$\text{s.t. } \mathbb{E}_\phi[a_t u(p_t)] \geq \delta_A \mathbb{E}_\phi[a_{t+1} u(p_{t+1})], \forall t \tag{3.8}$$

This will be the problem we work with below. For any t, we will refer the constraint in (3.8) as IC_t. Intuitively, one can understand it as a constraint preventing cohort t from mimicking cohort $t+1$, since it just requires $U_t \geq \delta_A U_{t+1}$ under ϕ.

It is also helpful to highlight the simplified transition rule of p_t here. Notice that under a no-pause policy, at any t either *one* cohort is present and adopts (when $a_t = 1$), or no one adopts (when $a_t = 0$). Transition rule (2.1) – (2.2) then becomes:

$$p_1 \sim \mu_1 \tag{3.9}$$

$$p_{t+1}|p_t, a_t \sim a_t G(\cdot|p_t) + (1 - a_t) D(\cdot|p_t) \tag{3.10}$$

where $D(\cdot|p_t)$ is the Dirac measure at p_t. Evidently, $(p_t)_{t=1}^\infty$ now follows a Markov process controlled by $(a_t)_{t=1}^\infty$. This makes optimization (3.7) – (3.8) a *constrained Markov decision process*. It is *constrained* because IC constraint (3.8) involves integrating over the belief states. It is well known that this kind of constraints cannot be directly handled with dynamic programming. To characterize the optimal design, we thus adopt a Lagrangian duality approach below.

Before proceeding, we first provide the existence result here:

Proposition 3. *Designer's optimization (3.7) – (3.8) admits a solution.*

The proof draws upon classical results in Schäl (1975) and Balder (1989), which are generally useful for proving existence related to Markov decision processes.

3.4 The dual characterization

Given any $\phi \in \Phi^\dagger$ and (current value) multiplier $\lambda = (\lambda_t)_{t=1}^\infty \in \mathbb{R}_+^\infty$, we define the Lagrangian function as:

$$\mathcal{L}(\phi, \lambda) = \sum_{t=1}^{\infty} \delta_D^t \mathbb{E}_\phi[a_t u(p_t)] + \sum_{t=1}^{\infty} \delta_D^t \lambda_t \Big[\mathbb{E}_\phi[a_t u(p_t)] - \delta_A \mathbb{E}_\phi[a_{t+1} u(p_{t+1})]\Big] \tag{3.11}$$

For simplicity, define $\lambda_0 := 0$. Then we can more compactly write:

$$\mathcal{L}(\phi, \lambda) = \sum_{t=1}^{\infty} \delta_D^t \big(1 + \lambda_t - \frac{\delta_A}{\delta_D} \lambda_{t-1}\big) \mathbb{E}_\phi[a_t u(p_t)] \tag{3.12}$$

The duality result is stated as follows.

Theorem 1. *Let v^* denote the designer's optimal value. Then,*

$$v^* = \min_{\lambda \in \mathbb{R}_+^\infty} \sup_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda) \tag{3.13}$$

The minimization has a solution, and any λ^* solving it has finitely many entries being non-zero. Given such a λ^*, ϕ^* solves designer's problem (3.7) – (3.8) if and only if:

(i) $\phi^* \in \arg\max_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda^*)$;

(ii) $\lambda_t^* \big[\mathbb{E}_{\phi^*}[a_t u(p_t)] - \delta_A \mathbb{E}_{\phi^*}[a_{t+1} u(p_{t+1})]\big] = 0$, $\forall t$;

(iii) ϕ^* *satisfies the IC constraint (3.8).*

This characterization is useful in two ways. First, it allows us to analytically examine features of the optimal policy by looking into an unconstrained problem $\max_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda^*)$ (together with the complementary-slackness condition (ii)), which can be handled by standard dynamic programming. Second, it enables numerically solving the optimal policy by first solving the dual problem, which is always convex and can be solved relatively efficiently. We purse these analyses in the next section. Some technical numerical details are discussed in Appendix B.

4 General Features of the Optimal Design

4.1 Over-recommend v.s. under-recommend

As is mentioned in the introduction, to incentivize early adoption the designer needs to restrain the future information flow. Naturally this can be done through:

A. Under-recommending: when the past signals are positive enough to make it individually optimal to adopt, sometimes recommend waiting instead. In this way, some positive news are pooled with negative news to induce individually suboptimal waiting (possibly permanent waiting).

B. Over-recommending: when the past signals are negative enough to make it individually optimal to wait, sometimes recommend adoption instead. In this way, some negative news are pooled with positive news to induce individually suboptimal adoption.

We will formally show that an optimal no-pause policy never under-recommend below. In particular, as long as $u(p_t) \geq 0$, adoption should happen.

We first provide the following lemma regarding the immediate value of improving a cohort's expected utility, which is useful for showing several results later.

Lemma 1. *For any λ^* solving the dual problem (3.13), we have* $1 + \lambda_t^* - \frac{\delta_A}{\delta_D} \lambda_{t-1}^* \geq 0$, $\forall t$.

To see the lemma's economic content, notice that when we improve cohort t's expected utility U_t, besides dynamic or informational consequences, there are three immediate effects. First, cohort t becomes better off; second, IC_t in (3.8) becomes slacker; third, IC_{t-1} becomes tighter as waiting towards period t becomes more advantageous. The first two effects are beneficial, while the third one is adverse. Despite of this conflict, Lemma 1 implies that their total marginal impact on the designer's value must be non-negative at the optimum, which is measured by $1 + \lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^*$. The lemma can be proved as a consequence of conditions (i) and (ii) in Theorem 1. In particular, if $1 + \lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^* < 0$, then whenever $u(p_t) < 0$, adoption should be continued (supposing it has not stopped) because it will be beneficial not only dynamically but also immediately in the Lagrangian optimization $\max_{\phi\in\Phi^\dagger} \mathcal{L}(\phi, \lambda^*)$. In that case, however, U_t cannot exceed U_{t-1}, which makes IC_{t-1} slack. We then must have $\lambda_{t-1}^* = 0$, contradicting with $1 + \lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^* < 0$.

Recall $\bar{p}$ is the myopically indifferent belief. We have the following result:

Proposition 4. *Under an optimal no-pause policy,* $a_{t-1} = 1$ *and* $p_t \geq \bar{p} \overset{a.s.}{\Longrightarrow} a_t = 1, \forall t$.

That is, it is optimal to keep recommending adoption as long as $u(p_t) \geq 0$. This implies that the optimal no-pause policy never restrains information flow by under-recommending adoption.[8] Given Lemma 1, a rough intuition behind is simple to see. When $u(p_t) \geq 0$, the immediate impact of adoption on the designer's value captured by $1 + \lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^*$ must be non-negative; meanwhile, the dynamic value of keeping adoption unstopped is always positive. Thus adoption should be optimally continued.

Notice Proposition 4 implies that when $a_t = 0$, $u(p_t)$ must be negative under the optimal no-pause policy. Thus constraint (3.2) we omitted before is indeed satisfied.

4.2 Threshold structure and its dynamic features

We now show the optimal policy has a simple threshold structure.

Definition 3. $\phi \in \Phi^\dagger$ is a *threshold* no-pause policy if there exist thresholds $(\eta_t)_{t=1}^\infty$ s.t. $\mathbb{P}_\phi(a_{t-1} = 1, p_t > \eta_t, a_t = 0) = \mathbb{P}_\phi(a_{t-1} = 1, p_t < \eta_t, a_t = 1) = 0\,, \forall t$.

Intuitively, a threshold no-pause policy with thresholds $(\eta_t)_{t=1}^\infty$ continues adoption when $p_t > \eta_t$, and stops adoption when p_t drops below η_t. Notice the definition does not restrict how the policy may choose or randomize when $p_t = \eta_t$.

Proposition 5. *There exists an optimal threshold no-pause policy* ϕ^**. Moreover, if* $(s_i)_{i=1}^\infty$ *are upwardly unbounded (Definition 2), then every optimal no-pause policy is equivalent to a threshold one with the same thresholds as* ϕ^**.*

[8] Under the conditions of Proposition 2, any optimal policy must be no-pause. Proposition 4 then further implies that it is sub-optimal to under-recommend even when all policies are considered.

While threshold policies are intuitively appealing and easy to use in practice, we note that their optimality is not straightforward in our setting due to the complex incentive compatibility issues. Key to the proposition's proof is still Lemma 1, which tells us that even after considering all those IC constraints through looking into the Lagrangian optimization $\max_{\phi\in\Phi^\dagger}\mathcal{L}(\phi,\lambda^*)$, inducing higher expected utility for each cohort t is still directly beneficial. This implies that both the dynamic value and the immediate value of continuing adoption in the Lagrangian optimization are increasing in p_t. It is thus optimal to use threshold policies.

Next, we look into dynamic features of the optimal thresholds and compare them with the first-best design. By first best, we mean the optimal design when agents will unconditionally follow recommendations published by the designer. Formally, it solves problem (3.7) without IC constraints in (3.8). Notice this unconstrained problem is a standard two-arm bandit problem, whose solution is well known to feature a constant stopping threshold, which we will denote as η^F.

Figure 4.1 presents a numerical example where the signals $(s_i)_{i=0}^{\infty}$ equal to the true θ plus some normal noises (see the figure caption for details). The blue bars represent the optimal thresholds, and the black dotted line plots the first-best threshold. One can see that the optimal threshold initially starts above η^F, and goes below it later, but will eventually come back and stay at η^F. Intuitively, by setting a high recommendation standard at the beginning, the designer boosts up early agents' expected utility from following an adoption recommendation; by choosing low standards later, the designer restrains later information flow and thus reduces the value of earlier waiting. These both serve to incentivize early adoption and deter waiting when the first-best design is not incentive compatible. Notice that the optimal threshold must eventually go back to η^F because only finitely many IC constraints in (3.8) can be restrictive (since U_t is bounded). Once no further IC_t is binding, the optimal threshold will equal to η^F forever.

The key features in Figure 4.1 mentioned above can be formally proved to hold more generally. For simplicity, let us assume s_i admits a density $f_{s_i|\theta}$ conditional on θ (w.r.t. some common measure over the signal realization space) for both $i = 0$ and $i \geq 1$, and impose the following assumption:

Assumption 2. For both $i = 0$ and $i \geq 1$, the log-likelihood ratio $\log(\frac{f_{s_i|H}(s_i)}{f_{s_i|L}(s_i)})$ as a random variable is finite, atomless and having full support over $\mathbb{R}$.

Under the assumption, p_t will always be atomless and have full support over $(0, 1)$ while adoption is continuing. It is atomless so that we can just focus on the threshold without specifying possible randomization at the threshold; it has full support so that the optimal thresholds will be unique.

We say s_i is *log-concave* if its log-likelihood ratio $\log(\frac{f_{s_i|H}(s_i)}{f_{s_i|L}(s_i)})$ has a log-concave density

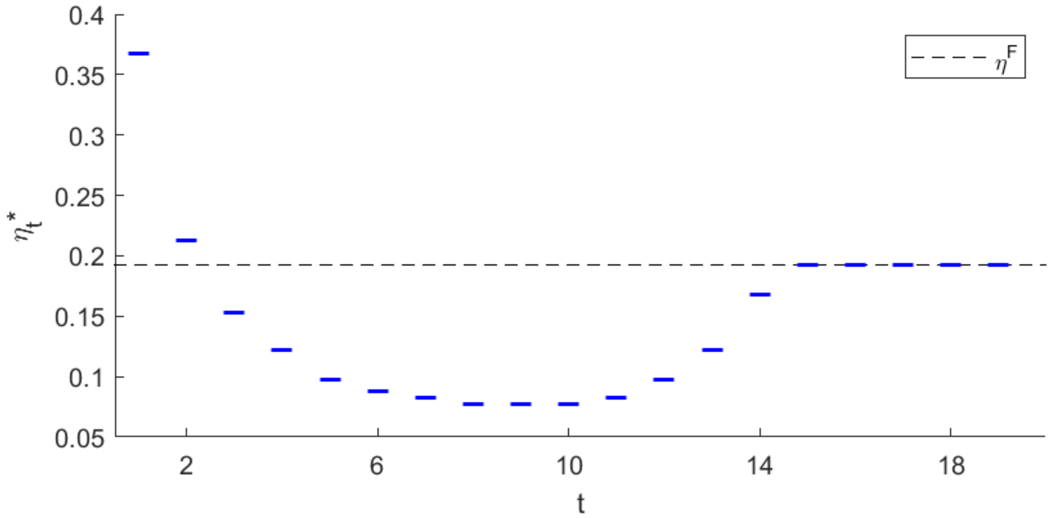

Figure 4.1: The figure compares $(\eta_t^*)_{t=1}^{\infty}$ with η^F for a numerical example, where $\delta_D = \delta_A = 0.95$, $H = 1$, $L = -1$, $p_0 = 0.6$, $s_i = \theta + 3\epsilon_i$ for all $i \geq 0$ with $\epsilon_i \overset{i.i.d.}{\sim} Normal(0, 1)$.

conditional on $\theta = L$.[9] Many common signals satisfy this property, including the one with normal noises in the example of figure 4.1. We have the following result:

Proposition 6. *Assume Assumption 2 holds. If the first-best design is not feasible, then the optimal thresholds $(\eta_t^*)_{t=1}^{\infty}$ are such that:*

(a) If $(s_i, i \geq 1)$ are log-concave, then $\eta_1^* > \eta^F$.

(b) There exists $\bar{t}$ such that $\eta_{\bar{t}}^* < \eta^F$ and $\eta_t^* = \eta^F$ for all $t > \bar{t}$.

Part (a) confirms that the optimal threshold starts above η^F under relatively general conditions. Part (b) confirms that it will later drop below η^F and then go back to eventually coincide with η^F.

One particular implication of part (b) above is that it can be optimal to continue adoption even when the designer's belief has dropped below the first-best threshold. This is a distinctive feature of our optimal design compared to those in previous papers focusing on short-lived agents (Kremer et al., 2014; Che & Hörner, 2018; Lyu, 2023). Intuitively, when agents can wait, further lowering the recommendation standard can help to reduce the value of information released and thereby deter previous waiting. This benefit is present neither when the agents are short-lived nor in the first-best problem, which explains the distinction. Actually, from the proof of Proposition 6(b), one can infer that the only motivation for the designer to set $\eta_{\bar{t}} < \eta^F$ given $\eta_t = \eta^F, \forall t > \bar{t}$ is to relax a binding IC constraint of cohort $\bar{t} - 1$.

[9] It is equivalent to require the density to be log-concave conditional on $\theta = H$, because the conditional densities given $\theta = H$ and $\theta = L$ only differ by a log-linear factor.

5 Learning with Conclusive Good News

We now specialize analysis to settings in which the designer learns via conclusive news. Such learning environments are often studied in the literature (e.g., Che & Hörner (2018)) due to its tractability and ease to interpret. We study good-news learning below. The case of bad-news learning is much simpler and will be considered in the next section.

5.1 The setting

Throughout this section, we assume $(s_i)_{i=0}^{\infty}$ take values in $\{g, null\}$ with distributions:

- $\mathbb{P}[s_0 = g|\theta = H] = \kappa_0 \geq 0$ and $\mathbb{P}[s_0 = g|\theta = L] = 0$;
- $\mathbb{P}[s_i = g|\theta = H] = \kappa > 0$ and $\mathbb{P}[s_i = g|\theta = L] = 0$, for all $i \geq 1$.

Intuitively, each signal may realize to be good news (g) or null. Good news may arrive only when $\theta = H$. Its arrival thus conclusively reveals the technology's value being high. On the other hand, persistent absence of good news will make the designer more pessimistic about θ. The central question is when the designer should give up experimenting on the new technology without good news arrival. Notice this learning environment satisfies conditions in Proposition 2. Thus only no-pause policies can be optimal.

Let p_t^n denote the designer's posterior belief about $\theta = H$ given $s_0 = ... = s_{t-1} = null$. Then by the Bayes formula $(p_t^n)_{t=1}^{\infty}$ is a decreasing sequence with:

$$p_1^n = \frac{p_0(1-\kappa_0)}{1-p_0+p_0(1-\kappa_0)};\ p_{t+1}^n = \frac{p_t^n(1-\kappa)}{1-p_t^n+p_t^n(1-\kappa)}\ \forall t \geq 1$$

In this environment, it is well known that the first-best design features a threshold $\eta^F \in (0, \bar{p})$ such that adoption is stopped absent good news once p_t^n falls below η^F. For simplicity, assume $\eta^F \leq p_1^n$ and $\eta^F \neq p_t^n\ \forall t$.[10] We define some useful times:

- $t^m = \min\{t : p_t^n < \bar{p}\}$. It is the first time adoption without good news becomes myopically suboptimal.
- $t^F = \min\{t : p_t^n < \eta^F\}$. It is the first time the first-best design stops adoption without good news.

Then $t^m \leq t^F$ since $\bar{p} > \eta^F$. To keep things interesting, we also assume the first-best policy is not incentive compatible throughout this section.

Assumption 3. The first-best policy violates IC constraint (3.8).

With any optimal no-pause policy ϕ^*, by Proposition 4 we know that adoption must be recommended forever once good news has arrived. To pin down ϕ^*, it thus suffices to

[10]If $\eta^F > p_1^n$, then the optimal design is trivially the first-best design that stops adoption as long as $s_0 \neq g$. If $\eta^F = p_t^n$ for some t, then at that t the first-best design can randomize arbitrarily absent good news. Our assumption $\eta^F \neq p_t^n\ \forall t$ eases the analysis by ruling out this indeterminacy. It holds for generic choices of δ_D given other parameters fixed.

characterize the probability of continuing adoption without good news. To ease notations, let $a_0 := 1$. For any $t \geq 1$, we define:

$$\alpha_t^{\phi} := \begin{cases} \mathbb{P}_{\phi}(a_t = 1 | a_{t-1} = 1, p_t = p_t^n) & \text{if } \mathbb{P}_{\phi}(a_{t-1} = 1, p_t = p_t^n) > 0 \\ 0 & \text{if } \mathbb{P}_{\phi}(a_{t-1} = 1, p_t = p_t^n) = 0 \end{cases} \tag{5.1}$$

It then suffices to find the optimal sequence of $(\alpha_t^{\phi})_{t=1}^{\infty}$.

5.2 Optimal design: the three-phase structure

For any ϕ, let $t_a^{\phi} := \min\{t \geq 1 : \alpha_t^{\phi} < 1\}$ and $t_b^{\phi} := \max\{t \geq 1 : \alpha_t^{\phi} > 0\}$. Intuitively, t_a^{ϕ} is the first time adoption absent good news may be stopped, and t_b^{ϕ} is the last time adoption absent good news may be continued. The following proposition shows that the optimal design generally features 3 phases, with a non-trivial randomization phase in the middle.

Proposition 7. *Let ϕ^* be an optimal no-pause policy. Then $t^m \leq t_a^{\phi^*} \leq t_b^{\phi^*} < \infty$ and:*

$$\textit{(i) } \alpha_t^{\phi^*} = 1, \forall t < t_a^{\phi^*}; \quad \textit{(ii) } \alpha_t^{\phi^*} \in (0,1), \forall t \in [t_a^{\phi^*}, t_b^{\phi^*}]; \quad \textit{(iii) } \alpha_t^{\phi^*} = 0, \forall t > t_b^{\phi^*}$$

Moreover, $[t_a^{\phi^*}, t_b^{\phi^*}] \cap \{t^F - 1, t^F\} \neq \emptyset$, and IC_t is binding under ϕ^* for all $t \in [t_a^{\phi^*}, t^F - 1] \cup [t^F - 1, t_b^{\phi^*} - 1]$.

The proposition implies that the optimal design in general has three phases. The first phase ($t < t_a^{\phi^*}$) features full exploration in the sense that adoption is always continued; the second phase ($t \in [t_a^{\phi^*}, t_b^{\phi^*}]$) features *partial* exploration as the policy may randomly stop adoption and learning absent good news; the third phase features full exploitation as adoption continues only if good news has arrived. Notice that compared to the first-best design, which stops exploration suddenly for sure at t^F, having the randomized middle phase is the most notable feature of the optimal design. This is illustrated by Figure 5.1 with a numerical example. Intuitively, the extended randomizing phase makes information release around t^F more spread over time, which softens the agents' incentive problems. In particular, starting random termination of exploration before t^F helps to boost up U_t of early cohorts in the randomizing phase, while randomly keeping exploration continued even after t^F helps to suppress U_t of the later cohorts. These together restrict the growth rate of U_t to be no greater than $\frac{1}{\delta_A}$, which makes the IC constraints in (3.8) satisfied.

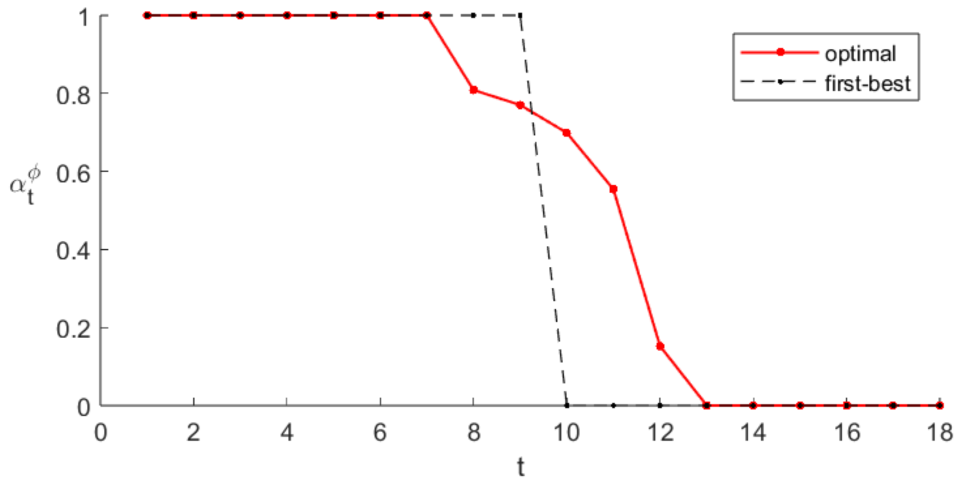

Figure 5.1: The figure compares the optimal design and the first-best design in terms of their $(\alpha_t^\phi)_{t=1}^\infty$ for a numerical example, where $\delta_D = 0.92$, $\delta_A = 0.9$, $H = 1$, $L = -1$, $p_0 = \frac{2}{3}$, $\kappa_0 = 0$, $\kappa = 0.2$.

5.3 Optimal design: the full characterization

Notice the last statement in Proposition 7 implies that for all $t \geq t_a^{\phi^*}$, either $\alpha_{t+1}^{\phi^*} = 0$, or $\alpha_{t+1}^{\phi^*}$ is positive and just large enough to make IC_t binding. This allows us to inductively compute $(\alpha_t^{\phi^*} : t > t_a^{\phi^*})$ once $(t_a^{\phi^*}, \alpha_{t_a^{\phi^*}}^{\phi^*})$ has been decided. Specifically, for any optimal no-pause ϕ^*:

Corollary 1. *Given $(t_a^{\phi^*}, \alpha_{t_a^{\phi^*}}^{\phi^*})$, $(\alpha_t^{\phi^*})_{t \neq t_a^{\phi^*}}$ can be uniquely determined as follows:*

(a) $\alpha_t^{\phi^*} = 1$ for all $t < t_a^{\phi^*}$.

(b) For each $t \geq t_a^{\phi^*}$ in sequence, given $\alpha_{\leq t}^{\phi^*}$: set $\alpha_{t+1}^{\phi^*} = 0$ if this satisfies IC_t; otherwise, set $\alpha_{t+1}^{\phi^*} > 0$ to be such that IC_t holds as equality.

The proof is provided in Appendix A.7, where we also provide detailed formula for computing $\alpha_{t+1}^{\phi^*}$ in part (b).

The corollary reduces the designer's problem into a decision problem just regarding two parameters – $(t_a^\phi, \alpha_{t_a^\phi}^\phi)$. The optimal choice of these parameters can be characterized by a single-dimension optimization over a continuous piecewise-linear quasi-concave function, which we develop in Appendix A.8. This makes it easier to compute numerical examples and enables comparative statics with respect to the agents' patience below.

5.4 Comparative statics w.r.t. δ_A

A central factor affecting the optimal design is how patient the agents are. The more patient are they, the more inclined they will be to wait for future information, which tightens the designer's IC constraints.

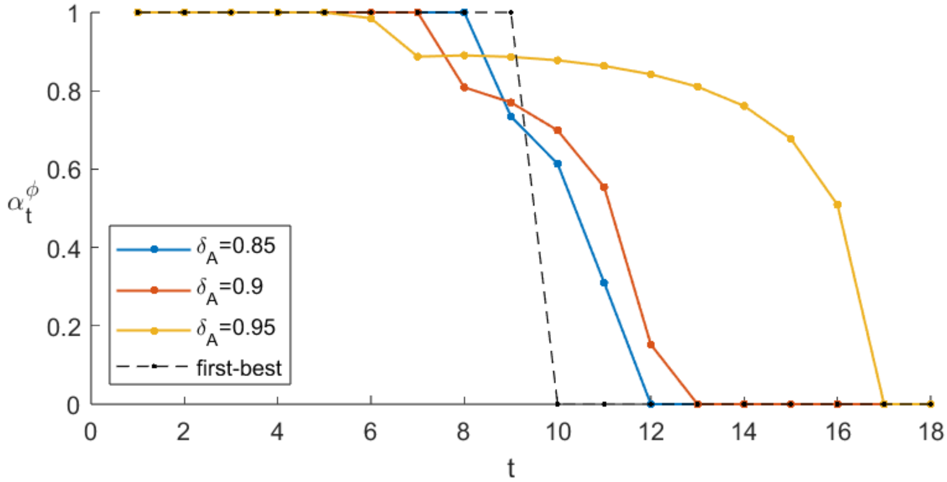

Figure 5.2: The figure compares the optimal design with different δ_A. Except for δ_A, the other parameters are the same as in Figure 5.1.

Figure 5.2 compares the optimal designs with different δ_A for a numerical example. We can see that when agents become more patient (i.e., δ_A increases), the optimal design typically needs to involve a more extended phase of partial exploration. This helps to release information more gradually before terminating experimental adoption for sure. Particularly, one phenomenon one can see in Figure 5.2 is that when δ_A increases, the partial-exploration phase will optimally end (weakly) later. This is true in general:

Proposition 8. *For any $\delta'_A < \delta''_A$, let ϕ' and ϕ'' denote the corresponding optimal no-pause policies. Then $t_b^{\phi'} \leq t_b^{\phi''}$.*

5.5 Learning outcome as $t \to \infty$

An interesting question is how well the optimal design will perform in the limit compared to the first-best design. In particular, given the true θ, how do they compare in terms of the chance of making the right recommendation when $t \to \infty$? Since the optimal design needs to obey IC constraints, one may conjecture that it has worse limit performance. However, this is not true in general.

Notice in the current good news environment, when $\theta = L$, both designs must eventually stop adoption, which leads to the same limit performance. When $\theta = H$, the chance of correctly choosing $a_t = 1$ is eventually just proportional to U_t. Figure 5.3 illustrates the path of U_t for the numerical examples in Figure 5.2. One can see that in the limit, U_t under the optimal design may be either higher or lower than that under the first best, and there is no monotone pattern regarding how the limit changes in δ_A. The optimal design may thus perform either better or worse than the first best in the limit. Intuitively, due to the randomizing phase around t^F, the optimal design may terminate exploration either earlier or later than the first best. The comparison about overall limit

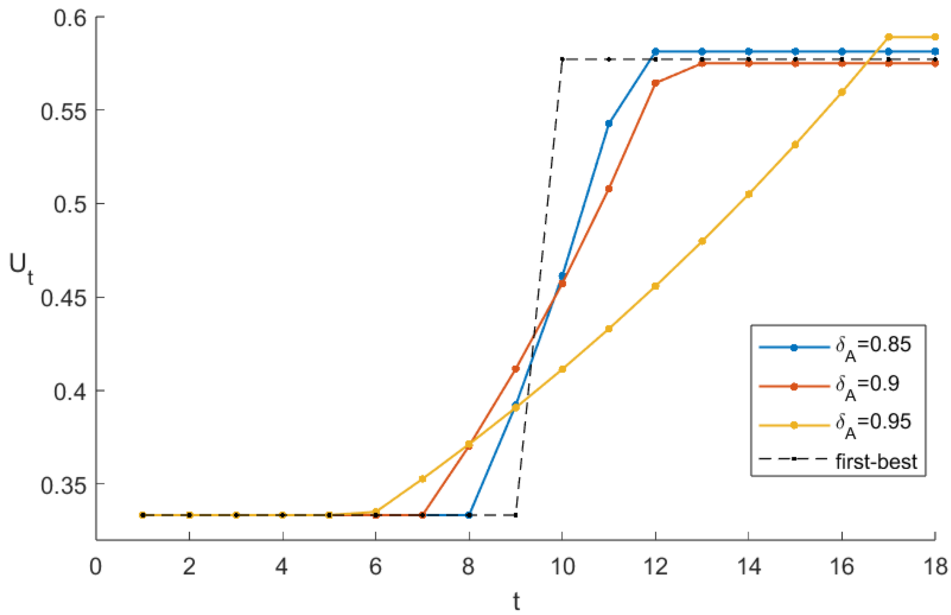

Figure 5.3: The figure compares the optimal path of U_t with different δ_A. The parameters are the same as in Figure 5.2.

learning qualities is thus unclear, which depends on detailed parameters in the model. Notice that this contrasts with what can happen when agents are short-lived (e.g., in Che & Hörner (2018)). With short-lived agents, the optimal design always has worse limit learning performance because it never stops exploration later than the first-best design.

6 Learning with Conclusive Bad News

We now turn to study the case of learning with conclusive bad news. Specifically, we assume $(s_i)_{i=1}^{\infty}$ take values in $\{b, null\}$ with distribution:

- $\mathbb{P}[s_0 = b|\theta = H] = 0$ and $\mathbb{P}[s_0 = b|\theta = L] = \kappa_0 \geq 0$;
- $\mathbb{P}[s_i = b|\theta = H] = 0$ and $\mathbb{P}[s_i = b|\theta = L] = \kappa > 0$, for all $i \geq 1$.

Intuitively, once bad news has arrived, θ is revealed to be low; absent bad news, the designer's belief will gradually become more optimistic. For simplicity, we will still focus on no-pause policies.[11]

Let p_t^n denote the designer's posterior belief about $\theta = H$ given $s_0 = ... = s_{t-1} = null$. Then by the Bayes formula, $(p_t^n)_{t=1}^{\infty}$ is an increasing sequence with $p_1^n = \frac{p_0}{p_0+(1-p_0)(1-\kappa_0)}$ and $p_{t+1}^n = \frac{p_t}{p_t+(1-p_t)(1-\kappa)}$ $\forall t \geq 1$. To obey Assumption 1(a), we assume $p_1^n > \bar{p}$. Proposition 4 then directly implies that adoption should be continued as long as no bad news has realized. It then suffices to pin down the adoption probability after bad news arrival.

[11]Notice this environment does not satisfy conditions in Proposition 2. There can be other policies with pauses that are also optimal.

Still let $a_0 := 1$. For any $t \geq 1$, we define:

$$\beta_t^\phi := \begin{cases} \mathbb{P}_\phi(a_t = 1 | a_{t-1} = 1, p_t = 0) & \text{if } \mathbb{P}_\phi(a_{t-1} = 1, p_t = 0) > 0 \\ 0 & \text{if } \mathbb{P}_\phi(a_{t-1} = 1, p_t = 0) = 0 \end{cases} \tag{6.1}$$

β_t^ϕ then denotes the probability of continuing adoption in period t after bad news arrival. The optimal $(\beta_t^\phi)_{t=1}^\infty$ can be computed sequentially to keep the IC constraints satisfied.

Proposition 9. *The optimal $(\beta_t^\phi)_{t=1}^\infty$ can be uniquely computed as follows:*

(a) $\beta_1^\phi = 0$.

(b) For every $t \geq 1$, given $\beta_{\leq t}^\phi$: set $\beta_{t+1}^\phi = 0$ if this satisfies IC_t; otherwise, set $\beta_{t+1}^\phi > 0$ to be such that IC_t holds as equality.

The proof is provided in Appendix A.9, where we also provide detailed induction formula for deciding β_{t+1}^ϕ in part (b).

The most notable feature of the optimal design here, compared to the first best, is that it may keep recommending adoption even after bad news arrival. This can never be optimal without the IC constraints, since bad news concludes learning with $\theta = L$. With the IC constraints, however, the optimal design needs to deter early waiting by slowing down future information release. The result above suggests withholding bad news is an efficient way to do so. Notice by the detailed rule in Proposition 9, $\beta_t^\phi > 0$ is optimal only if it makes IC_{t-1} binding. This implies β_t^ϕ can never equal to 1. Thus any withholding of bad news should be done in a random manner.

7 Discussion and Extension

7.1 Another option for adopting

In our model, each active agent chooses between waiting (i.e., keeping his status-quo situation) and adopting a new technology. In some applications, there might also be another technology that can be adopted in a once-for-all manner. For example, a patient with a chronic disease may have three options: (1) keeping taking a traditional medicine (the status quo); (2) accepting a newly developed surgery; (3) accepting a traditional surgery. When the third option exists, the situation does not directly match our model. However, when this third option's value is already well known, our result can be immediately extended. The key is to notice that, under our optimal no-pause design, an agent either adopts immediately or never adopts, which makes his option value from waiting never really valuable. Thus when the third option's expected value exceeds the value from permanent waiting, it is always favored and will replace the role of permanent waiting when designing the optimal policy.

Formally, consider the same setting as in Section 2 except for that in every period, an active agent may also choose to adopt another old technology and leave with utility u_0 — we will call this the *two-tech* setting. To keep things non-trivial, assume $u_0 > 0$ and $\mathbb{P}\big(u(p_1) > u_0\big) > 0$. Also consider another *auxiliary setting*, which is the same as that in Section 2 except for that the flow utility of waiting is set to $(1 - \delta_A)u_0$ instead of being normalized to zero. As we mentioned in footnote 2, our previous analyses easily accommodate this. Let ϕ^{au} denote an optimal no-pause policy in the auxiliary setting. We then have the following result:

Proposition 10. *In the two-tech setting, an optimal policy ϕ^* is described as follows: recommend adopting the new technology whenever ϕ^{au} does so; otherwise, recommend adopting the old technology.*

Intuitive Proof. First, notice that under any information design, an agent must be weakly better off in the auxiliary setting than in the two-tech setting. This is because waiting is more profitable in the auxiliary setting and permanent waiting there is surplus-equivalent to adopting the old technology. We then have two implications:

(a) The designer's optimal value in the auxiliary setting must be weakly higher.

(b) Every agent in the auxiliary setting under ϕ^{au} must be weakly better off than in the two-tech setting under ϕ^* proposed above (since the two policies are informationally equivalent).

However, notice that in the two-tech setting under ϕ^*, if one follows the designer's recommendations, he will receive the same expected utility as in the auxiliary setting under ϕ^{au}. Thus ϕ^* is incentive-compatible and indeed optimal for the designer in the two-tech setting. *Q.E.D.*

The analysis above relies on the assumption that the second technology's value is already known. If the second technology one may adopt is also new and of uncertain value, the design problem will become much more complex — one needs to properly trade off between explorations of the two technologies while dealing with the forward-looking agents' incentive problem. We leave this interesting problem to future research.

7.2 Information about past messages and design robustness

Our analyses have assumed all active agents know all past messages sent by the designer. What if they only have limited information about past messages, due to limited memory or non-trivial cost in obtaining the full public record? In that situation, the agents' incentive problem will become less severe for two reasons: (1) multiple IC constraints corresponding to different message histories can be consolidated into a single one; (2) waiting becomes less valuable when past memory may be lost. Because of these, if the designer knows exactly what information the agents will have about past messages and can tailor the design correspondingly, she may in general be able to achieve better performance than that under the optimal design in our setting.

In practice, however, it is typically hard for the designer to know exactly what information agents will actually have about past messages. If this is the case, then our design will regain its optimality in an informationally robust sense. When agents cannot perfectly observe or recall past messages, since their incentive constraints will just become more relaxed for reasons mentioned above, our optimal design will remain incentive compatible and implement the same allocation no matter what agents actually know about past messages.[12] It thus provides the best performance guarantee when the designer is unsure about how well agents can observe or memorize past messages.

8 Concluding Remarks

We have studied how an information intermediary can improve social learning among sequentially coming agents via its dynamic information publishing policy, where the agents are long-lived and have a tendency to wait for information generated by others. To incentivize socially beneficial early adoption, the optimal policy restrains future information flow via over-recommending adoption, but in general not via under-recommending or inducing temporary waiting. The optimal policy features a simple threshold stopping structure. It stops adoption once the intermediary's current belief becomes more pessimistic than a time-varying threshold, and otherwise keeps recommending for adoption. One interesting property is that the optimal recommendation threshold in general needs to go even below the first-best threshold for some periods, which is never optimal if the agents are short-lived.

There are many directions remaining for future research. Besides the one mentioned at the end of Section 7.1, another direction we want to highlight is to study detail-robust designs. Currently, our optimal design depends on details about the learning environment, including the signal structures and agents' detailed preference. In practice, knowing these accurately may be hard, and we often want to have a design that is overall performing well, at least when the true environment is not too extreme. This makes robustness important. Theoretically, it is also interesting to see how a robust design (in whatever sense) will compare with the exact optimal design in our setting, in terms of their performances and qualitative features. We think our study may serve as a natural benchmark for this practical research agenda.

[12] A direct way to check this is to notice that under our design: if the latest recommendation is "adopt", agents know all past messages must be "adopt" and there is no need to see them directly; if the latest recommendation is "waiting", then current $u(p_t) < 0$ by Proposition 4 and one definitely wants to wait.

A Proofs

A.1 Proof for Proposition 1

To prove Proposition 1, it is more convenient to work with general policies directly conditioning on past signals and certain randomization devices (instead of restricting to belief-based ones). We thus first define these general policies below.

A.1.1 General policies

Fix an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ that is general enough to host all random objects to be introduced below.

First, θ and $(s_i)_{i=0}^{\infty}$ are defined as in the model.

Second, two sequences of randomization devices $(\epsilon_t^1)_{t=1}^{\infty}$ and $(\epsilon_t^2)_{t=1}^{\infty}$ are defined. All these terms are i.i.d. draws from Uniform$[0,1]$ and are independent from all other random objects.

Finally, a (general) recommendation policy is defined as a process $(a_t)_{t=1}^{\infty}$ taking values in $\{0,1\}^{\infty}$ such that:

- $$\forall t \text{ and } x_{\le t} \in \{0,1\}^t,\ \{a_{\le t} = x_{\le t}\} \in \underbrace{\sigma\Big((s_i)_{i=0}^{m(x_{<t})}, (\epsilon_t^1)_{t=1}^{\infty}, (\epsilon_t^2)_{t=1}^{\infty}\Big)}_{=:\mathcal{F}(x_{<t})} \tag{A.1}$$

where $m(x_{<t}) := \max\Big\{0, \max\{m < t : x_m = 1\}\Big\}$. Intuitively, given any history $a_{<t} = x_{<t}$, $m(x_{<t})$ tells how many cohorts have adopted before. The measurability requirement then guarantees that the choice of a_t can only depend on signals generated from cohorts having already adopted (as well as the randomization devices), but not on signals not generated yet. To ease notations, for any $x_{<t}$, we denote the σ-field $\sigma\Big((s_i)_{i=0}^{m(x_{<t})}, (\epsilon_t^1)_{t=1}^{\infty}, (\epsilon_t^2)_{t=1}^{\infty}\Big)$ as $\mathcal{F}(x_{<t})$.

We note that since the designer needs to randomize no more than once every period, referring to one sequence of randomization devices is sufficient (for inducing any joint distribution over θ, $(s_i)_{i=0}^{\infty}$ and the adoption process). However, having another sequence of randomization device will be helpful in the proof later.

A.1.2 The proof

Fix any policy $(a_t^1)_{t=1}^{\infty}$ satisfying IC constraint (3.1) and featuring pausing, i.e., $\mathbb{P}(a_{t'}^1 = 0, a_{t''}^1 = 1) > 0$ for some $t'' > t'$. W.l.g., assume $(a_t^1)_{t=1}^{\infty}$ only uses $(\epsilon_t^1)_{t=1}^{\infty}$ for randomization and is independent from $(\epsilon_t^2)_{t=1}^{\infty}$. Construct another policy $(a_t^2)_{t=1}^{\infty}$ as follows:

- *For any t, $a_t^2 = \mathbb{1}\{a_{t'}^1 = 1$ for some $t' \ge t\}$ ($\mathbb{1}$ is the indicator function).*

Then $(a_t^2)_{t=1}^{\infty}$ features no pause since by construction $a_t^2 = 0 \Rightarrow a_{t'}^1 = 0\, \forall t' \ge t \Rightarrow a_{t'}^2 = 0\, \forall t' \ge t$. To see it is informationally feasible, we need to check the measurability condition (A.1). We do this by induction. When $t = 1$, we need to check $\{a_1^2 = 1\} \in \sigma\Big(s_0, (\epsilon_t^1)_{t=1}^\infty, (\epsilon_t^2)_{t=1}^\infty\Big)$. To see this, notice by construction $\{a_1^2 = 1\} = \cup_{t' \geq 1}\{a_1^1 = ... = a_{t'-1}^1 = 0,\, a_{t'}^1 = 1\}$. Since $(a_t^1)_{t=1}^\infty$ satisfies condition (A.1), each of these sub-events belongs to $\sigma\Big(s_0, (\epsilon_t^1)_{t=1}^\infty, (\epsilon_t^2)_{t=1}^\infty\Big)$. Thus so is $\{a_1^2 = 1\}$. Now, assume the condition holds for period $t-1$. Then $\{a_{<t}^2 = x_{<t}\} \in \mathcal{F}(x_{<t})$. It suffices to check $\{a_{<t}^2 = x_{<t}, a_t^2 = 1\} \in \mathcal{F}(x_{<t})$ for any $x_{<t}$. Notice:

$$
\begin{aligned}
\{a_{<t}^2 = x_{<t}, a_t^2 = 1\} &= \bigcup_{t' \geq t}\Big(\{a_{<t}^2 = x_{<t}\}\bigcap\{a_t^1 = ... = a_{t'-1}^1 = 0,\, a_{t'}^1 = 1\}\Big) \\
&= \bigcup_{t' \geq t}\Big(\{a_{<t}^2 = x_{<t}\}\bigcap\big(\cup_{x'_{<t} \leq x_{<t}}\{a_{<t}^1 = x'_{<t}, a_t^1 = ... = a_{t'-1}^1 = 0,\, a_{t'}^1 = 1\}\big)\Big)
\end{aligned}
$$

The first equality holds by the construction of a_t^2. The second equality holds because by construction $a_t^2 \geq a_t^1\, \forall t$, and thus on the event $\{a_{<t}^2 = x_{<t}\}$ we must have $a_{<t}^1$ equals to some $x'_{<t} \leq x_{<t}$.[13] Now, notice $(a_t^1)_{t=1}^\infty$ satisfying (A.1) implies that for any t', $\{a_{<t}^1 = x'_{<t}, a_t^1 = ... = a_{t'-1}^1 = 0,\, a_{t'}^1 = 1\} \in \mathcal{F}(x'_{<t}) \subset \mathcal{F}(x_{<t})$, where the set inclusion holds because $x'_{<t} \leq x_{<t} \Rightarrow m(x'_{<t}) \leq m(x_{<t})$. Thus the expression above is indeed in $\mathcal{F}(x_{<t})$, which implies condition (A.1) holds for period t.

Now, let U_t^i denote cohort i's expected utility under $(a_t^i)_{t=1}^\infty$. Notice incentive compatibility of $(a_t^1)_{t=1}^\infty$ implies $U_t^1 \geq 0\, \forall t$; the fact that $(a_t^2)_{t=1}^\infty$ features no pause implies $U_t^2 = \mathbb{E}[\theta a_t^2]\, \forall t$. We argue $U_t^2 \geq U_t^1\, \forall t$. To see this, notice:

$$
\begin{aligned}
U_t^2 = \mathbb{E}[\theta a_t^2] = \mathbb{E}[\theta \mathbb{1}\{a_{t'}^1 = 1 \text{ for some } t' \geq t\}] &= \mathbb{E}[\sum_{t' \geq t} \theta \mathbb{1}\{a_t^1 = ... = a_{t'-1}^1 = 0,\, a_{t'}^1 = 1\}] \\
&= \sum_{t' \geq t} \mathbb{E}[\theta \mathbb{1}\{a_t^1 = ... = a_{t'-1}^1 = 0,\, a_{t'}^1 = 1\}] \\
&\geq \sum_{t' \geq t} \delta_A^{t'-t}\mathbb{E}[\theta \mathbb{1}\{a_t^1 = ... = a_{t'-1}^1 = 0,\, a_{t'}^1 = 1\}] = U_t^1 \qquad \text{(A.2)}
\end{aligned}
$$

The second equality holds by the definition of $(a_t^2)_{t=1}^\infty$. The third equality is a trivial identity. The fourth equality holds because the infinite summation is uniformly bounded by $[L, H]$ and thus can be interchanged with the expectation. The inequality holds because $(a_t^1)_{t=1}^\infty$ being incentive compatible implies that $\mathbb{E}[\theta \mathbb{1}\{a_t^1 = ... = a_{t'-1}^1 = 0,\, a_{t'}^1 = 1\}] \geq 0\, \forall t' \geq t$.[14]

Now, notice that $(a_t^2)_{t=1}^\infty$ may not satisfy constraint (3.1). To find an incentive compatible improvement over $(a_t^1)_{t=1}^\infty$, define another policy $(a_t^3)_{t=1}^\infty$ inductively as follows.

- *Let $a_1^3 = a_1^2$. Given $a_{<t}^3$ having been well defined:*
    - *If $\delta_A \mathbb{E}[\theta a_t^2] \leq \mathbb{E}[\theta a_{t-1}^3]$, then define $a_t^3 = a_t^2$;*

[13] By inequality between two vectors, we mean component-wise inequality.

[14] Otherwise, we must have $\mathbb{P}(a_t^1 = ... = a_{t'-1}^1 = 0,\, a_{t'}^1 = 1) > 0$ and after observing $a_t^1 = ... = a_{t'-1}^1 = 0,\, a_{t'}^1 = 1$ it is overall better for the agents not to adopt at t'.

– *If $\delta_A \mathbb{E}[\theta a_t^2] > \mathbb{E}[\theta a_{t-1}^3]$, then define $a_t^3 = \mathbb{1}\{a_t^2 = 1\} + \mathbb{1}\{a_t^2 = 0, a_{t-1}^3 = 1, \epsilon_t^2 \leq y_t\}$, where $y_t \in (0,1]$ is chosen such that $\delta_A \mathbb{E}[\theta a_t^3] = \mathbb{E}[\theta a_{t-1}^3]$. To see such an y_t exists, notice:*
  - *(i) Since ϵ_t^2 is an independent draw from $U[0,1]$, the value of $\delta_A \mathbb{E}[\theta a_t^3]$ is continuous in y_t.*
  - *(ii) When $y_t = 0$, $a_t^3 = a_t^2$ and we have $\delta_A \mathbb{E}[\theta a_t^3] > \mathbb{E}[\theta a_{t-1}^3]$.*
  - *(iii) When $y_t = 1$, $a_t^3 = \mathbb{1}\{a_t^2 = 1 \text{ or } a_{t-1}^3 = 1\} = a_{t-1}^3$ a.s., where the second equality holds because $a_t^2 = 1 \Rightarrow a_{t-1}^2 = 1$ a.s. (since $(a_t^2)_{t=1}^\infty$ features no pause) and $a_{t-1}^2 = 1 \Rightarrow a_{t-1}^3 = 1$ (by previous construction). Thus when $y_t = 1$, $\delta_A \mathbb{E}[\theta a_t^3] = \delta_A \mathbb{E}[\theta a_{t-1}^3]$. Moreover, by previous construction, for any $t' < t$, $\mathbb{E}[\theta a_{t'}^3]$ equals to either $\mathbb{E}[\theta a_{t'}^2]$ or $\frac{1}{\delta_A}\mathbb{E}[\theta a_{t'-1}^3]$. The fact that $\mathbb{E}[\theta a_{t'}^2] = U_{t'}^2 \geq 0 \, \forall t'$ then implies $\mathbb{E}[\theta a_{t'}^3] \geq 0 \, \forall t' < t$. Thus $\delta_A \mathbb{E}[\theta a_t^3] = \delta_A \mathbb{E}[\theta a_{t-1}^3] \leq \mathbb{E}[\theta a_{t-1}^3]$.*

Let U_t^3 denote cohort t's expected utility under $(a_t^3)_{t=1}^\infty$. With this construction, we can show the following.

First, $(a_t^3)_{t=1}^\infty$ is informationally feasible, i.e., it satisfies condition (A.1). The inductive proof is similar to that for $(a_t^2)_{t=1}^\infty$. When $t = 1$, (A.1) holds trivially since $a_1^3 = a_1^2$. Now assume the condition holds for period $t-1$. Then $\{a_{<t}^3 = x_{<t}\} \in \mathcal{F}(x_{<t})$. It suffices to check $\{a_{<t}^3 = x_{<t}, a_t^3 = 1\} \in \mathcal{F}(x_{<t})$. There are two possible cases:

1. $\delta_A \mathbb{E}[\theta a_t^2] \leq \mathbb{E}[\theta a_{t-1}^3]$. In this case, $a_t^3 = a_t^2$ and thus

$$
\begin{aligned}
\{a_{<t}^3 = x_{<t}, a_t^3 = 1\} &= \{a_{<t}^3 = x_{<t}\} \bigcap \{a_t^2 = 1\} \\
&= \{a_{<t}^3 = x_{<t}\} \bigcap \big(\cup_{x'_{<t} \leq x_{<t}} \{a_{<t}^2 = x'_{<t}, a_t^2 = 1\} \big) \qquad \text{(A.3)}
\end{aligned}
$$

   The second equality holds because by construction $a_t^3 \geq a_t^2$ and thus in the event $\{a_{<t}^3 = x_{<t}\}$ we must have $a_{<t}^2$ equals to some $x'_{<t} \leq x_{<t}$. Notice for any $x'_{<t} \leq x_{<t}$, $\{a_{<t}^2 = x'_{<t}, a_t^2 = 1\} \in \mathcal{F}(x'_{<t}) \subset \mathcal{F}(x_{<t})$. The expression in (A.3) thus belongs to $\mathcal{F}(x_{<t})$.

2. $\delta_A \mathbb{E}[\theta a_t^2] > \mathbb{E}[\theta a_{t-1}^3]$. In this case by construction we have:

$$
\begin{aligned}
\{a_{<t}^3 = x_{<t}, a_t^3 = 1\} = &\Big(\{a_{<t}^3 = x_{<t}\} \cap \{a_t^2 = 1\}\Big) \\
&\bigcup \Big(\{a_{<t}^3 = x_{<t}\} \cap \{a_t^2 = 0\} \cap \{a_{t-1}^3 = 1\} \cap \{\epsilon_t^2 \leq y_t\}\Big) \quad \text{(A.4)}
\end{aligned}
$$

   The argument in case 1 has already shown $\big(\{a_{<t}^3 = x_{<t}\} \cap \{a_t^2 = 1\}\big) \in \mathcal{F}(x_{<t})$. This further implies $\big(\{a_{<t}^3 = x_{<t}\} \cap \{a_t^2 = 0\}\big) \in \mathcal{F}(x_{<t})$. Moreover, the induction hypothesis implies $\{a_{<t}^3 = x_{<t}\} \cap \{a_{t-1}^3 = 1\} \in \mathcal{F}(x_{<t})$. Finally $\{\epsilon_t^2 \leq y_t\} \in \mathcal{F}(x_{<t})$ trivially. The expression in (A.4) thus belongs to $\mathcal{F}(x_{<t})$.

Second, $(a_t^3)_{t=1}^\infty$ features no pause. To see this, notice $a_{t-1}^3 = 0$ implies $a_{t-1}^2 = 0$, which further implies $a_t^2 = 0\, a.s.$ since $(a_t^2)_{t=1}^\infty$ has no pause. Thus whenever $a_{t-1}^3 = 0$ we also have $a_t^2 = 0$. a_t^3 then must be 0 by its construction.

Third, $(a_t^3)_{t=1}^\infty$ satisfies constraint (3.1). Since $(a_t^3)_{t=1}^\infty$ features no pause, as we show in Section 3.3, constraint (3.1) boils down to $\delta_A \mathbb{E}[\theta a_t^3] \leq \mathbb{E}[\theta a_{t-1}^3]$. This is satisfied since when defining a_t^3, either $\delta_A \mathbb{E}[\theta a_t^3] = \delta_A \mathbb{E}[\theta a_t^2] \leq \mathbb{E}[\theta a_{t-1}^3]$ or $\delta_A \mathbb{E}[\theta a_t^3] = \mathbb{E}[\theta a_{t-1}^3]$.

Finally, we have $U_t^3 \geq U_t^1\, \forall t$. Since $(a_t^3)_{t=1}^\infty$ has no pause, $U_t^3 = \mathbb{E}[\theta a_t^3]$. The result can be shown by induction. For $t = 1$, $a_1^3 = a_1^2$ by construction and thus $U_1^3 = U_1^2 \geq U_1^1$. Assume $U_{t-1}^3 \geq U_{t-1}^1$ and consider cohort t. We have two possible cases:

1. $\delta_A \mathbb{E}[\theta a_t^2] \leq \mathbb{E}[\theta a_{t-1}^3]$. In this case, $a_t^3 = a_t^2$ and thus we again have $U_1^3 = U_1^2 \geq U_1^1$.
2. $\delta_A \mathbb{E}[\theta a_t^2] > \mathbb{E}[\theta a_{t-1}^3]$. In this case, we have $U_t^3 = \frac{1}{\delta_A} U_{t-1}^3 \geq \frac{1}{\delta_A} U_{t-1}^1 \geq U_t^1$, where the second inequality holds because $(a_t^1)_{t=1}^\infty$ satisfies constraint (3.1).[15]

These together imply that $(a_t^3)_{t=1}^\infty$ is indeed a no-pause policy that satisfies constraint (3.1) and performs weakly better than $(a_t^1)_{t=1}^\infty$ for every cohort.

A.2 Proof for Proposition 2

We note that in this proof, we will refer to some later formulation and result characterizing the optimal no-pause policy. In particular, we will refer to optimization (3.5) – (3.6) in Section 3.3 and Proposition 4 in Section 4.1. This is legitimate because their derivations do not rely on Proposition 2.

Consider the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ in Appendix A.1.1, on which general policies were defined. Suppose general policy $(a_t^1)_{t=1}^\infty$ is optimal and features pauses. Given $(a_t^1)_{t=1}^\infty$, define $(a_t^2)_{t=1}^\infty$ and $(a_t^3)_{t=1}^\infty$ as in Appendix A.1.2. Then $(a_t^2)_{t=1}^\infty$ and $(a_t^3)_{t=1}^\infty$ are no-pause policies, and $(a_t^3)_{t=1}^\infty$ is also optimal. Let U_t^i denote cohort t's expected utility under $(a_t^i)_{t=1}^\infty$. Also define

$$p_t^i = \sum_{k=0}^{t-1} \mathbb{P}(\theta = H | s_0, ... s_k) \mathbb{1}\{m(a_{<t}^i) = k\},\ \forall t$$

where recall $m(a_{<t}^i) := \max\left\{0, \max\{m < t : a_m^i = 1\}\right\}$ records how many cohorts have adopted before t under $(a_t^i)_{t=1}^\infty$. Then $(p_t^i)_{t=1}^\infty$ is the designer's belief process under $(a_t^i)_{t=1}^\infty$.

We first argue that $(a_t^3)_{t=1}^\infty$ must coincide with $(a_t^2)_{t=1}^\infty$ almost surely. Suppose not. Notice by construction $a_t^3 \geq a_t^2$ for all t. We then must have $\mathbb{P}(a_t^2 = 0, a_t^3 = 1) > 0$ for some t. Notice in the event $\{a_t^2 = 0, a_t^3 = 1\}$, adoption has not stopped in period t under $(a_t^3)_{t=1}^\infty$. Our assumptions that $(s_i)_{i=0}^\infty$ are negatively-inconclusive and $(s_i)_{i=1}^\infty$ are upwardly unbounded then imply $\mathbb{P}(a_t^2 = 0, a_t^3 = 1, p_{t+1}^3 > \bar{p}) > 0$.[16] Since $(a_t^2)_{t=1}^\infty$

[15] If $\frac{1}{\delta_A} U_{t-1}^1 < U_t^1$, then under $(a_t^1)_{t=1}^\infty$ cohort $t-1$ would become strictly better off by always waiting in period $t-1$ and then subsequently mimicking cohort t. Constraint (3.1) then must be violated.

[16] Specifically, $\mathbb{P}(p_{t+1}^3 > \bar{p} | a_t^2 = 0, a_t^3 = 1) = \mathbb{E}\Big[\mathbb{P}\Big(p_{t+1}^3 > \bar{p} | (s_i)_{i=0}^{t-1}, (\epsilon_t^1)_{t=1}^\infty, (\epsilon_t^2)_{t=1}^\infty, a_t^3\Big) \Big| a_t^2 = 0, a_t^3 =$

features no pause, this further implies $\mathbb{P}(a_{t+1}^2 = 0, a_t^3 = 1, p_{t+1}^3 > \bar{p}) > 0$. Now, notice the fact that $(a_t^3)_{t=1}^\infty$ is an optimal no-pause policy implies that $p_{t+1}^3 > \bar{p} \overset{a.s.}{\Rightarrow} a_{t+1}^3 = 1$ by Proposition 4 in Section 4.1. We then must have $a_{t+1}^2 = 0, a_t^3 = 1, p_{t+1}^3 > \bar{p} \overset{a.s.}{\Rightarrow} a_{t+1}^2 = 0, a_t^3 = 1, a_{t+1}^3 = 1$. Moreover, by the construction of $(a_t^3)_{t=1}^\infty$, $\{a_{t+1}^2 = 0, a_t^3 = 1, a_{t+1}^3 = 1\}$ can happen only if we have $\delta_A \mathbb{E}[\theta a_{t+1}^2] > \mathbb{E}[\theta a_t^3]$. In this case, the construction implies $a_{t+1}^3 = \mathbb{1}\{a_{t+1}^2 = 1\} + \mathbb{1}\{a_{t+1}^2 = 0, a_t^3 = 1, \epsilon_{t+1}^2 \leq y_{t+1}\}$, where $y_{t+1} \in (0, 1]$ is chosen such that $\delta_A \mathbb{E}[\theta a_{t+1}^3] = \mathbb{E}[\theta a_t^3]$. We thus have:

$$a_{t+1}^2 = 0, a_t^3 = 1, p_{t+1}^3 > \bar{p} \overset{a.s.}{\Rightarrow} a_{t+1}^2 = 0, a_t^3 = 1, a_{t+1}^3 = 1 \overset{a.s.}{\Rightarrow} \epsilon_{t+1}^2 \leq y_{t+1}$$

Notice ϵ_{t+1}^2 is independent from $(a_{t+1}^2, a_t^3, p_{t+1}^3)$ in our construction. Thus the only way to make $a_{t+1}^2 = 0, a_t^3 = 1, p_{t+1}^3 > \bar{p} \overset{a.s.}{\Rightarrow} \epsilon_{t+1}^2 \leq y_{t+1}$ is to have $y_{t+1} = 1$. By construction, this then implies $a_{t+1}^3 = 1$ whenever $a_t^3 = 1$, and thus $a_{t+1}^3 = a_t^3$ (recall $(a_t^3)_{t=1}^\infty$ has no pause). This further implies $\mathbb{E}[\theta a_{t+1}^3] = \mathbb{E}[\theta a_t^3]$. Meanwhile, we also have $\delta_A \mathbb{E}[\theta a_{t+1}^3] = \mathbb{E}[\theta a_t^3]$ by the choice of y_{t+1}. Since $\delta_A < 1$, these together imply $U_t^3 = \mathbb{E}[\theta a_t^3] = 0$. Since cohort t can at least mimic cohort 1's choice, $U_t^3 = 0$ must imply $U_1^3 = 0$. The IC constraint for no-pause policies we specify in (3.6) then implies $U_{t'}^3 = 0$ for all t'. The designer's value is thus zero under $(a_t^3)_{t=1}^\infty$, which contradicts with its optimality since the designer can always gain strictly positive value under Assumption 1(i). Therefore, we must have $(a_t^3)_{t=1}^\infty = (a_t^2)_{t=1}^\infty$ almost surely.

The above analysis implies $(a_t^2)_{t=1}^\infty$ is also an optimal policy. Recall that our proof for Proposition 1 has shown (see equation (A.2)):

$$\begin{aligned} U_t^2 &= \sum_{t' \geq t} \mathbb{E}[\theta \mathbb{1}\{a_t^1 = \ldots = a_{t'-1}^1 = 0,\ a_{t'}^1 = 1\}] \\ &\geq \sum_{t' \geq t} \delta_A^{t'-t} \mathbb{E}[\theta \mathbb{1}\{a_t^1 = \ldots = a_{t'-1}^1 = 0,\ a_{t'}^1 = 1\}] = U_t^1,\ \forall t \end{aligned}$$

For $(a_t^1)_{t=1}^\infty$ to be optimal, we then must have the inequality to hold as equality for all t. Since $\delta_A < 1$, we then must have $\mathbb{E}[\theta \mathbb{1}\{a_t^1 = \ldots = a_{t'-1}^1 = 0,\ a_{t'}^1 = 1\}] = 0$ for all $t' > t$ for all t. This implies $U_t^2 = \mathbb{E}[\theta a_t^1] = U_t^1$ for all t.

Notice $a_t^1 = \mathbb{1}\{a_i^1 = 1, \forall i \leq t\} + \sum_{k=2}^t \mathbb{1}\{a_{k-1}^1 = 0, a_k^1 = \ldots = a_t^1 = 1\}$. For every $k \in [2, t]$, we argue $\mathbb{E}[\theta \mathbb{1}\{a_{k-1}^1 = 0, a_k^1 = \ldots = a_t^1 = 1\}] = 0$. It cannot be < 0 because that will on average violate cohort t's IC constraint in event $\{a_{k-1}^1 = 0, a_k^1 = \ldots = a_t^1 = 1\}$; it cannot be > 0 because that will allow cohort k to obtain strictly positive expected surplus (possibly by waiting in period k) in the event $\{a_{k-1}^1 = 0, a_k^1 = 1\}$, which contradicts

$1\Big] = \mathbb{E}\Big[\mathbb{P}(p_{t+1}^3 > \bar{p} | p_t^3, a_t^3)\Big| a_t^2 = 0, a_t^3 = 1\Big] = \mathbb{E}\Big[G\big((\bar{p}, 1] | p_t^3\big)\Big| a_t^2 = 0, a_t^3 = 1\Big] > 0$. The first equality holds because $a_t^2 \in \sigma\big((s_i)_{i=0}^{t-1}, (\epsilon_t^1)_{t=1}^\infty, (\epsilon_t^2)_{t=1}^\infty\big)$; the second equality holds because given $a_t^3 = 1$, p_t^3 summarizes $(s_i)_{i=0}^{t-1}$ for information about θ; the third equality holds by the definition of G; the final inequality holds because $G\big((\bar{p}, 1] | p_t^3\big) > 0$ a.s. under my assumption.

with $\mathbb{E}[\theta \mathbb{1}\{a^1_{k-1} = 0, a^1_k = 1\}] = 0$ that previous analysis has implied. We then have $U^2_t = U^1_t = \mathbb{E}[\theta a^1_t] = \mathbb{E}[\theta \mathbb{1}\{a^1_i = 1, \forall i \leq t\}]$ for all t.

Now, define another policy $(a^0_t)^\infty_{t=1}$ to be such that $a^0_t = \mathbb{1}\{a^1_{t'} = 1, \forall t' \leq t\}$. Then $(a^0_t)^\infty_{t=1}$ is a no-pause policy which stops adoption when $(a^1_t)^\infty_{t=1}$ recommends waiting for the first time. One can easily check $(a^0_t)^\infty_{t=1}$ is informationally feasible. The previous analysis then implies $U^2_t = U^1_t = U^0_t$ for all t, and thus $(a^0_t)^\infty_{t=1}$ is also an optimal no-pause policy. Since $(a^1_t)^\infty_{t=1}$ features pauses, there must be some $t_1 < t_2$ such that:

$$\mathbb{P}(\underbrace{a^1_1 = ... = a^1_{t_1-1} = 1, a^1_{t_1} = ... = a^1_{t_2-1} = 0, a^1_{t_2} = 1}_{\text{event } A}) > 0$$

On this event A, by our construction $a^0_{t_1} = 0$, which further almost surely implies $p^0_{t_1} < \bar{p}$ by Proposition 4 since $(a^0_t)^\infty_{t=1}$ is an optimal no-pause policy. Moreover, on this event one can see $p^0_{t_1} = p^1_{t_1} = p^1_{t_2}$. We thus must have $p^1_{t_2} < \bar{p}$ on it. This further implies $\mathbb{E}[\theta|A] = \mathbb{E}[u(p^1_{t_2})|A] < 0$. This violates cohort t_2's IC for following $a^1_{t_2} = 1$ after the history of A – contradicting with the optimality of $(a^1_t)^\infty_{t=1}$!

A.3 Proof for Proposition 3

The proof largely follows the proof of Lemma 1 in Feinberg & Piunovskiy (2000), which draws upon results in Balder (1989). In the proof, we formulate the decision problem in an equivalent but slightly different manner. The difference is that instead of directly requiring policies to have no pause, we introduce another state variable to indicate whether adoption has stopped before and use it to restrict current recommendation. Specifically, consider the following controlled Markov process:

- State $(p_t, e_t) \in X := [0,1] \times \{0,1\}$. Here e_t is a state variable indicating whether adoption has ended before.
- Action $a_t \in A := \{0,1\}$.
- $C : X \to \mathcal{P}(A)$ specifies the feasible actions given the current state. We require $C(p_t, e_t) = \{0,1\}$ if $e_t = 0$ and $C(p_t, e_t) = \{0\}$ if $e_t = 1$. (This guarantees no pause.)
- The transition rule of p_t is as specified in (3.9) – (3.10) regardless of e_t. The transition rule of e_t is: $e_1 = 0$; $e_{t+1} = \mathbb{1}\{e_t = 1 \text{ or } a_t = 0\}\ \forall t$.

Let $\bar{\Phi}$ denote the set of all measurable policies. For any $\phi \in \bar{\Phi}$, let P_ϕ denote the induced probability measure over the trajectory space $H^\infty := (X \times A)^\infty$. Let $\mathcal{D} := \{P_\phi : \phi \in \bar{\Phi}\}$, i.e., the set of all strategic measures. Let $\mathcal{U}_t : \mathcal{D} \to \mathbb{R}$ be such that $\mathcal{U}_t(P) = \int a_t u(p_t) dP$.

Any $P \in \mathcal{D}$ is evaluated by the following criteria

$$J^*(P) = \sum_{t=1}^{\infty} \delta_D^t \mathcal{U}_t(P) \tag{A.5}$$

$$J^t(P) = \mathcal{U}_t(P) - \delta_A \mathcal{U}_{t+1}(P) \quad \forall t = 1, 2, ... \tag{A.6}$$

Let $\mathcal{V} := \{(J^*(P), J^1(P), J^2(P), ...) : P \in \mathcal{D}\}$. We ultimately want to show the compactness of $\mathcal{V}$ in $\mathbb{R}^\infty$ (with product topology). We first show three conditions hold.

- **C1:** $C(\cdot)$ has compact range, is compact-valued and is upper-hemicontinuous.
  The first two claims are obvious. Given them, the upper-hemicontinuity is equivalent to $C(\cdot)$ having closed graph.[17] The graph is indeed closed since $Gr(C) = \{(p, e, a) : a \in C(p, e)\} = ([0, 1] \times \{0\} \times \{0, 1\}) \cup ([0, 1] \times \{1\} \times \{0\})$.
- **C2:** Let $\Gamma(\cdot|p_t, e_t, a_t)$ denote the CDF of (p_{t+1}, e_{t+1}) given (p_t, e_t, a_t). Then $\Gamma(\cdot|p_t, e_t, a_t)$ is weakly continuous in (p_t, e_t, a_t) on $Gr(C)$.
  Since the domain of e_t and a_t are discrete, it suffices to check weak continuity in p_t given any possible combination of e_t and a_t on $Gr(C)$.
  - When $(e_t, a_t) = (0, 1)$, $\Gamma(p, e|p_t, e_t, a_t) = G(p|p_t)\mathbb{1}\{e \geq 0\}$. Pick any $(y_k)_k \to p_t$ and fix (p, e) as a continuous point of $\Gamma(\cdot, \cdot|p_t, 0, 1)$. If $e < 0$, $\Gamma(p, e|y_k, 0, 1) \overset{\forall k}{=} 0 = \Gamma(p, e|p_t, 0, 1)$. If $e \geq 0$, $\Gamma(p, e|p_t, 0, 1) = G(p|p_t)$ and thus p must be a continuous point of $G(\cdot|p_t)$. Since $G(\cdot|p_t)$ reflects Bayesian updating, it is well known that it is weakly continuous in the prior p_t (see, e.g., Lyu (2023) Lemma 1).[18] Thus we have $\Gamma(p, e|y_k, 0, 1) = G(p|y_k) \overset{k\to\infty}{\to} G(p|p_t) = \Gamma(p, e|p_t, 0, 1)$.
  - When $(e_t, a_t) = (0, 0)$ or $(1, 0)$, $\Gamma(p, e|p_t, e_t, a_t) = \mathbb{1}\{p \geq p_t, e \geq 1\}$. For any (p, e) not on the boundary of set $\{(p, e) : p \geq p_t, e \geq 1\}$, we obviously have $\mathbb{1}\{p \geq y_k, e \geq 1\} \overset{(y_k)_k\to p_t}{\to} \mathbb{1}\{p \geq p_t, e \geq 1\}$.
- **C3:** $u(p_t)a_t$ is continuous on $Gr(C)$ for any t.
  This is trivial since $u(p_t)a_t = (Hp_t + L(1 - p_t))a_t$.

Now, we note that condition C1 implies conditions (C1a) and (C1b) in Balder (1989);

[17] See, e.g., Sundaram (1996) Proposition 9.8 and the discussion at the end of the same section.

[18] The proof in Lyu (2023) only deals with non-conclusive signals. For s_i that may have conclusive realizations, we can separate those realizations out. Specifically, assume that a bad (resp., good) conclusive news $s_i = \ell$ ($s_i = h$) realizes with probability π_ℓ (π_h) when $\theta = L$ ($\theta = H$). If no conclusive news realizes, s_i will equal to some never-conclusive signal $\hat{s}_i$. Let $\hat{G}(\cdot|\cdot)$ denote the belief updating kernel when one can only observe $\hat{s}_i$. Then $G(\cdot|p_t)$ satisfies:

$$G(p|p_t) = (1 - p_t)\pi_\ell \mathbb{1}\{p \geq 0\} + p_t\pi_h \mathbb{1}\{p \geq 1\} + (1 - p_t\pi_h - (1 - p_t)\pi_\ell)\hat{G}(p|q(p_t)) \tag{A.7}$$

where $q(p_t) := \frac{p_t(1-\pi_h)}{1-p_t\pi_h-(1-p_t)\pi_\ell}$ is the posterior after updating only based on the fact that no conclusive news realizes. Fix any $p_t \in [0, 1]$ and let p be a continuous point of $G(\cdot|p_t)$. If $1 - p_t\pi_h - (1 - p_t)\pi_\ell > 0$, $G(\cdot|p_t)$ being continuous at p implies $\hat{G}(\cdot|q(p_t))$ being continuous at p and we thus have $\hat{G}(p|q(y_k)) \overset{(y_k)_k\to p_t}{\to} \hat{G}(p|q(p_t))$. (A.7) then implies $G(p|y_k) \overset{(y_k)_k\to p_t}{\to} G(p|p_t)$. If $1 - p_t\pi_h - (1 - p_t)\pi_\ell = 0$, we have $1 - y_k\pi_h - (1 - y_k)\pi_\ell \overset{(y_k)_k\to p_t}{\to} 0 = 1 - p_t\pi_h - (1 - p_t)\pi_\ell$. (A.7) then again implies $G(p|y_k) \overset{(y_k)_k\to p_t}{\to} G(p|p_t)$.

C2 implies (C2) in that paper; C3 implies (C3) in the same paper. Theorems 2.1 and 2.2 in Balder (1989) then imply that there is a topology on H^∞ such that when $\mathcal{D}$ is endowed with the corresponding weak topology, we have: (1) $\mathcal{D}$ is compact; (2) $\mathcal{U}_t$ is continuous on $\mathcal{D}$ for all t.[19] The later further implies that $\big(J^*(\cdot), J^1(\cdot), J^2(\cdot), ...\big) : \mathcal{D} \to \mathbb{R}^\infty$ is continuous,[20] whose image $\mathcal{V}$ then must be compact by the compactness of $\mathcal{D}$.

Notice that given the definition of $\mathcal{V}$, designer's problem (3.7) – (3.8) is just equivalent to $\max\{Proj_1(\mathcal{V} \cap (\mathbb{R} \times \mathbb{R}_+^\infty))\}$, where $Proj_1$ projects the set to its first dimension. Since $\mathcal{V}$ is compact, $\mathcal{V} \cap (\mathbb{R} \times \mathbb{R}_+^\infty)$ is also compact, and thus the projection is also compact. Thus the maximum exists.

A.4 Proof for Theorem 1

For any $\phi \in \Phi^\dagger$, we use U_t^ϕ to denote cohort t's expected utility under ϕ, i.e., $U_t^\phi = \mathbb{E}_\phi[a_t u(p_t)]$. First we show a simple lemma:

Lemma A.1. *If ϕ is an optimal no-pause policy, then $0 < U_1^\phi \leq U_2^\phi \leq U_3^\phi \leq ...$*

Proof. First, suppose $U_1^\phi = 0$. Then we must have $U_t^\phi = 0$ for all t by the IC constraints in (3.8). Then the designer's optimal value must be zero. Under Assumption 1(i), however, the designer can generate strictly positive value by recommending adoption in every period when and only when $u(p_1) > 0$ – contradiction!

Second, suppose $U_t^\phi > U_{t+1}^\phi$. we can modify the recommendation policy in period $t+1$ to be such that $a_{t+1} = 1$ as long as $a_t = 1$. This increases cohort $(t+1)$'s expected utility from U_{t+1}^ϕ to U_t^ϕ. At the same time, (weakly) more information is generated in period $t+1$ and it is possible to keep the original expected utility for all later cohorts unchanged. The original ϕ thus cannot be optimal. *Q.E.D.*

Now we start the main proof. The trickiness is that because we have countably many IC constraints, the standard Lagrangian duality result cannot be directly applied due to the *Slater's conundrum* (see, e.g., Martin et al. (2016)). To circumvent it, we first show only finitely many of those constraints are truly restrictive, and then consider duality on the relaxed problem.

Pick any optimal no-pause policy ϕ^*. Let $\epsilon := \frac{1}{2}(1-\delta_A)U_0^{\phi^*}$. Then by the previous lemma $\epsilon > 0$, and we can show: $U_t^{\phi^*} - \delta_A U_{t+1}^{\phi^*} < \epsilon$ for at most finitely many t. To see

[19] While Balder (1989) deals with more general settings, for our purposes the topology on H^∞ can simply be chosen as the usual product topology.

[20] To see $J^*(\cdot)$ is continuous, pick any net $(P_\alpha)_{\alpha\in\mathcal{A}} \overset{w}{\to} P$ in $\mathcal{D}$. For any $\epsilon > 0$ we can pick T such that $\sum_{t=T+1}^\infty \delta_D^t \max\{H, -L\} < \frac{1}{3}\epsilon$. Then $|J^*(P) - J^*(P_\alpha)| \leq |\sum_{t=1}^T \delta_D^t \mathcal{U}_t(P) - \sum_{t=1}^T \delta_D^t \mathcal{U}_t(P_\alpha)| + \frac{2}{3}\epsilon$, which must be less than ϵ when α exceeds some $\alpha_0 \in \mathcal{A}$ since $\mathcal{U}_t(P_\alpha) \overset{P_\alpha \to P}{\to} \mathcal{U}_t(P)$.

this, notice

$$U_t^{\phi^*} - \delta_A U_{t+1}^{\phi^*} < \epsilon \;\Rightarrow\; U_{t+1}^{\phi^*} > \frac{1}{\delta_A}(U_t^{\phi^*} - \epsilon) \;\Rightarrow\; U_{t+1}^{\phi^*} > \underbrace{\frac{1}{\delta_A}(U_t^{\phi^*} - \frac{1}{2}(1-\delta_A)U_t^{\phi^*})}_{=(\frac{1}{2}+\frac{1}{2\delta_A})U_t^{\phi^*}}$$

where the second "⇒" holds because $U_0^{\phi^*} \le U_t^{\phi^*}$ by Lemma A.1. Since $(U_t^{\phi^*})_t$ are bounded and strictly positive by Lemma A.1, the last inequality can only hold for finitely many t.

Now, let $T := \min\{t : U_{t'}^{\phi^*} - \delta_A U_{t'+1}^{\phi^*} \ge \epsilon,\ \forall t' \ge t\}$. Then we know constraint (3.8) is slack by more than ϵ for all $t \ge T$ under ϕ^*. We can define a relaxed problem by dropping those slack constraints:

$$\max_{\phi \in \Phi^\dagger} \sum_{t=1}^{\infty} \delta_D^t \mathbb{E}_\phi[a_t u(p_t)] \tag{A.8}$$

$$\text{s.t. } \mathbb{E}_\phi[a_t u(p_t)] \ge \delta_A \mathbb{E}_\phi[a_{t+1} u(p_{t+1})],\ \forall t < T \tag{A.9}$$

Then it is easy to see that ϕ^* is also optimal for this relaxed problem.[21] By applying the standard Lagrangian duality theorem (e.g., Theorem 1 in Section 8.6 of Luenberger (1997)) to this relaxed problem (we justify the theorem's applicability below), we then obtain the observation below, where $[\mathbb{R}_+^\infty]_T := \{x \in \mathbb{R}_+^\infty : x_t = 0\, \forall t \ge T\}$ and $\mathcal{L}(\cdot,\cdot)$ is as defined in (3.11) in the main text.

Observation. We have $v^* = \min_{\lambda \in [\mathbb{R}_+^\infty]_T} \sup_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda)$ and there exists $\hat{\lambda} \in [\mathbb{R}_+^\infty]_T$ solving the minimization.

Proof for the observation. Notice when λ is restricted to be within $[\mathbb{R}_+^\infty]_T$, $\mathcal{L}(\phi,\lambda)$ is exactly the Lagrangian function for relaxed optimization (A.8) – (A.9). The observation thus holds as long as the Lagrangian duality theorem applies to the relaxed problem. For its applicability, there are two points worth mentioning.

First, the relaxed problem can indeed be considered as a convex (actually linear) program. This can be seen most clearly via the formulation in Section A.3. Let $\mathcal{D}$ be the set of strategic measures as defined in Section A.3, and let $J^*(\cdot), J^1(\cdot), J^2(\cdot), \ldots$ also be as defined there. Then the relaxed problem is equivalent to: $\max_{P \in \mathcal{D}} J^*(P)$ s.t. $J^t(P) \ge 0,\ \forall t < T$. $J^*(\cdot), J^1(\cdot), J^2(\cdot), \ldots$ are obviously linear and the set $\mathcal{D}$ is well known

[21] Suppose not. There exists another policy $\phi' \in \Phi^\dagger$ yielding higher designer's value and:

$$\mathbb{E}_{\phi'}[a_t u(p_t)] - \delta_A \mathbb{E}_{\phi'}[a_{t+1} u(p_{t+1})] \ge 0,\ \forall t < T$$
$$\mathbb{E}_{\phi'}[a_t u(p_t)] - \delta_A \mathbb{E}_{\phi'}[a_{t+1} u(p_{t+1})] \ge L - H,\ \forall t \ge T$$

(Notice the second inequality is trivial since adoption utility is within $[L, H]$.) Then we can consider another mixed policy ϕ^m which follows ϕ^* with probability $\frac{H-L}{H-L+\epsilon}$ and follows ϕ' with probability $\frac{\epsilon}{H-L+\epsilon}$. Then ϕ^m yields higher designer's value than ϕ^* and satisfies constraint (3.8) for all t. Notice that since both ϕ^* and ϕ' feature no pause, ϕ^m also features no pause. Thus ϕ^m dominates ϕ^* in the designer's problem (3.7) – (3.8), which contradicts with ϕ^*'s optimality.

to be convex (see, e.g., section 5.5 in Dynkin & Yushkevich (1979)). This is thus indeed a linear program.

Second, the Slater's condition holds. To see this, notice that under Assumption 1(i), if adoption is recommended in all periods when and only when $u(p_1) > 0$, then all cohorts have the same strictly positive expected utility. All IC constraints are then slack. *Q.E.D.*

Now, pick $\hat{\lambda} \in \arg\min_{\lambda \in [\mathbb{R}_+^\infty]_T} \sup_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda)$. Then for any $\lambda \in \mathbb{R}_+^\infty$, we have:

$$\sup_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \hat{\lambda}) = v^* = \sum_{t=1}^{\infty} \delta_D^t \mathbb{E}_{\phi^*}[a_t u(p_t)] \leq \mathcal{L}(\phi^*, \lambda) \leq \sup_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda) \tag{A.10}$$

where the first equality is due to the above observation; the second equality is by the optimality of ϕ^*; the first inequality holds because ϕ^* satisfies constraint (3.8) and $\lambda \geq 0$. (A.10) implies that $\hat{\lambda}$ also solves $\min_{\lambda \in \mathbb{R}_+^\infty} \sup_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda)$. We thus have shown the existence of λ^* in the proposition, and indeed $v^* = \min_{\lambda \in [\mathbb{R}_+^\infty]_T} \sup_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda) = \min_{\lambda \in \mathbb{R}_+^\infty} \sup_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda)$.

Now, fix any $\lambda^* \in \arg\min_{\lambda \in \mathbb{R}_+^\infty} \sup_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda)$ and pick any $\phi' \in \Phi^\dagger$ satisfying constraint (3.8). We can show that ϕ' is optimal if and only if: (i) $\phi' \in \arg\max_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda^*)$; (ii) $\lambda_t^* \left[\mathbb{E}_{\phi'}[a_t u(p_t)] - \delta_A \mathbb{E}_{\phi'}[a_{t+1} u(p_{t+1})]\right] = 0$, $\forall t$. To show this, notice we have:

$$v^* = \sup_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda^*) \geq \mathcal{L}(\phi', \lambda^*) \geq \sum_{t=1}^{\infty} \delta_D^t \mathbb{E}_{\phi'}[a_t u(p_t)]$$

where the second inequality holds because $\lambda_t^* \left[\mathbb{E}_{\phi'}[a_t u(p_t)] - \delta_A \mathbb{E}_{\phi'}[a_{t+1} u(p_{t+1})]\right] \geq 0$, $\forall t$. Notice that the first inequality holds as equality if and only if (i) holds; the second inequality holds as equality if and only if (ii) holds. Thus ϕ' is optimal if and only if those two conditions hold.

Finally, recall that as we have argued at the beginning, under any optimal ϕ^* constraint (3.8) must be slack except for finitely many periods. ϕ^* satisfying condition (ii) then implies that λ^* can only have finitely many non-zero entries.

A.5 Proofs in Section 4

As a preparation, we first provide the dynamic programming formulation and results for $\max_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda)$.

A.5.1 Dynamic programming for the Lagrangian optimization

Let $\widehat{\mathbb{R}_+^\infty} := \{x \in \mathbb{R}_+^\infty : \exists t \text{ s.t. } x_{t'} = 0 \, \forall t' \geq t\}$. We first set up the Bellman equation for any $\lambda \in \widehat{\mathbb{R}_+^\infty}$.

Let $C[0,1]$ denote the set of continuous functions on $[0,1]$ endowed with $||\cdot||_{sup}$. For any $V \in C[0,1]$, define $\Upsilon(V)$ as

$$\Upsilon(V)(p) = \max\left\{0, u(p) + \delta_D \int V(p')G(dp'|p)\right\} \forall p \in [0,1] \tag{A.11}$$

Then it is easy to see that Υ is a contraction mapping from $C[0,1]$ to itself and thus has a unique fixed point.[22] We define $V^0(\cdot)$ to be this fixed point, i.e., $V^0 = \Upsilon(V^0)$. Intuitively, notice $V^0 = \Upsilon(V^0)$ is just the Bellman equation for $\max_{\phi\in\Phi^\dagger} \mathcal{L}(\phi, \mathbf{0})$, which is optimization (3.7) without constraint (3.8), i.e., the first-best design problem.

Now, given any $\lambda \in \widehat{\mathbb{R}^\infty_+}$, let $M(\lambda) := \min\{t : \lambda_{t'} = 0 \, \forall t' \geq t\}$. We define value functions $\left(V_t^\lambda(\cdot)\right)_{t=1}^\infty$ as follows:

$$V_t^\lambda(\cdot) = V^0(\cdot), \; \forall t > M(\lambda) \tag{A.12}$$

$$V_t^\lambda(p) = \max\left\{0, (1 + \lambda_t - \frac{\delta_A}{\delta_D}\lambda_{t-1})u(p) + \delta_D \int V_{t+1}^\lambda(p')G(dp'|p)\right\} \forall p, \; \forall t \leq M(\lambda) \tag{A.13}$$

(where recall $\lambda_0 := 0$). Notice that by this definition, equation (A.13) holds even for $t > M(\lambda)$, since for those t, $\lambda_t = \lambda_{t-1} = 0$. Moreover, $V_t^\lambda(\cdot)$ is continuous for all t (see footnote 22).

Define $H^0(p) := u(p) + \delta_D \int V^0(p')G(dp'|p)$ and

$$H_t^\lambda(p) := (1 + \lambda_t - \frac{\delta_A}{\delta_D}\lambda_{t-1})u(p) + \delta_D \int V_{t+1}^\lambda(p')G(dp'|p), \; \forall t \tag{A.14}$$

Then $H_t^\lambda(\cdot)$ is also continuous. To ease notations, define $a_0 := 1$. We have the standard dynamic programming result:

Lemma A.2. *Given any $\lambda \in \widehat{\mathbb{R}^\infty_+}$ and initial belief distribution μ_1, we have:*

(a) No-pause policy $\phi^\lambda \in \arg\max_{\phi\in\Phi^\dagger} \mathcal{L}(\phi, \lambda)$ if and only if: $\mathbb{P}_{\phi^\lambda}\big(a_{t-1} = 1, a_t = 0, H_t^\lambda(p_t) > 0\big) = \mathbb{P}_{\phi^\lambda}\big(a_{t-1} = 1, a_t = 1, H_t^\lambda(p_t) < 0\big) = 0 \, , \forall t = 1, 2, ...$

(b) $\max_{\phi\in\Phi^\dagger} \mathcal{L}(\phi, \lambda) = \int \delta_D V_1^\lambda(p)\mu_1(dp)$

Proof. Fix any $\phi \in \Phi^\dagger$. By the Bellman equation (A.13) (recall it holds even for $t > M(\lambda)$ as we mentioned earlier), we have $\mathbb{P}_\phi$-a.s.:

$$a_{t-1}V_t^\lambda(p_t) \geq (1 + \lambda_t - \frac{\delta_A}{\delta_D}\lambda_{t-1})a_t u(p_t) + \delta_D \int a_t V_{t+1}^\lambda(p')G(dp'|p_t), \; \forall t \geq 1 \tag{A.15}$$

(If $a_{t-1} = 0$, the no-pause feature implies $a_t = 0$ $\mathbb{P}_\phi$-a.s., and the inequality holds as equality trivially; if $a_{t-1} = 1$, the inequality holds for both $a_t = 0$ and $a_t = 1$ by

[22]Given any $V \in C[0,1]$, to show $\Upsilon(V)$ is continuous it suffices to show $\int V(p')G(dp'|p)$ is continuous in p. This is true because $G(\cdot|p)$ is weakly continuous in p as we have mentioned in the proof of Proposition 3 (see discussion before footnote 18).

(A.13).) Notice that when $a_{t-1} = 1$ and $a_t = 0$, the inequality holds as equality if and only if $H_t^\lambda(p_t) \le 0$; when $a_{t-1} = a_t = 1$, the inequality holds as equality if and only if $H_t^\lambda(p_t) \ge 0$. Thus the inequality holds as equality $\mathbb{P}_\phi$-a.s. if and only if: $\mathbb{P}_\phi\big(a_{t-1} = 1, a_t = 0, H_t^\lambda(p_t) > 0\big) = \mathbb{P}_\phi\big(a_{t-1} = 1, a_t = 1, H_t^\lambda(p_t) < 0\big) = 0$.

Referring to the above inequality repeatedly, one can see:

$$\int \delta_D V_1^\lambda(p)\mu_1(dp) = \mathbb{E}_\phi[\delta_D a_0 V_1^\lambda(p_1)] \ge \delta_D \mathbb{E}_\phi\Big[(1+\lambda_1 - \frac{\delta_A}{\delta_D}\lambda_0)a_1 u(p_1)\Big] + \delta_D^2 \mathbb{E}_\phi\big[a_1 V_2^\lambda(p_2)\big]$$
$$\ge \delta_D \mathbb{E}_\phi\Big[(1+\lambda_1 - \frac{\delta_A}{\delta_D}\lambda_0)a_1 u(p_1)\Big] + \delta_D^2 \mathbb{E}_\phi\Big[(1+\lambda_2 - \frac{\delta_A}{\delta_D}\lambda_1)a_2 u(p_2)\Big] + \delta_D^3 \mathbb{E}_\phi\big[a_2 V_3^\lambda(p_3)\big]$$
$$\ge \dots \ge \sum_{t=1}^{T} \delta_D^t \mathbb{E}_\phi\Big[(1+\lambda_t - \frac{\delta_A}{\delta_D}\lambda_{t-1})a_t u(p_t)\Big] + \delta_D^{T+1}\mathbb{E}_\phi\big[a_T V_{T+1}^\lambda(p_{T+1})\big]$$
$$\xrightarrow{T\to\infty} \sum_{t=1}^{\infty} \delta_D^t \mathbb{E}_\phi\Big[(1+\lambda_t - \frac{\delta_A}{\delta_D}\lambda_{t-1})a_t u(p_t)\Big] = \mathcal{L}(\phi, \lambda)$$

(Here, $\delta_D^{T+1}\mathbb{E}_\phi\big[a_T V_{T+1}^\lambda(p_{T+1})\big] \xrightarrow{T\to\infty} 0$ because when T is large enough $V_{T+1}^\lambda = V^0$, which is bounded.) This implies that $\int \delta_D V_1^\lambda(p)\mu_1(dp)$ is an upper bound for $\mathcal{L}(\phi, \lambda)$. This upper bound is achieved if and only if all inequalities involved above hold as equalities. This requires inequality (A.15) to hold as equality $\mathbb{P}_\phi$-a.s. for all t. Based on our previous discussion, this requires $\mathbb{P}_\phi\big(a_{t-1} = 1, a_t = 0, H_t^\lambda(p_t) > 0\big) = \mathbb{P}_\phi\big(a_{t-1} = 1, a_t = 1, H_t^\lambda(p_t) < 0\big) = 0$, $\forall t$.

Now it suffices to check there is indeed a measurable no-pause policy satisfying the condition. Such a no-pause policy can be constructed as: $\phi_t(p_{\le t}, a_{<t}) = \mathbb{1}\{a_{t-1} = 1, H_t^\lambda(p_t) \ge 0\}$, $\forall t$. It is measurable since $H_t^\lambda(\cdot)$ is continuous. *Q.E.D.*

A.5.2 Proof for Lemma 1

Pick any $\lambda^* \in \arg\min_{\lambda \in \mathbb{R}_+^\infty} \sup_{\phi \in \Phi^\dagger} \mathcal{L}(\phi, \lambda)$ and let ϕ^* be an optimal no-pause policy. Suppose $1 + \lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^* < 0$. Then $t > 1$, and $(1 + \lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^*)u(p_t) > 0$ whenever $u(p_t) < 0$. This implies $H_t^{\lambda^*}(p_t) > 0$ whenever $u(p_t) < 0$. By Theorem 1 and Lemma A.13, we then must have $a_{t-1} = 1, u(p_t) < 0 \Rightarrow a_t = 1$ $\mathbb{P}_{\phi^*}$-a.s. This implies that when $a_{t-1} = 1$, no matter whether $u(p_t) < 0$ or not, we have $a_t u(p_t) \le u(p_t)$ $\mathbb{P}_{\phi^*}$-a.s. We then have:

$$\mathbb{E}_{\phi^*}[a_t u(p_t)] = \mathbb{E}_{\phi^*}[a_{t-1} a_t u(p_t)] \le \mathbb{E}_{\phi^*}[a_{t-1} u(p_t)] = \mathbb{E}_{\phi^*}[a_{t-1} u(p_{t-1})]$$

where the last equality holds because $\mathbb{E}[u(p_t)|p_{t-1}, a_{t-1}] = u(p_{t-1})$. Notice Lemma A.1 implies $\mathbb{E}_{\phi^*}[a_{t-1}u(p_{t-1})] > 0$. The above result then implies $\mathbb{E}_{\phi^*}[a_{t-1}u(p_{t-1})] - \delta_A \mathbb{E}_{\phi^*}[a_t u(p_t)] > 0$, which implies $\lambda_{t-1}^* = 0$ by condition (ii) in Theorem 1. This contradicts with $1 + \lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^* < 0$.

A.5.3 Proof for Proposition 4

Pick any $\lambda^* \in \arg\min_{\lambda\in\mathbb{R}_+^\infty} \sup_{\phi\in\Phi^\dagger} \mathcal{L}(\phi,\lambda)$. By the dynamic programming result (Lemma A.2), it suffices to show $p \geq \bar{p} \Rightarrow H_t^{\lambda^*}(p) > 0, \forall t$.

Fix $t \geq 1$ and $p \geq \bar{p}$. Recall that $M(\lambda^*) := \min\{t : \lambda^*_{t'} = 0 \,\forall t' \geq t\}$. Define $T := \max\{t, M(\lambda^*)\}$. We have:

$$
\begin{aligned}
H_t^{\lambda^*}(p) =& (1+\lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^*)u(p) + \delta_D \int V_{t+1}^{\lambda^*}(p')G(dp'|p) \\
\geq& (1+\lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^*)u(p) + \delta_D \int (1+\lambda_{t+1}^* - \frac{\delta_A}{\delta_D}\lambda_t^*)u(p')G(dp'|p) \\
&+ \delta_D \int \delta_D \int V_{t+2}^{\lambda^*}(p'')G(dp''|p')G(dp'|p) \\
\geq \dots \geq& \sum_{k=t}^{T} \delta_D^{k-t} \int (1+\lambda_k^* - \frac{\delta_A}{\delta_D}\lambda_{k-1}^*)u(p')G^{k-t}(dp'|p) \\
&+ \delta_D^{T+1-t} \int V_{T+1}^{\lambda^*}(p')G^{T+1-t}(dp'|p) \qquad \text{(A.16)}
\end{aligned}
$$

(where recall $G^K(\cdot|\cdot)$ reflects Bayesian updating following signals of K cohorts).[23] The inequalities holds because by the Bellman equation (A.13), $V_t^{\lambda^*}(p) \geq (1+\lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^*)u(p) + \delta_D \int V_{t+1}^{\lambda^*}(p')G(dp'|p), \forall t$. Since $p \geq \bar{p}$, by the law of iterated expectation (LIE) we have $\int u(p')G^{k-t}(dp'|p) \geq 0$. The fact that $(1+\lambda_k^* - \frac{\delta_A}{\delta_D}\lambda_{k-1}^*) \geq 0$ (by Lemma 1) then implies the first term in expression (A.16) is non-negative. It thus suffices to show $\int V_{T+1}^{\lambda^*}(p')G^{T+1-t}(dp'|p) > 0$. To see this, notice $\int V_{T+1}^{\lambda^*}(p')G^{T+1-t}(dp'|p)$ equals to

$$
\int \max\left\{0, (1+\lambda_{T+1}^* - \frac{\delta_A}{\delta_D}\lambda_T^*)u(p') + \delta_D \int V_{T+2}^{\lambda^*}(p'')G(dp''|p')\right\} G^{T+1-t}(dp'|p) \qquad \text{(A.17)}
$$

Since $T \geq M(\lambda)$, $\lambda_T^* = \lambda_{T+1}^* = 0$. Also notice $V_{T+2}^{\lambda^*}(\cdot) \geq 0$. The integrand in (A.17) is thus > 0 for any $p' > \bar{p}$. Moreover, Assumption 1(ii) and LIE imply that we must have $G^{T+1-t}((\bar{p},1]|p) > 0$ for any $p \geq \bar{p}$. Thus the integration in (A.17) is indeed > 0.

A.5.4 Proof for Proposition 5

The key are the following properties about $V_t^{\lambda^*}(\cdot)$ and $H_t^{\lambda^*}(\cdot)$:

Lemma A.3. *Fix $\lambda^* \in \arg\min_{\lambda\in\mathbb{R}_+^\infty} \sup_{\phi\in\Phi^\dagger} \mathcal{L}(\phi,\lambda)$. For any t, $V_t^{\lambda^*}(\cdot)$ and $H_t^{\lambda^*}(\cdot)$ are weakly increasing, and $H_t^{\lambda^*}(\cdot)$ has a root in $[0,1]$. Moreover, when $(s_i)_{i=1}^\infty$ are upwardly unbounded, $H_t^{\lambda^*}(\cdot)$'s root is unique.*

Proof. Let $C^I[0,1]$ denote the set of weakly increasing functions in $C[0,1]$. Since $G(\cdot|\cdot)$ reflects the belief transition following Bayesian updating, it is well known that $G(\cdot|p)$

[23] Formally, $G^0(\cdot|x_0) = D(\cdot|x_0)$; $G^K(A|x_0) = \int \dots \int \mathbb{1}\{x_K \in A\}G(dx_K|x_{K-1})\dots G(dx_1|x_0)$.

increases in first-order stochastic dominance in p. Υ defined in (A.11) thus maps $C^I[0,1]$ into itself. Since $C^I[0,1]$ is closed in $C[0,1]$, Υ's fixed point $V^0 \in C^I[0,1]$. Thus $V_t^{\lambda^*} \in C^I[0,1]$ for all $t > M(\lambda^*)$. Moreover, since $1+\lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^* \geq 0$ by Lemma 1, the functional mapping $V_{t+1}^{\lambda} \mapsto V_t^{\lambda}$ defined in (A.13) also maps $C^I[0,1]$ into itself with $\lambda = \lambda^*$. Thus $V_t^{\lambda^*} \in C^I[0,1]$ also for all $t \leq M(\lambda^*)$. Given $1+\lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^* \geq 0$ and $V_{t+1}^{\lambda^*}(\cdot)$ being weakly increasing, $H_t^{\lambda^*}(\cdot)$ is also weakly increasing by its definition (see equation (A.14)).

Since $V_t^{\lambda^*}(\cdot)$ reflects the continuation value of the Lagrangian maximization problem, it is easy to see $V_t^{\lambda^*}(0) = 0$ and $V_t^{\lambda^*}(1) \geq 0$. Also notice $u(0) < 0$ and $u(1) > 0$. The definition of $H_t^{\lambda^*}(\cdot)$ then implies $H_t^{\lambda^*}(0) \leq 0$ and $H_t^{\lambda^*}(1) \geq 0$. The continuity of $H_t^{\lambda^*}(\cdot)$ then implies it has a root.

Now, assume $(s_i)_{i=1}^{\infty}$ are upwardly unbounded. Suppose there exist p^a, p^b s.t. $p^a < p^b$ but $H_t^{\lambda^*}(p^a) = H_t^{\lambda^*}(p^b) = 0$. Then by equation (A.14), for $p = p^a, p^b$, we have:

$$(1 + \lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^*)u(p) + \delta_D \int V_{t+1}^{\lambda^*}(p')G(dp'|p) = 0 \tag{A.18}$$

Since $\int V_{t+1}^{\lambda^*}(p')G(dp'|p^b) \geq \int V_{t+1}^{\lambda^*}(p')G(dp'|p^a)$, we must have $(1 + \lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^*)u(p^b) \leq (1+\lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^*)u(p^a)$, which can hold only when $1+\lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^* = 0$. Equation (A.18) then implies $\int V_{t+1}^{\lambda^*}(p')G(dp'|p^b) = 0$, which by the Bellman equation (A.13) further implies:

$$\int \max\{0, H_{t+1}^{\lambda^*}(p')\}G(dp'|p^b) = 0$$

This implies that we must have $H_{t+1}^{\lambda^*}(p') \leq 0$ for $G(\cdot|p^b)$-a.e. p'. However, $(s_i)_{i=1}^{\infty}$ being upwardly unbounded implies $G((\bar{p}, 1]|p^b) > 0$, while for any $p' \in (\bar{p}, 1]$ we must have $H_{t+1}^{\lambda^*}(p') > 0$ according to our proof for Proposition 4 (see Section A.5.3)– contradiction! Thus $H_t^{\lambda^*}(\cdot)$'s root is unique. *Q.E.D.*

Now we go to the main proof for Proposition 5. Given λ^* solving the dual problem, for all t, pick η_t to be a root of $H_t^{\lambda^*}(\cdot)$. Let ϕ^* be a threshold no-pause policy with thresholds $(\eta_t)_{t=1}^{\infty}$. Since $H_t^{\lambda^*}(\cdot)$ is weakly increasing, we then have $\mathbb{P}_{\phi^*}\big(a_{t-1} = 1, a_t = 0, H_t^{\lambda^*}(p_t) > 0\big) = \mathbb{P}_{\phi^*}\big(a_{t-1} = 1, a_t = 1, H_t^{\lambda^*}(p_t) < 0\big) = 0\,, \forall t$. Lemma A.2 then implies ϕ^* is optimal.

When $(s_i)_{i=1}^{\infty}$ are upwardly unbounded, $H_t^{\lambda^*}(\cdot)$ has a unique root. Then, the choice of $(\eta_t)_{t=1}^{\infty}$ above is unique, and $H_t^{\lambda^*}(p_t) > (<) 0 \iff p_t > (<) \eta_t$. Thus any no-pause policy ϕ satisfying $\mathbb{P}_{\phi}\big(a_{t-1} = 1, a_t = 0, H_t^{\lambda^*}(p_t) > 0\big) = \mathbb{P}_{\phi}\big(a_{t-1} = 1, a_t = 1, H_t^{\lambda^*}(p_t) < 0\big) = 0\,, \forall t$ must be a threshold policy with $(\eta_t)_{t=1}^{\infty}$ being the thresholds.

A.5.5 Proof for Proposition 6

Part (a): Fix λ^* solving the dual problem. Let ϕ^* denote the optimal threshold policy with thresholds $(\eta_t^*)_t$ and let ϕ^F denote the first-best policy. By the dynamic programming result and the continuity of $H_1^{\lambda^*}$, we must have $H_1^{\lambda^*}(\eta_1^*) \geq 0$. Now, suppose $\eta_1^* \leq \eta^F$.

Given any distribution F_{p_1} of p_1, we use $\mathbb{P}_{F_{p_1},\phi}$ to denote the distribution over $(p_t, a_t)_{t=1}^{\infty}$ induced by policy ϕ, and use $\mathbb{E}_{F_{p_1},\phi}$ to denote the corresponding expectation operator. If F_{p_1} supports on a singleton $\{x\}$, we also use notations $\mathbb{P}_{x,\phi}$ and $\mathbb{E}_{x,\phi}$ instead.

Let $\phi^\dagger$ denote the policy that chooses $a_1 = 1$ for sure and then subsequently follows the thresholds of ϕ^*. By the construction of H^λ and the optimality of ϕ^*, one can see:

$$H_1^{\lambda^*}(\eta_1^*) = \sum_{t=1}^{\infty} \delta_D^t \mathbb{E}_{\eta_1^*,\phi^\dagger}[a_t u(p_t)] + \sum_{t=1}^{\infty} \delta_D^t \lambda_t^* \mathbb{E}_{\eta_1^*,\phi^\dagger}\left[a_t u(p_t) - \delta_A a_{t+1} u(p_{t+1})\right]$$

Since the first-best is not achievable, η_t^* must sometimes deviate from η^F. Thus the first term $\sum_{t=1}^{\infty} \delta_D^t \mathbb{E}_{\eta_1^*,\phi^\dagger}[a_t u(p_t)] < \sum_{t=1}^{\infty} \delta_D^t \mathbb{E}_{\eta_1^*,\phi^F}[a_t u(p_t)] = 0.$[24] $H_1^{\lambda^*}(\eta_1^*) \geq 0$ then implies the second term above must be > 0. This in particular implies there exists some t such that $\lambda_t^* \mathbb{E}_{\eta_1^*,\phi^\dagger}\left[a_t u(p_t) - \delta_A a_{t+1} u(p_{t+1})\right] > 0$. Notice by condition (ii) in Theorem 1, we have $\lambda_t^* \mathbb{E}_{\mu_1,\phi^*}\left[a_t u(p_t) - \delta_A a_{t+1} u(p_{t+1})\right] = 0$, which is equivalent to $\lambda_t^* \int_{x \geq \eta_1^*} \mathbb{E}_{x,\phi^\dagger}\left[a_t u(p_t) - \delta_A a_{t+1} u(p_{t+1})\right] \mu_1(dx) = 0$. If we can show $\mathbb{E}_{x,\phi^\dagger}\left[a_t u(p_t) - \delta_A a_{t+1} u(p_{t+1})\right]$ single-crosses zero from below in x, then this will contradict with $\lambda_t^* \mathbb{E}_{\eta_1^*,\phi^\dagger}\left[a_t u(p_t) - \delta_A a_{t+1} u(p_{t+1})\right] > 0$. It thus suffices for us to show that single-crossing property in the rest of the proof.

Notice:

$$\begin{aligned}
\mathbb{E}_{x,\phi^\dagger}\left[a_t u(p_t) - \delta_A a_{t+1} u(p_{t+1})\right] &= \mathbb{E}_{x,\phi^\dagger}\left[a_t u(p_{t+1}) - \delta_A a_{t+1} u(p_{t+1})\right] \\
&= \mathbb{E}_{x,\phi^\dagger}\left[a_t (1 - \delta_A \mathbb{1}\{p_{t+1} \geq \eta_{t+1}^*\}) u(p_{t+1})\right] \\
&= \mathbb{P}_{x,\phi^\dagger}(a_t = 1) \mathbb{E}_{x,\phi^\dagger}\left[(1 - \delta_A \mathbb{1}\{p_{t+1} \geq \eta_{t+1}^*\}) u(p_{t+1}) \middle| a_t = 1\right]
\end{aligned}$$

where the second equality holds by the definition of $\phi^\dagger$. It then suffices to show $\mathbb{E}_{x,\phi^\dagger}\left[(1 - \delta_A \mathbb{1}\{p_{t+1} \geq \eta_{t+1}^*\}) u(p_{t+1}) \middle| a_t = 1\right]$ is (weakly) increasing in x for all t.

For any $t \geq 2$, let $f_t^x := f_{p_t | p_{t-1} \geq \eta_{t-1}^*, \ldots, p_2 \geq \eta_2^*, p_1 = x}$ denote the density of p_t under $\mathbb{P}_{x,\phi^\dagger}$ conditional on $p_{t-1} \geq \eta_{t-1}^*, \ldots, p_2 \geq \eta_2^*$. Since $a_t = 1 \overset{a.s.}{\Longleftrightarrow} p_{t-1} \geq \eta_{t-1}^*, \ldots, p_2 \geq \eta_2^*$ under $\phi^\dagger$, $\mathbb{E}_{x,\phi^\dagger}\left[(1 - \delta_A \mathbb{1}\{p_{t+1} \geq \eta_{t+1}^*\}) u(p_{t+1}) \middle| a_t = 1\right] = \int_p (1 - \delta_A \mathbb{1}\{p \geq \eta_{t+1}^*\}) u(p) f_{t+1}^x(p) dp$. Notice that with $\eta_{t+1}^* \leq \bar{p}$ (implied by Proposition 4), term $(1 - \delta_A \mathbb{1}\{p \geq \eta_{t+1}^*\}) u(p)$ is increasing in p. It thus suffices for us to show f_t^x is increasing in x in the likelihood ratio order ($\succeq_{LR}$) for all $t \geq 2$. To show this, it is more convenient to work with the log-likelihood ratio transformation of p_t. Formally, for any $y \in (0,1)$, let $lr(y) := \log(\frac{y}{1-y})$. Define $\ell_t := lr(p_t)$ and let $\hat{f}_t^x := \hat{f}_{\ell_t | \ell_{t-1} \geq lr(\eta_{t-1}^*), \ldots, \ell_2 \geq lr(\eta_2^*), \ell_1 = lr(x)}$ denote the density of ℓ_t under $\mathbb{P}_{x,\phi^\dagger}$ conditional on $\ell_{t-1} \geq lr(\eta_{t-1}^*), \ldots, \ell_2 \geq lr(\eta_2^*)$. It then suffices to show $\hat{f}_t^x$ is increasing in x in $\succeq_{LR}$ for all $t \geq 2$.

For $t \geq 1$, let $\xi_t := \log(\frac{f_{s_t|H}(s_t)}{f_{s_t|L}(s_t)})$, and let j_θ denote its density given θ. It is then well

[24] Notice the first-best threshold does not depend on the initial distribution of p_1. We have the inequality strict even when $\eta_1^* = \eta^F$ because under our assumption about the signals, p_t ($t > 1$) given adoption having not stopped always has full support, and thus any deviation from η^F will strictly decrease the designer's value. The designer's first-best value given $p_1 = \eta_1^*$ equals to 0 since $\eta_1^* \leq \eta^F$.

known that $j_H(y) = e^y j_L(y)$.[25] Also, the Bayes formula implies $\ell_{t+1} = \ell_t + \xi_t$ if $a_t = 1$. We now show $\hat{f}_t^x$ is increasing in x in $\succeq_{LR}$ for all $t \geq 2$ by induction.

- For $t = 2$, notice under $\phi^\dagger$, $\ell_2 = y$ if and only if $\xi_1 = y - \ell_1$ given ℓ_1. By construction we then have

$$\begin{aligned}\hat{f}_2^x(y) &= \frac{1}{1+e^{\ell_1}} j_L(y-\ell_1) + \frac{e^{\ell_1}}{1+e^{\ell_1}} j_H(y-\ell_1)\Big|_{\ell_1 = lr(x)} \\ &= \frac{1+e^y}{1+e^{\ell_1}} j_L(y-\ell_1)\Big|_{\ell_1=lr(x)}\end{aligned}$$

  where for the first equality notice $\frac{1}{1+e^{\ell_1}}$ just equals to the probability of $\theta = L$ given ℓ_1, and for the second equality we have used $j_H(y-\ell_1) = e^{y-\ell_1} j_L(y-\ell_1)$. By our assumption, $j_L(\cdot)$ is log-concave, which implies $j_L(y-\ell_1)$ to be log-super-modular (log-SPM) in (y, ℓ_1). Moreover, term $\frac{1+e^y}{1+e^{\ell_1}}$ is obviously log-modular. Thus the expression above is log-SPM in (y, ℓ_1). Since $lr(\cdot)$ is increasing, this implies $\hat{f}_2^x(y)$ to be log-SPM in (y, x), which further implies $\hat{f}_2^x$ to be increasing in x in $\succeq_{LR}$.
- Given the result holding for $\hat{f}_t^x$, consider $\hat{f}_{t+1}^x$. Notice under $\phi^\dagger$ when $\ell_t \geq lr(\eta_t^*), ..., \ell_2 \geq lr(\eta_2^*)$, we have $\ell_{t+1} = y$ if and only if $\xi_t = y - \ell_t$. Thus by construction we have

$$\begin{aligned}\hat{f}_{t+1}^x(y) &= \frac{\int_{\ell_t} \left[\frac{1}{1+e^{\ell_t}} j_L(y-\ell_t) + \frac{e^{\ell_t}}{1+e^{\ell_t}} j_H(y-\ell_t)\right] \mathbb{1}\{\ell_t \geq lr(\eta_t^*)\} \hat{f}_t^x(\ell_t) d\ell_t}{\int_{\ell_t \geq lr(\eta_t^*)} \hat{f}_t^x(\ell_t) d\ell_t} \\ &= \frac{\int_{\ell_t} \frac{1+e^y}{1+e^{\ell_t}} j_L(y-\ell_t) \mathbb{1}\{\ell_t \geq lr(\eta_t^*)\} \hat{f}_t^x(\ell_t) d\ell_t}{\int_{\ell_t \geq lr(\eta_t^*)} \hat{f}_t^x(\ell_t) d\ell_t}\end{aligned}$$

  Notice under our assumption and induction condition, terms $\frac{1+e^y}{1+e^{\ell_t}}$, $j_L(y-\ell_t)$, $\mathbb{1}\{\ell_t \geq lr(\eta_t^*)\}$ and $\hat{f}_t^x(\ell_t)$ are all log-SPM in (y, ℓ_t, x), and thus so is their product. The numerator above is thus log-SPM in (y, x) since this property is preserved after partial integration.[26] Since the denominator only depends on x, this implies $\hat{f}_{t+1}^x(y)$ to be log-SPM in (y, x), which is equivalent to $\hat{f}_{t+1}^x$ being increasing in x in $\succeq_{LR}$.

Part (b): Still let λ^* solve the dual problem. Define $\bar{t} = \min\{t : \lambda_{t'}^* = 0 \, \forall t' \geq t\}$. We have $\bar{t} < \infty$ by Theorem 1 and $\bar{t} > 1$ because the first-best is not feasible. Since $\lambda_t^* = 0$ for all $t \geq \bar{t}$, $V_t^{\lambda^*} \equiv V^0$ for all $t > \bar{t}$ by the Bellman equation (A.12). Thus $\eta_t^* = \eta^F$ for all $t > \bar{t}$. Moreover, by the definition of H^λ we have:

$$\begin{aligned}H_{\bar{t}}^{\lambda^*}(\eta^F) &= (1 - \frac{\delta_A}{\delta_D} \lambda_{\bar{t}-1}^*) u(\eta^F) + \delta_D \int V^0(p') G(dp'|\eta^F) \\ &< u(\eta^F) + \delta_D \int V^0(p') G(dp'|\eta^F) = 0\end{aligned}$$

where the strict inequality holds because $\lambda_{\bar{t}-1}^* > 0$ and $u(\eta^F) < 0$. The last expression

[25]For a proof, see, e.g., Claim (a) in the proof of Lemma 1 in Lyu (2023).
[26]See Karlin & Rinott (1980), where log-SPM is called MTP_2.

equals to zero since η^F is the first-best threshold. By the dynamic programming result (Lemma A.2), $H_t^{\lambda^*}(\eta^F) < 0$ then implies $\eta_t^* > \eta^F$.

A.6 Proof for Proposition 7

Consider learning with conclusive good news. Recall that in Section A.5.1 we have defined $V^0(\cdot)$, $H^0(\cdot)$, $(V_t^\lambda(\cdot))_{t=1}^\infty$ and $(H_t^\lambda(\cdot))_{t=1}^\infty$. Intuitively, $(V_t^\lambda(\cdot))_{t=1}^\infty$ are the value functions for $\max_{\phi\in\Phi^\dagger}\mathcal{L}(\phi,\lambda)$, and $V^0(\cdot)$ is the value function for $\max_{\phi\in\Phi^\dagger}\mathcal{L}(\phi,\mathbf{0})$, which is the same as the first-best problem. Assume the first-best policy is not incentive compatible. Fix ϕ^* as an optimal no-pause policy. We first prove several lemmas below.

Lemma A.4. *(a) For any $t \le t^F - 1$, if $\alpha_t^{\phi^*} < 1$, then IC_t is binding under ϕ^*; (b) For any $t \ge t^F$, if $\alpha_t^{\phi^*} > 0$, then IC_{t-1} is binding under ϕ^*.*

Proof. Fix λ^* solving the dual problem (3.13). We prove the two parts in turn.
Part (a): First, we argue that

Observation (i). $\lambda_{t-1}^* = 0 \Rightarrow V_t^{\lambda^*}(1) \ge V^0(1)$.

Subproof. Since $1 + \lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^* \ge 0$ by Lemma 1 and $u(1) > 0$, the Bellman equation (A.13) implies $V_t^{\lambda^*}(1) = (1 + \lambda_t^* - \frac{\delta_A}{\delta_D}\lambda_{t-1}^*)u(1) + \delta_D V_{t+1}^{\lambda^*}(1)$, $\forall t$, which further implies $V_t^{\lambda^*}(1) = \sum_{t'=t}^\infty \delta_D^{t'-t}(1 + \lambda_{t'}^* - \frac{\delta_A}{\delta_D}\lambda_{t'-1}^*)u(1)$. Similarly, $V^0(1) = \sum_{t'=t}^\infty \delta_D^{t'-t}u(1)$. Thus $V_t^{\lambda^*}(1) - V^0(1) = \sum_{t'=t}^\infty \delta_D^{t'-t}(\lambda_{t'}^* - \frac{\delta_A}{\delta_D}\lambda_{t'-1}^*)u(1) = -\frac{\delta_A}{\delta_D}\lambda_{t-1}^* u(1) + \sum_{t'=t}^\infty \delta_D^{t'-t}(1-\delta_A)\lambda_{t'}^* u(1)$, which, when $\lambda_{t-1}^* = 0$, equals to $\sum_{t'=t}^\infty \delta_D^{t'-t}(1-\delta_A)\lambda_{t'}^* u(1) \ge 0$. □

Second, we argue that:

Observation (ii). For all $t \le t^F - 1$, $u(p_t^n) + \delta_D p_t^n \kappa V^0(1) > 0$.

Subproof. $t \le t^F - 1$ implies that without IC constraints it is strictly optimal to continue adoption when $p_t = p_t^n$. This implies $H^0(p_t^n) > 0$. Thus $V^0(p_t^n) > 0$ and

$$V^0(p_t^n) = u(p_t^n) + \delta_D p_t^n \kappa V^0(1) + \delta_D(1 - p_t^n\kappa)V^0(p_{t+1}^n) \tag{A.19}$$

Since $p_{t+1}^n < p_t^n$, $V^0(p_{t+1}^n) \le V^0(p_t^n)$. Equation (A.19) then implies

$$V^0(p_t^n) \le u(p_t^n) + \delta_D p_t^n \kappa V^0(1) + \delta_D(1 - p_t^n\kappa)V^0(p_t^n)$$

which implies $u(p_t^n) + \delta_D p_t^n \kappa V^0(1) \ge \big(1 - \delta_D(1 - p_t^n\kappa)\big)V^0(p_t^n)$. Since $V^0(p_t^n) > 0$, we then have $u(p_t^n) + \delta_D p_t^n \kappa V^0(1) > 0$.

□

Now, under ϕ^* suppose for some $t \leq t^F - 1$ we have $\alpha_t^{\phi^*} < 1$ but IC_t is non-binding. Let t_0 be the smallest t satisfying these. Then we must have $\mathbb{P}_{\phi^*}(a_{t_0-1} = 1, p_{t_0} = p_{t_0}^n) > 0$.[27] Then $\alpha_{t_0}^{\phi^*} < 1 \Rightarrow u(p_{t_0}^n) < 0$ by Proposition 4, and IC_{t_0} being non-binding implies $\lambda_{t_0}^* = 0$. We argue that the following are true:

$$\begin{aligned} V_{t_0}^{\lambda^*}(p_{t_0}^n) &\geq (1 + \lambda_{t_0}^* - \frac{\delta_A}{\delta_D}\lambda_{t_0-1}^*)u(p_{t_0}^n) + \delta_D p_{t_0}^n \kappa V_{t_0+1}^{\lambda^*}(1) + \delta_D(1 - p_{t_0}^n\kappa)V_{t_0+1}^{\lambda^*}(p_{t_0+1}^n) \\ &\geq (1 + \lambda_{t_0}^* - \frac{\delta_A}{\delta_D}\lambda_{t_0-1}^*)u(p_{t_0}^n) + \delta_D p_{t_0}^n \kappa V_{t_0+1}^{\lambda^*}(1) \\ &\geq u(p_{t_0}^n) + \delta_D p_{t_0}^n \kappa V_{t_0+1}^{\lambda^*}(1) \\ &\geq u(p_{t_0}^n) + \delta_D p_{t_0}^n \kappa V^0(1) > 0 \end{aligned}$$

The first inequality is due to the Bellman equation (A.13); the second inequality holds since $V_{t_0+1}^{\lambda^*}(\cdot) \geq 0$; the third inequality holds because $\lambda_{t_0}^* = 0$, $\lambda_{t_0-1}^* \geq 0$ and $u(p_{t_0}^n) < 0$; the fourth inequality holds because Observation (i) above implies $V_{t_0+1}^{\lambda^*}(1) \geq V^0(1)$ when $\lambda_{t_0}^* = 0$; finally, the last line above is > 0 by Observation (ii). In sum, we have $V_{t_0}^{\lambda^*}(p_{t_0}^n) > 0$ and thus $H_{t_0}^{\lambda^*}(p_{t_0}^n) > 0$. Since $\mathbb{P}_{\phi^*}(a_{t_0-1} = 1, p_{t_0} = p_{t_0}^n) > 0$, Lemma A.2 then implies $\alpha_{t_0}^{\phi^*} = 1$ – contradiction! This concludes the proof for part (a).

Part (b): Suppose for some $t_0 \geq t^F$ we have $\alpha_{t_0}^{\phi^*} > 0$ but IC_{t_0-1} is non-binding. By our definition about $(\alpha_t)_{t=1}^{\infty}$, this implies $\mathbb{P}_{\phi^*}(a_{t_0-1} = 1, p_{t_0} = p_{t_0}^n) > 0$. Define another no-pause policy ϕ' such that: $\phi_t'(\cdot) = \phi_t^*(\cdot)\, \forall t < t_0$; $\phi_t'(p_{\leq t}, a_{<t}) = \mathbb{1}\{p_t = 1, a_{t-1} = 1\}\, \forall t \geq t_0$. Then ϕ' follows ϕ^* for $t < t_0$ and subsequently follows the first-best policy. Let U_t^{ϕ} be cohort-t's expected utility under ϕ. Then we have $U_t^{\phi^*} = U_t^{\phi'}\, \forall t < t_0$ and $\sum_{t'=t_0}^{\infty} \delta_D^{t'} U_{t'}^{\phi^*} < \sum_{t'=t_0}^{\infty} \delta_D^{t'} U_{t'}^{\phi'}$.[28] Thus ϕ' generates strictly higher designer's value than ϕ^*

Since ϕ^* is incentive compatible with IC_{t_0-1} being slack, we have $U_{t-1}^{\phi^*} \geq \delta_A U_t^{\phi^*}$ for all t and $U_{t_0-1}^{\phi^*} > \delta_A U_{t_0}^{\phi^*}$. Since ϕ' follows ϕ^* before t_0 and subsequently recommend only when good news has arrived, we have $U_{t-1}^{\phi'} \geq \delta_A U_t^{\phi'}$ for both $t < t_0$ and $t > t_0$. These imply that for some $\epsilon > 0$ small enough, a mixed policy ϕ^{ϵ} which follows ϕ^* with probability $1 - \epsilon$ and follows ϕ' with probability ϵ will satisfy all IC constraints in (3.8). Since both ϕ^* and ϕ' have no pause, ϕ^{ϵ} also has no pause. Moreover, since ϕ' generates strictly higher designer's value than ϕ^*, so does ϕ^{ϵ}. This violates the optimality of ϕ^* – contradiction! *Q.E.D.*

Lemma A.5. *(a) IC_{t^F-1} is binding under ϕ^*; (b) $\alpha_{t^F-1}^{\phi^*} \in (0, 1)$ or $\alpha_{t^F}^{\phi^*} \in (0, 1)$.*

[27] Otherwise, for some $t < t_0$ we must have $\alpha_t^{\phi^*} = 0$. Given this, we must also have $U_t = U_{t+1} > 0$ and thus IC_t is non-binding. t_0 is then not the smallest t satisfying the aforementioned conditions.

[28] To see this strict inequality more clearly, notice that by construction ϕ' and ϕ^* induce the same distribution over (p_{t_0}, a_{t_0-1}) and $\mathbb{P}_{\phi^*}(a_{t_0-1} = 1, p_{t_0} = p_{t_0}^n) > 0$. Given this distribution, since $p_{t_0}^n < \eta^F$, the no-pause policy maximizing $\sum_{t'=t_0}^{\infty} \delta_D^{t'} U_{t'}^{\phi}$ (without IC constraints) requires keep adopting from t_0 on if and only if $a_{t_0-1} = 1$ and $p_{t_0} = 1$. ϕ' constructed above satisfies this, while ϕ^* violates it because $\alpha_{t_0}^{\phi^*} > 0$.

Proof. For part (a), suppose IC_{t^F-1} is slack. Then Lemma A.4 implies $\alpha^{\phi^*}_{t^F-1} = 1$ and $\alpha^{\phi^*}_{t^F} = 0$. Notice $\alpha^{\phi^*}_{t^F} = 0$ means that adoption continues after t^F only when good news has arrived. Thus $0 < U_{t^F} = U_{t^F+1} = ...$, which implies IC_t to be slack for all $t \geq t^F$. Moreover, $\alpha^{\phi^*}_{t^F-1} = 1$ implies that no new information will be disclosed in period $t^F - 1$. Thus IC_{t^F-2} is also slack. This further implies $\alpha^{\phi^*}_{t^F-2} = 1$ by Lemma A.4(a), which by the same reasoning implies IC_{t^F-3} is slack. This reasoning in the end implies IC_t is slack for all $t < t^F - 1$. By the complementary slackness condition, the discussion above implies $\lambda^* = \mathbf{0}$ and thus the Lagrangian optimization $\max_{\phi\in\Phi^\dagger} \mathcal{L}(\phi, \lambda^*)$ is the same as the social planner's problem without IC constraints. This implies that the first-best outcome is feasible, which contradicts with our assumption. Thus IC_{t^F-1} is binding.

To show part (b), notice IC_{t^F-1} being binding implies $\alpha^{\phi^*}_{t^F-1} > 0$ and $\alpha^{\phi^*}_{t^F} < 1$ (since otherwise we must have $U_{t^F-1} = U_{t^F}$ under ϕ^*). Suppose $\alpha^{\phi^*}_{t^F-1}, \alpha^{\phi^*}_{t^F} \notin (0, 1)$. Then we must have $\alpha^{\phi^*}_{t^F-1} = 1$ and $\alpha^{\phi^*}_{t^F} = 0$. By the same reasoning as in the proof for part (a), $\alpha^{\phi^*}_{t^F-1} = 1$ further implies $\alpha^{\phi^*}_{t} = 1$ for all $t \leq t^F - 1$. Together with $\alpha^{\phi^*}_{t^F} = 0$, we know ϕ^* coincides with the first-best policy. This contradicts with our assumption that the first-best is not incentive compatible. *Q.E.D.*

We now prove Proposition 7. By definition, $t^{\phi^*}_a := \min\{t \geq 1 : \alpha^{\phi^*}_t < 1\}$ and $t^{\phi^*}_b := \max\{t \geq 1 : \alpha^{\phi^*}_t > 0\}$. Proposition 4 implies $t^{\phi^*}_a \geq t^m$. Recall the proof for Theorem 1 has shown only finitely many IC constraints can be binding under ϕ^*. Thus Lemma A.4(b) implies $t^{\phi^*}_b < +\infty$. To see $t^{\phi^*}_a \leq t^{\phi^*}_b$, notice if not, then for any t either $t < t^{\phi^*}_a$ or $t > t^{\phi^*}_b$, which implies $\alpha^{\phi^*}_t$ to be either 1 or 0. This violates Lemma A.5(b). Moreover, notice Lemma A.5(b) implies $t^{\phi^*}_a \leq t^F$ and $t^{\phi^*}_b \geq t^F - 1$, or equivalently, $[t^{\phi^*}_a, t^{\phi^*}_b] \cap \{t^F - 1, t^F\} \neq \emptyset$.

Notice (i) and (iii) are obvious by definitions of $t^{\phi^*}_a$ and $t^{\phi^*}_b$. For (ii), notice:

- Suppose $\alpha^{\phi^*}_{t_0} = 0$ for some $t_0 \in [t^{\phi^*}_a, t^{\phi^*}_b]$, then by the definition of $(\alpha^{\phi^*}_t)_t$ (equation 5.1), $\alpha^{\phi^*}_t = 0$ for all $t \geq t_0$. Thus $\alpha^{\phi^*}_{t^{\phi^*}_b} = 0$ – contradicting with the definition of $t^{\phi^*}_b$.
- Suppose $\alpha^{\phi^*}_{t_0} = 1$ for some $t_0 \in [t^{\phi^*}_a, t^{\phi^*}_b]$. Notice IC_{t_0-1} must be slack since no new information will be provided in period t_0. If $t_0 \leq t^F - 1$, then by repeatedly referring to Lemma A.4(a) as in the proof of Lemma A.5(a), we get $\alpha^{\phi^*}_t = 1$ for all $t \leq t_0$. In particular, $\alpha^{\phi^*}_{t^{\phi^*}_a} = 1$ – contradicting with the definition of $t^{\phi^*}_a$. If $t_0 \geq t^F$, then by Lemma A.4(b) we must have IC_{t_0-1} being binding – contradicting with IC_{t_0-1} being slack.

Thus we must have $\alpha^{\phi^*}_t \in (0, 1)$ for all $t \in [t^{\phi^*}_a, t^{\phi^*}_b]$. Finally, notice that this further implies IC_t being binding for all $t \in [t^{\phi^*}_a, t^F - 1]$ by Lemma A.4(a), and implies IC_t being binding for all $t \in [t^F - 1, t^{\phi^*}_b - 1]$ by Lemma A.4(b).

A.7 Proof for Corollary 1

Part (a) is directly implied by the definition of $t_a^{\phi^*}$. To see part (b), notice Proposition 7 implies that for all $t \geq t_a^{\phi^*}$, either (i) $\alpha_{t+1}^{\phi^*} = 0$ (i.e., $t+1 > t_b^{\phi^*}$), or (ii) $\alpha_{t+1}^{\phi^*} > 0$ and is just large enough to make IC_t binding (i.e., $t+1 \leq t_b^{\phi^*}$). If $\alpha_t^{\phi^*} = 0$ already, then by expression (5.1) we must have $\alpha_{t+1}^{\phi^*} = 0$ and it satisfies IC_t. This is consistent with the rule in (b). If $\alpha_t^{\phi^*} > 0$ and $\alpha_{t+1}^{\phi^*} = 0$ satisfies IC_t, then any positive α_{t+1} will make IC_t slack,[29] and thus (i) must be the case. If $\alpha_t^{\phi^*} > 0$ and $\alpha_{t+1}^{\phi^*} = 0$ violates IC_t, then (ii) must be the case. These imply that $\alpha_{t+1}^{\phi^*}$ must be determined by the rule in (b).

We now give a more explicit formula for computing $\alpha_{t+1}^{\phi^*}$ in (b). Define $\gamma_t^n = \mathbb{P}_{\phi^*}(p_t = p_t^n)$ (i.e., the probability of adoption not having stopped with no news arrival before t) and $\gamma_t^g = \mathbb{P}_{\phi^*}(p_t = 1)$. Then given $(\alpha_t^{\phi^*})_{t=1}^{\infty}$, sequences of $(\gamma_t^n)_{t=1}^{\infty}$, $(\gamma_t^g)_{t=1}^{\infty}$ and $(U_t)_{t=1}^{\infty}$ can be inductively computed by:

$$\gamma_1^n = 1 - \kappa_0 p_0, \; \gamma_1^g = \kappa_0 p_0 \tag{A.20}$$

$$\gamma_{t+1}^n = \gamma_t^n \alpha_t^{\phi^*}(1-\kappa p_t^n), \; \gamma_{t+1}^g = \gamma_t^g + \gamma_t^n \alpha_t^{\phi^*} \kappa p_t^n, \; U_t = \gamma_t^g H + \gamma_t^n \alpha_t^{\phi^*} u(p_t^n), \; \forall t \geq 1 \tag{A.21}$$

For any $t \geq t_a^{\phi^*}$, given $\alpha_{\leq t}^{\phi^*}$, one could have inductively computed $\gamma_{\leq t+1}^n$, $\gamma_{\leq t+1}^g$ and $U_{\leq t}$ using (A.20) – (A.21). Then, if $(\frac{1}{\delta_A})U_t \geq \gamma_{t+1}^g H$, $\alpha_{t+1}^{\phi^*} = 0$ is feasible and thus should be the case; otherwise, we must have $\gamma_{t+1}^n > 0$ and IC_t should be binding, which implies $\alpha_{t+1}^{\phi^*} = \frac{\frac{1}{\delta_A}U_t - \gamma_{t+1}^g H}{\gamma_{t+1}^n u(p_{t+1}^n)}$ (notice $u(p_{t+1}^n) < 0$ here since $t \geq t^m$). In sum, $\alpha_{t+1}^{\phi^*} = \max\{0, \frac{\frac{1}{\delta_A}U_t - \gamma_{t+1}^g H}{\gamma_{t+1}^n u(p_{t+1}^n)}\}$.

A.8 Optimal $(t_a^{\phi}, \alpha_{t_a^{\phi}}^{\phi})$ and the proof for Proposition 8

We first consider the case with $p_0 > \bar{p}$ (i.e., $u(p_0) > 0$). The other case is simpler but needs some minor modifications in the treatment, which we will explain in the end.

A.8.1 The case with $p_0 > \bar{p}$

- **Characterization for the optimal t_a^{ϕ} and $\alpha_{t_a^{\phi}}^{\phi}$**

We will construct a function $W(\cdot)$ such that $t_a \in \mathbb{Z}_+$ and $\alpha_{t_a} \in (0,1)$ are optimal if and only if $t_a - 1 + \alpha_{t_a} \in \arg\max_{\rho \in [t^m - 1, t^F)} W(\rho)$.

For any t, define $\pi_t = \mathbb{P}(\text{at least one of } s_0, s_1, ..., s_{t-1} \text{ is good})$. Then, by the LIE about p_t, we have $p_0 = \pi_t + (1-\pi_t)p_t^n$, or equivalently, $\pi_t = \frac{p_0 - p_t^n}{1 - p_t^n}$.

Define $\alpha_1^{min} = 0$ and $\alpha_t^{min} = \min\{\alpha \in [0,1] : \pi_t u(1) + (1-\pi_t)\alpha u(p_t^n) \leq \frac{1}{\delta_A} u(p_0)\} \; \forall t \geq 2$. Then α_t^{min} is the smallest value α_t can take without violating IC_{t-1} given $\alpha_{<t} = \mathbf{1}$. Since $\pi_t u(1) + (1-\pi_t)u(p_t^n) = u(p_0) < \frac{1}{\delta_A} u(p_0)$, we have $\alpha_t^{min} < 1$ for all t.

[29] This is because for all $t \geq t^m$ where adoption without good news hasn't fully stopped, the larger is α_{t+1}, the smaller is U_{t+1} and thus the slacker is IC_t.

To make things simpler, we will encode (t_a, α_{t_a}) into a single number. Specifically, for any $\rho \in [t^m - 1, t^F)$, we define ϕ^ρ to be the potentially optimal no-pause policy such that: (i) $t_a^{\phi^\rho} = \lfloor \rho \rfloor + 1$; (ii) $\alpha_{\lfloor \rho \rfloor+1}^{\phi^\rho} = de(\rho) + \big(1 - de(\rho)\big)\alpha_{\lfloor \rho \rfloor+1}^{min}$; (iii) $(\alpha_t^{\phi^\rho})_{t \neq t_a^{\phi^\rho}}$ are computed according to Corollary 1. Here, $de(\rho)$ is the decimal part of ρ and $\lfloor \rho \rfloor = \rho - de(\rho)$. We then just need to determine the optimal choice of $\rho \in [t^m - 1, t^F)$.

For any ρ, let U_t^ρ denote cohort t's expected utility and define $S_t^\rho := \mathbb{P}_{\phi^\rho}(p_t = 1)u(1)$. Intuitively, S_t^ρ *would* be cohort t's expected utility if adoption in period t is recommended only when good news has arrived then under ϕ^ρ. Then one can check $(S_t^\rho, U_t^\rho)_{t=1}^\infty$ satisfies:

$$(S_t^\rho, U_t^\rho) = \big(\pi_t u(1), u(p_0)\big), \ \forall t \leq \lfloor \rho \rfloor \tag{A.22}$$

$$S_{\lfloor \rho \rfloor+1}^\rho = \pi_{\lfloor \rho \rfloor+1} u(1) \tag{A.23}$$

$$U_{\lfloor \rho \rfloor+1}^\rho = \pi_{\lfloor \rho \rfloor+1} u(1) + (1 - \pi_{\lfloor \rho \rfloor+1})\Big[de(\rho) + \big(1 - de(\rho)\big)\alpha_{\lfloor \rho \rfloor+1}^{min}\Big] u(p_{\lfloor \rho \rfloor+1}^n) \tag{A.24}$$

$$S_{t+1}^\rho = S_t^\rho + \frac{U_t^\rho - S_t^\rho}{u(p_t^n)} p_t^n \kappa u(1), \ \forall t \geq t^m \tag{A.25}$$

$$U_{t+1}^\rho = \min\{\frac{1}{\delta_A} U_t^\rho, S_{t+1}^\rho\}, \ \forall t \geq \lfloor \rho \rfloor + 1 \tag{A.26}$$

(A.22) – (A.24) are true by the construction of ϕ^ρ (in particular notice under ϕ^ρ adoption is recommended for sure at any $t \leq \lfloor \rho \rfloor$ no matter whether good news has arrived, which implies $\mathbb{P}_{\phi^\rho}(p_t = 1) = \pi_t, \forall t \leq \lfloor \rho \rfloor + 1$). To see (A.25), notice in general $\mathbb{P}_{\phi^\rho}(p_{t+1} = 1) = \mathbb{P}_{\phi^\rho}(p_t = 1) + \mathbb{P}_{\phi^\rho}(p_t = p_t^n)\alpha_t^{\phi^\rho} p_t^n \kappa$. Thus

$$S_{t+1}^\rho = S_t^\rho + \mathbb{P}_{\phi^\rho}(p_t = p_t^n)\alpha_t^{\phi^\rho} p_t^n \kappa u(1) \tag{A.27}$$

Moreover, $U_t^\rho = \mathbb{P}_{\phi^\rho}(p_t = 1)u(1) + \mathbb{P}_{\phi^\rho}(p_t = p_t^n)\alpha_t^{\phi^\rho} u(p_t^n) = S_t^\rho + \mathbb{P}_{\phi^\rho}(p_t = p_t^n)\alpha_t^{\phi^\rho} u(p_t^n)$, which implies $\mathbb{P}_{\phi^\rho}(p_t = p_t^n)\alpha_t^{\phi^\rho} = \frac{U_t^\rho - S_t^\rho}{u(p_t^n)}$. Taking this back to (A.27), we get (A.25). To see (A.26), notice the rule in Corollary 1 implies that for any $t \geq t_a^{\phi^\rho} = \lfloor \rho \rfloor + 1$, either $\alpha_{t+1}^{\phi^\rho} = 0$ and hence $U_{t+1}^\rho = S_{t+1}^\rho$, or IC_t is binding and therefore $U_{t+1}^\rho = \frac{1}{\delta} U_t^\rho$. Thus we must have $U_{t+1}^\rho = \min\{\frac{1}{\delta_A} U_t^\rho, S_{t+1}^\rho\}$. Notice that equations (A.22) – (A.26) fully pin down the sequence of $(S_t^\rho, U_t^\rho)_{t=1}^\infty$.

Now, define

$$W(\rho) = \sum_{t=1}^\infty \delta_D^t U_t^\rho$$

Then $W(\rho)$ is the designer's value under policy ϕ^ρ. The optimal ρ can thus be found by $\max_{\rho \in [t^m-1, t^F)} W(\rho)$. We show this function $W(\cdot)$ is continuous, piecewise-linear, and quasi-concave below.

Proposition A.1. *(a) S_t^ρ and U_t^ρ are continuous in ρ for all t; (b) $W(\cdot)$ is continuous.*

Proof. Due to discounting and the boundedness of $(U_t^\rho)_{t=1}^\infty$, the continuity of $W(\cdot)$ is implied by continuity of U_t^ρ in ρ. It thus suffices to prove part (a).

First consider continuity over $[k, k+1)$ for some integer $k \in [t^m - 1, t^F)$. When $\rho \in [k, k+1)$, for all $t \leq k$, (A.22) implies S_t^ρ and U_t^ρ are constant thus continuous in ρ. For $t \geq k+1$, we can prove by induction:

- By (A.23), we have $S_{k+1}^\rho = \pi_{k+1}u(1)$, which is constant and thus continuous in ρ. By (A.24) we have $U_{k+1}^\rho = \pi_{k+1}u(1) + (1-\pi_{k+1})\Big[de(\rho) + \big(1 - de(\rho)\big)\alpha_{k+1}^{min}\Big]u(p_{k+1}^n)$, which is also continuous in ρ over $[k, k+1)$.
- Given that U_t^ρ and S_t^ρ are continuous in ρ, (A.25) and (A.26) imply U_{t+1}^ρ and S_{t+1}^ρ are also continuous in ρ.

Second, consider left-continuity at integer $k \in [t^m, t^F - 1]$. Pick any $(\rho_n)_{n=1}^\infty$ in $(k-1, k)$ converging to k. For all $t \leq k-1$, (A.22) implies $(S_t^{\rho_n}, U_t^{\rho_n}) \overset{\forall n}{=} (S_t^k, U_t^k)$. For $t = k$, notice (A.24) implies $U_k^{\rho_n} \overset{n\to\infty}{\to} \pi_k u(1) + (1-\pi_k)u(p_k^n)$, which equals to $u(p_0)$ since $u(\cdot)$ is affine. Thus $U_k^{\rho_n} \overset{n\to\infty}{\to} U_k^k$ since $U_k^k = u(p_0)$ by (A.22). Moreover, (A.22) and (A.23) imply $S_k^k = \pi_k u(1) = S_k^{\rho_n}$ for all n.

For $t \geq k+1$, we show $S_t^{\rho_n} \overset{n\to\infty}{\to} S_t^k$ and $U_t^{\rho_n} \overset{n\to\infty}{\to} U_t^k$ by induction:

- Notice by (A.23), we have $S_k^{\rho_n} \overset{\forall n}{=} \pi_k u(1)$, while by (A.22) we have $S_k^k = \pi_k u(1)$. Thus $S_k^{\rho_n} \overset{n\to\infty}{\to} S_k^k$. Together with $U_k^{\rho_n} \overset{n\to\infty}{\to} U_k^k$, this implies $S_{k+1}^{\rho_n} \overset{n\to\infty}{\to} S_{k+1}^k$ by (A.25). Moreover, by (A.26) we then have $U_{k+1}^{\rho_n} = \min\{\frac{1}{\delta_A}U_k^{\rho_n}, S_{k+1}^{\rho_n}\} \overset{n\to\infty}{\to} \min\{\frac{1}{\delta_A}U_k^k, S_{k+1}^k\} = \min\{\frac{1}{\delta_A}u(p_0), S_{k+1}^k\} = \min\{\frac{1}{\delta_A}u(p_0), \pi_{k+1}u(1)\}$, where the last equality holds because $S_{k+1}^k = \pi_{k+1}u(1)$ by (A.23). Meanwhile, (A.24) implies $U_{k+1}^k = \pi_{k+1}u(1) + (1-\pi_{k+1})\alpha_{k+1}^{min}u(p_{k+1}^n)$, which equals to $\min\{\frac{1}{\delta_A}u(p_0), \pi_{k+1}u(1)\}$ by the definition of α_{k+1}^{min}.[30] Thus $U_{k+1}^{\rho_n} \overset{n\to\infty}{\to} U_{k+1}^k$.
- Given $S_t^{\rho_n} \overset{n\to\infty}{\to} S_t^k$ and $U_t^{\rho_n} \overset{n\to\infty}{\to} U_t^k$, (A.25) and (A.26) imply $S_{t+1}^{\rho_n} \overset{n\to\infty}{\to} S_{t+1}^k$ and $U_{t+1}^{\rho_n} \overset{n\to\infty}{\to} U_{t+1}^k$.

Q.E.D.

For further results, define $\tau(\rho) = \min\{t \geq \lfloor \rho \rfloor + 2 : S_t^\rho \leq \frac{1}{\delta_A}U_{t-1}^\rho\}$. Intuitively, $\tau(\rho)$ is the first t after $t_a^{\phi^\rho}$ at which fully stopping no-good-news adoption won't violate IC_{t-1}. By Proposition 7 and Corollary 1, we know $\tau(\rho)$ just equals to $t_b^{\phi^\rho} + 1$.

Lemma A.6. *(a)* $U_t^\rho = (\frac{1}{\delta_A})^{t-(\lfloor\rho\rfloor+1)}U_{\lfloor\rho\rfloor+1}^\rho$ *for all* $t \in \big[\lfloor\rho\rfloor + 1, \tau(\rho) - 1\big]$*; (b)* $U_t^\rho = S_t^\rho = S_{\tau(\rho)}^\rho$ *for all* $t \geq \tau(\rho)$*; (c)* $\tau(\cdot) < \infty$*; (d)* $\tau(\cdot)$ *is weakly increasing.*

Proof. By the definition of $\tau(\rho)$, $S_t^\rho > \frac{1}{\delta_A}U_{t-1}^\rho$ for all $t \in [\lfloor\rho\rfloor + 2, \tau(\rho) - 1]$. (A.26) then implies $U_{t+1}^\rho = (\frac{1}{\delta_A})U_t^\rho$ for all $t \in [\lfloor\rho\rfloor + 1, \tau(\rho) - 2]$. Thus (a) is true. Also by the definition of $\tau(\rho)$, $S_{\tau(\rho)}^\rho \leq \frac{1}{\delta_A}U_{\tau(\rho)-1}^\rho$. (A.26) then implies $U_{\tau(\rho)}^\rho = S_{\tau(\rho)}^\rho$. (A.25) then implies $S_{\tau(\rho)+1}^\rho = S_{\tau(\rho)}^\rho$, which further implies $U_{\tau(\rho)+1}^\rho = U_{\tau(\rho)}^\rho$ by (A.26). The same reasoning implies all subsequent (S_t^ρ, U_t^ρ) are constant and thus (b) is true.

For part (c), notice adoption never stops when $p_t \geq \bar{p}$ implies $U_{\lfloor\rho\rfloor+1}^\rho \geq U_1^\rho > 0$. If $\tau(\rho) = \infty$, part (a) would imply $U_t^\rho \to \infty$, which cannot happen.

[30] By the definition of α_{k+1}^{min}, if $\pi_{k+1}u(1) \leq \frac{1}{\delta}u(p_0)$, it equals to 0; if $\pi_{k+1}u(1) > \frac{1}{\delta}u(p_0)$, it is such that $\pi_{k+1}u(1) + (1-\pi_{k+1})\alpha_{k+1}^{min}u(p_{k+1}^n) = \frac{1}{\delta}u(p_0)$.

For part (d), pick any $\rho'' > \rho'$ in $[t^m - 1, t^F)$. Suppose $\tau(\rho'') < \tau(\rho')$. Then we have $\lfloor \rho' \rfloor + 1 \leq \lfloor \rho'' \rfloor + 1 < \tau(\rho'') < \tau(\rho')$.

We first show $U_t^{\rho''} \leq U_t^{\rho'}$ for all $t \in [\lfloor \rho' \rfloor + 1, \tau(\rho'') - 1]$. For $t \in [\lfloor \rho' \rfloor + 1, \lfloor \rho'' \rfloor + 1]$, consider two cases:

- Case 1: $\lfloor \rho' \rfloor = \lfloor \rho'' \rfloor$. In this case, $de(\rho'') > de(\rho')$. (A.24) then implies $U_t^{\rho''} \leq U_t^{\rho'}$ (notice $u(p_t^n) < 0$ for $t \geq t^m$) for $t = \lfloor \rho' \rfloor + 1 (= \lfloor \rho'' \rfloor + 1)$.
- Case 2: $\lfloor \rho' \rfloor < \lfloor \rho'' \rfloor$.
  - For $t = \lfloor \rho' \rfloor + 1$, (A.22) implies $U_t^{\rho''} = u(p_0)$, while (A.24) implies $U_t^{\rho'} > \pi_t u(1) + (1 - \pi_t) u(p_t^n) = u(p_0)$ (again notice $u(p_t^n) < 0$ for $t \geq t^m$).
  - For $t \in (\lfloor \rho' \rfloor + 1, \lfloor \rho'' \rfloor + 1)$ (assuming its non-empty), (A.22) implies $U_t^{\rho''} = u(p_0)$, while result in part (a) implies $U_t^{\rho'} = (\frac{1}{\delta_A})^{t - (\lfloor \rho' \rfloor + 1)} U_{\lfloor \rho' \rfloor + 1}^{\rho'} \geq U_{\lfloor \rho' \rfloor + 1}^{\rho'} \geq u(p_0)$.
  - For $t = \lfloor \rho'' \rfloor + 1$, notice (A.24) implies $U_t^{\rho''} \leq \frac{1}{\delta_A} U_{t-1}^{\rho''}$, while part (a) implies $U_t^{\rho'} = \frac{1}{\delta_A} U_{t-1}^{\rho'}$. Thus $U_t^{\rho''} \leq U_t^{\rho'}$ is implied by $U_{t-1}^{\rho''} \leq U_{t-1}^{\rho'}$ established already.

For $t \in (\lfloor \rho'' \rfloor + 1, \tau(\rho'') - 1]$, notice part (a) implies $U_t^{\rho} = (\frac{1}{\delta_A})^{t - (\lfloor \rho'' \rfloor + 1)} U_{\lfloor \rho'' \rfloor + 1}^{\rho}$ for both $\rho = \rho'$ and $\rho = \rho''$. Thus $U_t^{\rho''} \leq U_t^{\rho'}$ is implied by $U_{\lfloor \rho'' \rfloor + 1}^{\rho''} \leq U_{\lfloor \rho'' \rfloor + 1}^{\rho'}$.

Next, we argue that $S_t^{\rho''} \geq S_t^{\rho'}$ for all $t \in [\lfloor \rho' \rfloor + 1, \tau(\rho'')]$. This can be seen by induction. For $t = \lfloor \rho' \rfloor + 1$, (A.22) and (A.23) imply $S_t^{\rho''} = \pi_t u(1) = S_t^{\rho'}$. Moreover, given $U_t^{\rho''} \leq U_t^{\rho'}$ shown above, by (A.25) it is easy to see $S_t^{\rho''} \geq S_t^{\rho'} \Rightarrow S_{t+1}^{\rho''} \geq S_{t+1}^{\rho'}$ (again notice $u(p_t^n) < 0$ for $t \geq t^m$).

The above discussion have proved $S_{\tau(\rho'')}^{\rho''} \geq S_{\tau(\rho'')}^{\rho'}$ and $U_{\tau(\rho'')-1}^{\rho''} \leq U_{\tau(\rho'')-1}^{\rho'}$. By the definition of $\tau(\rho'')$, we have $S_{\tau(\rho'')}^{\rho''} \leq \frac{1}{\delta_A} U_{\tau(\rho'')-1}^{\rho''}$. These together then imply $S_{\tau(\rho'')}^{\rho'} \leq \frac{1}{\delta_A} U_{\tau(\rho'')-1}^{\rho'}$. By the definition of $\tau(\rho')$, we then must have $\tau(\rho') \leq \tau(\rho'')$ – contradiction. Thus $\tau(\cdot)$ is weakly increasing. *Q.E.D.*

Now, we can provide a finer expression for $W(\cdot)$. Define $z_t := -\frac{u(1)}{u(p_t^n)} p_t^n \kappa$. Then for $t \geq t^m$, since $u(p_t^n) < 0$ and p_t^n strictly decreases in t, we know $z_t > 0$ and strictly decreases in t. Define $h_t^j := 1 - \frac{\delta_D z_t}{1 - \delta_D} \prod_{t'=t+1}^{j} \delta_D (1 + z_{t'})$, where we follow the convention that $\prod_{t=a}^{b} (\cdot) = 1$ if $b < a$. Given δ_A, we define

$$K(i, j; \delta_A) := \sum_{t=i}^{j} \delta_D^t (\frac{1}{\delta_A})^{t-i} h_t^j \tag{A.28}$$

Lemma A.7. *For any ρ we have:*

$$W(\rho) = \sum_{t=1}^{\lfloor \rho \rfloor} \delta_D^t u(p_0) + \frac{\delta_D^{\tau(\rho)}}{1 - \delta_D} \prod_{t=\lfloor \rho \rfloor + 1}^{\tau(\rho)-1} (1 + z_t) \pi_{\lfloor \rho \rfloor + 1} u(1) + K\Big(\lfloor \rho \rfloor + 1, \tau(\rho) - 1; \delta_A\Big) U_{\lfloor \rho \rfloor + 1}^{\rho}$$

Proof. (A.25) and Lemma A.6(a) imply for all $t \in \big[\lfloor\rho\rfloor + 1, \tau(\rho) - 1\big]$:

$$S^\rho_{t+1} = S^\rho_t + \frac{(\frac{1}{\delta_A})^{t-(\lfloor\rho\rfloor+1)} U^\rho_{\lfloor\rho\rfloor+1} - S^\rho_t}{u(p^n_t)} p^n_t \kappa u(1)$$

Rearranging terms, we get: $S^\rho_{t+1} = (1 + z_t)S^\rho_t - (\frac{1}{\delta_A})^{t-(\lfloor\rho\rfloor+1)} z_t U^\rho_{\lfloor\rho\rfloor+1}$. This is a linear difference equation on S^ρ_t. By the standard argument we can solve it explicitly:

$$S^\rho_t = \prod_{a=\lfloor\rho\rfloor+1}^{t-1} (1 + z_a) S^\rho_{\lfloor\rho\rfloor+1} - \sum_{a=\lfloor\rho\rfloor+1}^{t-1} \Big[\prod_{b=a+1}^{t-1} (1 + z_b)\Big] (\frac{1}{\delta_A})^{a-(\lfloor\rho\rfloor+1)} z_a \cdot U^\rho_{\lfloor\rho\rfloor+1} \tag{A.29}$$

for all $t \in \big[\lfloor\rho\rfloor + 2, \tau(\rho)\big]$.

Notice that Lemma A.6(a)(b) and (A.22) together imply

$$W(\rho) = \sum_{t=1}^{\lfloor\rho\rfloor} \delta^t_D u(p_0) + \sum_{t=\lfloor\rho\rfloor+1}^{\tau(\rho)-1} \delta^t_D (\frac{1}{\delta_A})^{t-(\lfloor\rho\rfloor+1)} U^\rho_{\lfloor\rho\rfloor+1} + \frac{\delta_D^{\tau(\rho)}}{1-\delta_D} S^\rho_{\tau(\rho)}$$

Substituting expression (A.29) for $S^\rho_{\tau(\rho)}$ in and rearranging, we get:

$$\begin{aligned} W(\rho) =& \sum_{t=1}^{\lfloor\rho\rfloor} \delta^t_D u(p_0) + \frac{\delta_D^{\tau(\rho)}}{1-\delta_D} \prod_{t=\lfloor\rho\rfloor+1}^{\tau(\rho)-1} (1 + z_t) S^\rho_{\lfloor\rho\rfloor+1} \\ &+ \sum_{t=\lfloor\rho\rfloor+1}^{\tau(\rho)-1} (\frac{1}{\delta_A})^{t-(\lfloor\rho\rfloor+1)} \underbrace{\Big[\delta^t_D - \frac{\delta_D^{\tau(\rho)}}{1-\delta_D} \Big[\prod_{t'=t+1}^{\tau(\rho)-1} (1 + z_{t'})\Big] z_t\Big]} U^\rho_{\lfloor\rho\rfloor+1} \end{aligned}$$

Notice the term with under-brace equals to $\delta^t_D h^{\tau(\rho)-1}_t$, and $S^\rho_{\lfloor\rho\rfloor+1} = \pi_{\lfloor\rho\rfloor+1} u(1)$ by (A.23). Thus the expression above equals to the desired one in the lemma. *Q.E.D.*

Next, we show some properties about h^j_t and $K(i, j; \delta_A)$.

Lemma A.8. *Fix $j \geq t^m$. For $t \in [t^m, j]$, h^j_t strictly single-crosses zero from below, i.e., $h^j_t \geq 0 \Rightarrow h^j_{t+1} > 0$, $\forall t \in [t^m, j - 1]$.*

Proof. Fix $t \in [t^m, j - 1]$. Suppose $h^j_{t+1} \leq 0$. By the definition of h^j_{t+1}, we have $1 - \frac{\delta_D z_{t+1}}{1-\delta_D} \prod_{t'=t+2}^{j} \delta_D(1 + z_{t'}) \leq 0$. We argue that this implies $\delta_D(1 + z_{t+1}) \geq 1$. Suppose not. Then $\delta_D(1 + z_{t+1}) < 1$, which implies $\frac{\delta_D z_{t+1}}{1-\delta_D} < 1$. Moreover, recall that $z_{t'}$ strictly decreases in t' for $t' \geq t^m$. Thus $\delta_D(1 + z_{t'+1}) < 1$ for all $t' > t$. Since $z_{t'} > 0$ for $t' \geq t^m$, these together imply $\frac{\delta_D z_{t+1}}{1-\delta_D} \prod_{t'=t+2}^{j} \delta_D(1+z_{t'}) < 1$ – contradiction. Thus $\delta_D(1+z_{t+1}) \geq 1$.

Notice by definition $h^j_t = 1 - (1 - h^j_{t+1}) \frac{z_t}{z_{t+1}} \delta_D(1 + z_{t+1})$. Since $h^j_{t+1} \leq 0$, $z_t > z_{t+1} > 0$ and $\delta_D(1 + z_{t+1}) \geq 1$ as is shown above, we then must have $h^j_t < 0$. Thus $h^j_t \geq 0 \Rightarrow h^j_{t+1} > 0$. *Q.E.D.*

Lemma A.9. *For any $j \geq i \geq t^m$, $K(i,j;\delta_A) \geq 0 \Rightarrow K(i',j';\delta_A) > 0$ for all $(i',j') \neq (i,j)$ such that $j' \geq i'$, $i' \geq i$ and $j' \geq j$.*

Proof. Fix $j \geq i \geq t^m$. Assume $K(i,j;\delta_A) \geq 0$ and pick any $(i',j') \neq (i,j)$ such that $j' \geq i'$, $i' \geq i$ and $j' \geq j$. By definition, $K(i',j';\delta_A) = \sum_{t=i'}^{j'} \delta_D^t (\frac{1}{\delta_A})^{t-i'} h_t^{j'}$. If $h_{i'}^{j'} > 0$, then Lemma A.8 implies $h_t^{j'} > 0$ for all $t \in [i',j']$ and thus $K(i',j';\delta_A) > 0$. It remains to consider the case where $h_{i'}^{j'} \leq 0$. In this case, we argue the following:

$$K(i',j';\delta_A) \geq K(i,j';\delta_A) = \sum_{t=i}^{j'} \delta_D^t (\frac{1}{\delta_A})^{t-i} h_t^{j'} \geq \sum_{t=i}^{j'} \delta_D^t (\frac{1}{\delta_A})^{t-i} h_t^{j} \geq \sum_{t=i}^{j} \delta_D^t (\frac{1}{\delta_A})^{t-i} h_t^{j}$$

(where the first inequality is strict if $i' > i$ and the second is strict if $j' > j$)

The first inequality holds because by Lemma A.8, $h_{i'}^{j'} \leq 0$ implies $h_i^{j'} \leq 0$ (strict if $i < i'$). To see the second inequality, notice $K(i,j;\delta_A) \geq 0$ implies $h_j^j \geq 0$ (otherwise Lemma A.8 would imply $h_t^j < 0$ for all $t \in [i,j]$), which is equivalent to $\delta_D(1+z_j) \leq 1$. This implies $\delta_D(1+z_t) < 1$ for all $t > j$. Thus for all $t \in [i,j']$, we have $h_t^j \leq h_t^{j'}$, which is strict if $j' > j$. The third inequality holds because for all $t > j$, by definition $h_t^j = 1 - \frac{\delta_D z_t}{1-\delta_D} \geq 1 - \frac{\delta_D z_j}{1-\delta_D} = h_j^j$, while $h_j^j \geq 0$ as is mentioned above. Notice the last expression above just equals to $K(i,j;\delta_A)$. Thus $K(i',j';\delta_A) > 0$. *Q.E.D.*

We are now ready to show:

Proposition A.2. *$W(\cdot)$ is piecewise-linear and quasi-concave. Moreover, $sign\big(W'(\rho)\big) = -sign\Big(K\big(\lfloor\rho\rfloor+1, \tau(\rho)-1;\delta_A\big)\Big)$ except for finitely many ρ.*

Proof. Let D_τ be the set of discountinuous points of $\tau(\cdot)$. Since $\tau(\cdot)$ is bounded and weakly increasing, D_τ is finite. Let $D := (D_\tau \cup \mathbb{N}) \cap [t^m - 1, t^F)$. Then D is finite and can be represented as $D = \{d_1, d_2, ..., d_{k^m}\}$ with $d_k < d_{k+1}, \forall k$. Then $\lfloor\cdot\rfloor$ and $\tau(\cdot)$ are constant over any (d_k, d_{k+1}), where we denote their corresponding values as x_k and τ_k. Then for $\rho \in (d_k, d_{k+1})$, Lemma A.7 implies

$$W(\rho) = \sum_{t=1}^{x_k} \delta_D^t u(p_0) + \frac{\delta_D^{\tau_k}}{1-\delta_D} \prod_{t=x_k+1}^{\tau_k - 1} (1+z_t)\pi_{x_k+1} u(1) + K\big(x_k+1, \tau_k - 1; \delta_A\big) U_{x_k+1}^{\rho}$$

Notice $U_{x_k+1}^{\rho}$ is linear in ρ when $\rho \in (d_k, d_{k+1})$ by (A.24), while all other terms above are constant. Thus $W(\cdot)$ is linear over (d_k, d_{k+1}). Given its continuity (Proposition A.1), $W(\cdot)$ is thus indeed piecewise-linear. Moreover, by (A.24) we know $\frac{\partial}{\partial\rho} U_{x_k+1}^{\rho} = (1-\pi_{x_k+1})(1-\alpha_{x_k+1}^{min})u(p_{x_k+1}^n)$, which is < 0 since $\alpha_{x_k+1}^{min} < 1$ and $p_t^n < \bar{p}$ for any $t \geq t^m$. Thus $sign\big(W'(\rho)\big) = -sign\Big(K\big(x_k+1, \tau_k - 1; \delta_A\big)\Big)$ for all $\rho \in (d_k, d_{k+1})$.

To show quasi-concavity, it suffices to show $W'(\cdot) \leq 0$ over (d_k, d_{k+1}) implies $W'(\cdot) < 0$ over (d_{k+1}, d_{k+2}), which is equivalent to $K\big(x_k+1, \tau_k - 1; \delta_A\big) \geq 0 \Rightarrow K\big(x_{k+1}+1, \tau_{k+1} -$

$1;\delta_A) > 0, \forall k$. Notice by construction $x_{k+1} \geq x_k$, $\tau_{k+1} \geq \tau_k$ and $(x_{k+1}, \tau_{k+1}) \neq (x_k, \tau_k)$. The desired result is thus implied by Lemma A.9. *Q.E.D.*

- **Proof for Proposition 8**

We first need several lemmas.

Lemma A.10. *For any $j > i \geq t^m$, $K(i,j;\delta_A) \geq 0 \Rightarrow K(i,j;\delta'_A) > 0$ for all $\delta'_A < \delta_A$.*

Proof. Fix $j > i \geq t^m$ and $\delta'_A < \delta_A$. Assume $K(i,j;\delta_A) \geq 0$. By definition, $K(i,j;\delta'_A) = \sum_{t=i}^{j} \delta_D^t (\frac{1}{\delta'_A})^{t-i} h_t^j$. When $h_i^j \geq 0$, Lemma A.8 implies $h_t^j > 0$ for all $t \in (i,j]$ and thus $K(i,j;\delta'_A) > 0$. When $h_i^j < 0$, by the single-crossing property in Lemma A.8, we can find *non-integer* $c \in (i,j)$ such that $h_t^j < 0 \, \forall t \in [i,c)$ and $h_t^j \geq 0 \, \forall t \in (c,j]$. Then we have:

$$K(i,j;\delta'_A) = \sum_{t=i}^{j} \delta_D^t (\frac{1}{\delta'_A})^{t-i} h_t^j = (\frac{\delta_A}{\delta'_A})^{c-i} \sum_{t=i}^{j} \delta_D^t (\frac{\delta_A}{\delta'_A})^{t-c} (\frac{1}{\delta_A})^{t-i} h_t^j$$
$$> (\frac{\delta_A}{\delta'_A})^{c-i} \sum_{t=i}^{j} \delta_D^t (\frac{1}{\delta_A})^{t-i} h_t^j = (\frac{\delta_A}{\delta'_A})^{c-i} K(i,j;\delta_A) \geq 0$$

The strict inequality holds because given $\frac{\delta_A}{\delta'_A} > 1$, $t \in [i,c)$ (which at least contains i) implies $(\frac{\delta_A}{\delta'_A})^{t-c} < 1$ and $h_t^j < 0$, while $t \in (c,j]$ implies $(\frac{\delta_A}{\delta'_A})^{t-c} > 1$ and $h_t^j \geq 0$. *Q.E.D.*

Recall that $\tau(\cdot)$ is weakly increasing. The following lemma further shows $\tau(\cdot)$ is left-continuous with jump size less than one, except for some points not interesting below.

Lemma A.11. *(a) For any $\rho \in [t^m - 1, t^F)$, $\lim_{x \downarrow \rho} \tau(x) \leq \tau(\rho) + 1$; (b) For any $\rho \in (t^m - 1, t^F)$ that is non-integer or satisfies $\alpha_{\lfloor \rho \rfloor + 1}^{\phi^\rho} > 0$, we have $\lim_{x \uparrow \rho} \tau(x) = \tau(\rho)$.*

Proof. For (a), suppose $x_n \downarrow \rho$ but $\lim \tau(x_n) \geq \tau(\rho) + 2$. Then for n large enough, by the definition of $\tau(\cdot)$, we have $S_{\tau(\rho)+1}^{x_n} > \frac{1}{\delta_A} U_{\tau(\rho)}^{x_n}$. With the continuity result in Proposition A.1(a), this implies $S_{\tau(\rho)+1}^{\rho} \geq \frac{1}{\delta_A} U_{\tau(\rho)}^{\rho}$. However, Lemma A.6(b) implies $U_{\tau(\rho)}^{\rho} = S_{\tau(\rho)+1}^{\rho}$. We then must have $S_{\tau(\rho)+1}^{\rho} \geq \frac{1}{\delta_A} S_{\tau(\rho)+1}^{\rho}$, which implies $S_{\tau(\rho)+1}^{\rho} \leq 0$. However, by definition $S_{\tau(\rho)+1}^{\rho} = \mathbb{P}_{\phi^\rho}(p_{\tau(\rho)+1} = 1) u(1) > 0$ – contradiction!

For (b), fix $\rho \in (t^m - 1, t^F)$ that is non-integer or satisfies $\alpha_{\lfloor \rho \rfloor + 1}^{\phi^\rho} > 0$. Let $x_n \uparrow \rho$.

- Case 1: $\tau(\rho) > \lfloor \rho \rfloor + 2$. In this case, definition of $\tau(\cdot)$ implies $S_{\tau(\rho)-1}^{\rho} > \frac{1}{\delta_A} U_{\tau(\rho)-2}^{\rho}$. By the continuity result in Proposition A.1(a), for n large enough, this implies $S_{\tau(\rho)-1}^{x_n} > \frac{1}{\delta_A} U_{\tau(\rho)-2}^{x_n}$ and thus $S_{\tau(\rho)-1}^{x_n} > U_{\tau(\rho)-1}^{x_n}$ by equation (A.26). By Lemma A.6(b), we then must have $\tau(x_n) > \tau(\rho) - 1$. Thus $\lim \tau(x_n) = \tau(\rho)$.
- Case 2: $\tau(\rho) = \lfloor \rho \rfloor + 2$. When ρ is non-integer, for all n large enough we have $\lfloor x_n \rfloor = \lfloor \rho \rfloor$. Then by definition we must have $\tau(x_n) \geq \lfloor x_n \rfloor + 2 = \tau(\rho)$. Thus $\tau(x_n) \to \tau(\rho)$. When ρ is an integer (which is relevant when $\rho \geq t^m$), let us suppose $\lim \tau(x_n) < \tau(\rho)$. Then we must have $\lim \tau(x_n) = \tau(\rho) - 1$ (if $\lim \tau(x_n) < \tau(\rho) - 1$,

then we must have $\lim\lfloor x_n \rfloor < \tau(\rho) - 3 = \lfloor \rho \rfloor - 1$, which cannot happen). Then for n large enough, we must have $S^{x_n}_{\tau(\rho)-1} \leq \frac{1}{\delta_A} U^{x_n}_{\tau(\rho)-2}$, which by Proposition A.1(a) implies $S^{\rho}_{\tau(\rho)-1} \leq \frac{1}{\delta_A} U^{\rho}_{\tau(\rho)-2}$. As $\tau(\rho) = \lfloor \rho \rfloor + 2$, we then have $S^{\rho}_{\lfloor \rho \rfloor+1} \leq \frac{1}{\delta_A} U^{\rho}_{\lfloor \rho \rfloor}$. By (A.22) and (A.23), we then have $\pi_{\lfloor \rho \rfloor+1} u(1) \leq \frac{1}{\delta_A} u(p_0)$. This implies $\alpha^{min}_{\lfloor \rho \rfloor+1} = 0$. Also when ρ is an integer, $de(\rho) = 0$. We then have $\alpha^{\phi^{\rho}}_{\lfloor \rho \rfloor+1} = 0$ – contradiction!

Q.E.D.

Given any δ_A, let $\tau(\rho; \delta_A)$ denote the corresponding $\tau(\rho)$. Then we have:

Lemma A.12. *$\tau(\rho; \delta_A)$ is weakly increasing in δ_A for any ρ.*

Proof. Fix $\rho \in [t^m - 1, t^F)$ and pick $\delta'_A < \delta''_A$. Suppose $\tau(\rho; \delta''_A) < \tau(\rho; \delta'_A)$. Let $\alpha^{min}_t(\delta_A)$, $U^{\rho}_t(\delta_A)$ and $S^{\rho}_t(\delta_A)$ denote the corresponding α^{min}_t, U^{ρ}_t and S^{ρ}_t given δ_A. Then by its definition, it is easy to see $\alpha^{min}_t(\cdot)$ is weakly increasing for all $t \geq t^m$. Equation (A.24) then implies $U^{\rho}_{\lfloor \rho \rfloor+1}(\delta''_A) \leq U^{\rho}_{\lfloor \rho \rfloor+1}(\delta'_A)$. Together with $\frac{1}{\delta''_A} < \frac{1}{\delta'_A}$, this implies $U^{\rho}_t(\delta''_A) \leq U^{\rho}_t(\delta'_A)$ for all $t \in [\lfloor \rho \rfloor + 1, \tau(\rho; \delta''_A) - 1]$ by Lemma A.6(a).

Moreover, we can show $S^{\rho}_t(\delta''_A) \geq S^{\rho}_t(\delta'_A)$ for all $t \in [\lfloor \rho \rfloor + 1, \tau(\rho; \delta''_A)]$ by induction: first, equation (A.23) implies $S^{\rho}_{\lfloor \rho \rfloor+1}(\delta''_A) = S^{\rho}_{\lfloor \rho \rfloor+1}(\delta'_A) = \pi_{\lfloor \rho \rfloor+1} u(1)$; second, for all $t \in [\lfloor \rho \rfloor + 1, \tau(\rho; \delta''_A) - 1]$, given $S^{\rho}_t(\delta''_A) \geq S^{\rho}_t(\delta'_A)$ and $U^{\rho}_t(\delta''_A) \leq U^{\rho}_t(\delta'_A)$, equation (A.25) implies $S^{\rho}_{t+1}(\delta''_A) \geq S^{\rho}_{t+1}(\delta'_A)$.

Now, notice the definition of $\tau(\rho; \delta''_A)$ implies $S^{\rho}_{\tau(\rho;\delta''_A)}(\delta''_A) \leq \frac{1}{\delta''_A} U^{\rho}_{\tau(\rho;\delta''_A)-1}(\delta''_A)$. Together with results above, we then have $S^{\rho}_{\tau(\rho;\delta''_A)}(\delta'_A) \leq \frac{1}{\delta'_A} U^{\rho}_{\tau(\rho;\delta''_A)-1}(\delta'_A)$. By the definition of $\tau(\rho; \delta'_A)$, this implies $\tau(\rho; \delta'_A) \leq \tau(\rho; \delta''_A)$ – contradiction! *Q.E.D.*

Now, we are ready to show Proposition 8. For any δ_A, let ϕ^{ρ,δ_A} and $W(\rho; \delta_A)$ denote the corresponding ϕ^{ρ} and $W(\rho)$. Pick any $\delta''_A > \delta'_A$. Let ρ'' and ρ' in $[t^m - 1, t^F)$ be the corresponding optimal values of ρ. Suppose $\tau(\rho''; \delta''_A) < \tau(\rho'; \delta'_A)$. Then because $\tau(\rho; \delta_A)$ increases in both ρ and δ_A (Lemma A.6(d) and A.12), we must have $\rho'' < \rho'$. Pick $x_k \searrow \rho''$ and $y_k \nearrow \rho'$.

Since $W(\rho; \delta''_A)$ is continuous and piecewise-linear in $\rho \in [t^m - 1, t^F)$ (Proposition A.1 and A.2), ρ'' being optimal implies $W(\cdot; \delta''_A)$ must be weakly sloping down to the right of ρ''. Proposition A.2 then implies $K\Big(\lfloor x_k \rfloor + 1, \tau(x_k; \delta''_A) - 1; \delta''_A\Big) \geq 0$ for k large enough. Notice when k is large enough, $\lfloor x_k \rfloor = \lfloor \rho'' \rfloor$ and $\tau(x_k; \delta''_A)$ goes down to either $\tau(\rho''; \delta''_A)$ or $\tau(\rho''; \delta''_A) + 1$ (Lemma A.11(a)). By Lemma A.9 we then at least have $K\Big(\lfloor \rho'' \rfloor + 1, \tau(\rho''; \delta''_A); \delta''_A\Big) \geq 0$, which implies $K\Big(\lfloor \rho'' \rfloor + 1, \tau(\rho''; \delta''_A); \delta'_A\Big) > 0$ by Lemma A.10. Now, distinguish between two cases:

- Case 1: ρ' is non-integer. In this case, when k is large enough, $\lfloor y_k \rfloor + 1 = \lfloor \rho' \rfloor + 1 \geq \lfloor \rho'' \rfloor + 1$, and by Lemma A.11(b) $\tau(y_k; \delta'_A) - 1 = \tau(\rho'; \delta'_A) - 1 \geq \tau(\rho''; \delta''_A)$. By Lemma A.9, $K\Big(\lfloor \rho'' \rfloor + 1, \tau(\rho''; \delta''_A); \delta'_A\Big) > 0$ then implies $K\Big(\lfloor y_k \rfloor + 1, \tau(y_k; \delta'_A) - 1; \delta'_A\Big) = K\Big(\lfloor \rho' \rfloor + 1, \tau(\rho'; \delta'_A) - 1; \delta'_A\Big) > 0$ for all k large enough. Proposition A.2 then implies $W'(y_k;\delta'_A) < 0$ for k large enough, which implies $W(\cdot;\delta'_A)$ being strictly decreasing to the left of ρ'. This violates the optimality of ρ' given δ'_A.

- Case 2: ρ' is an integer. In this case, $\alpha^{\phi^{\rho',\delta'_A}}_{\lfloor\rho'\rfloor+1}$ cannot be zero, since otherwise $\alpha^{\phi^{\rho',\delta'_A}}_t$ would be either 0 or 1 for all t, which contradicts with Proposition 7. When k is large enough, we then have $\lfloor y_k \rfloor + 1 = \rho' \geq \lfloor \rho'' \rfloor + 1$, and by Lemma A.11(b) $\tau(y_k;\delta'_A) - 1 = \tau(\rho';\delta'_A) - 1 \geq \tau(\rho'';\delta''_A)$. By Lemma A.9, $K\Big(\lfloor\rho''\rfloor+1, \tau(\rho'';\delta''_A); \delta'_A\Big) >$ 0 then implies $K\Big(\lfloor y_k\rfloor + 1, \tau(y_k;\delta'_A) - 1; \delta'_A\Big) = K\Big(\rho', \tau(\rho';\delta'_A) - 1; \delta'_A\Big) > 0$ for all k large enough. This again implies $W(\cdot;\delta'_A)$ being strictly decreasing to the left of ρ', which violates the optimality of ρ'.

The contradiction derived above implies that we cannot have $\tau(\rho'';\delta''_A) < \tau(\rho';\delta'_A)$. By the definition of τ, this implies $t''_b \geq t'_b$, which concludes the proof.

A.8.2 The case with $p_0 \leq \bar{p}$

When $p_0 \leq \bar{p}$, for Assumption 1(a) to hold, we must have $\kappa_0 > 0$. Notice randomized termination of adoption without good news must start in period 1. Otherwise $U_1 \leq 0$ and thus $U_t \leq 0\,\forall t$ by the IC constraints, which cannot be optimal under Assumption 1(a). This implies that we can restrict ρ to be within $(\underline{\rho}, 1)$, where $\underline{\rho}$ is such that $p_0\kappa_0 u(1) + (1-p_0\kappa_0)\underline{\rho}u(p^n_1) = 0$. This keeps $U^\rho_1 > 0$.

Once we restrict to $\rho \in (\underline{\rho}, 1)$, all previous proofs and results in A.8.1 remain valid. Just ignore any case under discussion that matters only when ρ may be outside $(\underline{\rho}, 1)$. In particular, one never needs to refer to $(\alpha^{min}_t : t > 1)$.

A.9 Proof for Proposition 9

Given any ϕ, we define $\gamma^n_t = \mathbb{P}_\phi(p_t = p^n_t)$ and $\gamma^b_t = \mathbb{P}_\phi(p_t = 0, a_{t-1} = 1)$ (i.e., the probability of having received bad news but adoption hasn't stopped at t). Then given $(\beta^\phi_t)^\infty_{t=1}$, sequences of $(\gamma^n_t)^\infty_{t=1}$, $(\gamma^b_t)^\infty_{t=1}$ and $(U_t)^\infty_{t=1}$ can be inductively computed by:

$$\gamma^n_1 = 1 - \kappa_0(1-p_0),\ \gamma^b_1 = \kappa_0(1-p_0) \tag{A.30}$$

$$\gamma^n_{t+1} = \gamma^n_t(1-\kappa(1-p^n_t)),\ \gamma^b_{t+1} = \gamma^b_t\beta^\phi_t + \gamma^n_t\kappa(1-p^n_t),\ \forall t \geq 1 \tag{A.31}$$

$$U_t = \gamma^n_t u(p^n_t) + \gamma^b_t\beta^\phi_t L,\ \forall t \geq 1 \tag{A.32}$$

Key to the proof is to notice that for any t_0, if we just decrease $\beta^\phi_{t_0}$, then U_t will (weakly) increase for all t and IC_t will remain satisfied for all $t \neq t_0 - 1$. To see this, notice $\beta^\phi_{t_0}$ does not affect U_t for $t < t_0$ and thus does not affect IC_t for $t < t_0 - 1$. It thus suffices to consider $t \geq t_0$. Notice (A.32) implies $U_t - \delta_A U_{t+1} = \gamma^n_t u(p^n_t) + \gamma^b_t\beta^\phi_t L -$

$\delta_A[\gamma^n_{t+1}u(p^n_{t+1}) + \gamma^b_{t+1}\beta^\phi_{t+1}L]$. Substituting with (A.31) and rearranging, we get:

$$U_t - \delta_A U_{t+1} = \gamma^n_t u(p^n_t) - \delta_A \gamma^n_{t+1} u(p^n_{t+1}) + (1-\delta_A\beta^\phi_{t+1})\gamma^b_t \beta^\phi_t L \\ - \delta_A \gamma^n_t \kappa(1-p^n_t)\beta^\phi_{t+1} L \tag{A.33}$$

By the induction rule of $(\gamma^n_t)^\infty_{t=1}$ and $(\gamma^b_t)^\infty_{t=1}$, it is easy to see that $(\gamma^n_t)^\infty_{t=1}$ does not depend on $(\beta^\phi_t)^\infty_{t=1}$, while $(\gamma^b_t)_{t>t_0}$ are increasing in $\beta^\phi_{t_0}$. Thus when $\beta^\phi_{t_0}$ decreases, for all $t \geq t_0$, the RHS of (A.32) and (A.33) will increase (notice $L < 0$), which implies U_t increases and IC_t becomes slacker.

Now, to see part (a), notice if $\beta^\phi_1 > 0$, then by decreasing it we can strictly increase U_1. At the same time, the argument above implies that no cohort will be harmed and all IC constraints will remain satisfied. Thus $\beta^\phi_1 > 0$ cannot be optimal.

For part (b), first consider the case where $\beta^\phi_{t+1} = 0$ satisfies IC_t. In this case, if $\beta^\phi_{t+1} > 0$, then we can decrease it a little bit without violating IC_t. This strictly increases U_{t+1}, and by the argument above won't harm any cohort or violate any other IC constraint. Thus we must have $\beta^\phi_{t+1} = 0$. In the case where $\beta^\phi_{t+1} = 0$ violates IC_t, we must have $\beta^\phi_{t+1} > 0$. If IC_t is non-binding, then similar argument implies a small decrease in β^ϕ_{t+1} will be feasible and strictly beneficial. Thus IC_t must hold as equality.

Finally, we give a more concrete formula for computing β^ϕ_{t+1} in (b). Given $\beta^\phi_{\leq t}$, one could have inductively computed $\gamma^n_{\leq t+1}$, $\gamma^b_{\leq t+1}$ and $U_{\leq t}$ using (A.30) – (A.32). Then, if $(\frac{1}{\delta_A})U_t \geq \gamma^n_{t+1}u(p^n_{t+1})$, $\beta^\phi_{t+1} = 0$ is feasible and thus should be the case; otherwise, IC_t should be binding, which implies $\beta^\phi_{t+1} = \frac{\frac{1}{\delta_A}U_t - \gamma^n_{t+1}u(p^n_{t+1})}{\gamma^b_{t+1}L}$. To sum up, $\beta^\phi_{t+1} = \max\{0, \frac{\frac{1}{\delta_A}U_t - \gamma^n_{t+1}u(p^n_{t+1})}{\gamma^b_{t+1}L}\}$.[31]

B Notes on Numerically Solving the Dual (3.13)

One tricky part of numerically solving the dual problem (3.13) is that the control variable $(\lambda_t)^\infty_{t=1}$ has infinite dimensions. Fortunately, as we have shown in Theorem 1, the solution can only have finitely many non-zero entries. This motivates a simple guess-and-verify approach. We first guess a horizon T beyond which the IC constraints in (3.8) will be non-restrictive, and solve the dual problem with $\lambda_t = 0$ for all $t > T$. This T-dimensional convex optimization can be solved relatively easily. Then, with the multipliers found, we can find a candidate optimal solution $\hat{\phi}$ by solving the Lagrangian maximization using standard dynamic programming. Finally, we check whether all IC constraints are indeed satisfied. If not, we can go back to extend the horizon T.

[31] We note that the computation can be stopped once $u(p^n_t) > \delta_A H$ and $\beta^\phi_t = 0$, because thereafter the IC constraint will always be slack with $\beta^\phi_{t'} = 0 \, \forall t' > t$. This must happen at some finite t, since $u(p^n_t) \uparrow H$ and only finitely many IC constraints can be binding.

One issue here is how we can check all IC constraints beyond period T since there are infinitely many of them. The trick we take is to stop checking once the following condition holds at the current $t > T$:

$$\mathbb{E}_{\hat{\phi}}\big[\big(u(p_t) - \delta_A H\big)a_t\big] \geq 0 \tag{B.1}$$

Notice H is the highest utility an agent can receive and $a_t \geq a_{t'}, \forall t' > t$ under a no-pause policy. Thus the condition guarantees $\mathbb{E}_{\hat{\phi}}[a_t u(p_t)] \geq \delta_A \mathbb{E}_{\hat{\phi}}[a_{t'} u(p_{t'})]$ for all $t' > t$. Also notice $\hat{\phi}$ must follow the first-best since period $T+2$, which never stops adoption when $u(p_t) > 0$. Thus $\mathbb{E}_{\hat{\phi}}[a_t u(p_t)]$ must be increasing in $t > T$. These together imply that $IC_{t'}$ must be satisfied for all $t' > t$ and there is no need to check them directly.

To see condition (B.1) must be satisfied at some finite t when $\hat{\phi}$ is indeed incentive-compatible (thus optimal), notice that by the martingale convergence theorem, $(p_t)_{t=1}^{\infty}$ almost surely converge. When it converges to some $p^* \in (0,1)$, a_t must a.s. converges to 0, since otherwise Bayesian updating will drive p_t away from its limit; when it converges to 0, a_t also must converge to 0 since $\hat{\phi}$ eventually follows the first-best; when $(p_t)_{t=1}^{\infty}$ converges to 1, we have $u(p_t) \to H$. These together imply $\big(u(p_t) - H\big)a_t \to 0$ a.s. By the bounded convergence theorem, we then have $\mathbb{E}_{\hat{\phi}}\big[\big(u(p_t) - H\big)a_t\big] \overset{t\to\infty}{\to} 0$. When $\hat{\phi}$ is indeed optimal, Assumption 1(i) and Proposition 4 together imply $\mathbb{P}_{\hat{\phi}}(a_t = 1)$ must be uniformly bounded away from zero, since otherwise $(p_t)_{t=1}^{\infty}$ cannot satisfy the martingale property. Thus $\mathbb{E}_{\hat{\phi}}\big[\big(u(p_t) - H\big)a_t\big] \overset{t\to\infty}{\to} 0$ implies that we must have $\mathbb{E}_{\hat{\phi}}\big[\big(u(p_t) - \delta_A H\big)a_t\big] > 0$ from some t on.

References

Balder, E. J. (1989). On compactness of the space of policies in stochastic dynamic programming. *Stochastic processes and their applications*, *32*(1), 141–150.

Ball, I. (2019). Dynamic information provision: Rewarding the past and guiding the future. *Available at SSRN 3103127*.

Che, Y.-K., & Hörner, J. (2018). Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics*, *133*(2), 871–925.

Chen, C.-H., Eraslan, H., Ishida, J., & Yamashita, T. (2024). Optimal feedback dynamics against free-riding in collective experimentation. *ISER DP*(1247).

Dynkin, E. B., & Yushkevich, A. A. (1979). *Controlled markov processes* (Vol. 235). Springer.

Ely, J. C. (2017). Beeps. *American Economic Review*, *107*(1), 31–53.

Ely, J. C., & Szydlowski, M. (2020). Moving the goalposts. *Journal of Political Economy*, *128*(2), 468–506.

Feinberg, E. A., & Piunovskiy, A. B. (2000). Multiple objective nonatomic markov decision processes with total reward criteria. *Journal of mathematical analysis and applications*, *247*(1), 45–66.

Frick, M., & Ishii, Y. (2024). Innovation adoption by forward-looking social learners. *Theoretical Economics*, *19*(4), 1505–1541.

Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, *101*(6), 2590–2615.

Karlin, S., & Rinott, Y. (1980). Classes of orderings of measures and related correlation inequalities. i. multivariate totally positive distributions. *Journal of Multivariate Analysis*, *10*(4), 467–498.

Knoepfle, J., & Salmi, J. (2024). Dynamic evidence disclosure: Delay the good to accelerate the bad. *arXiv preprint arXiv:2406.11728*.

Kremer, I., Mansour, Y., & Perry, M. (2014). Implementing the "wisdom of the crowd". *Journal of Political Economy*, *122*(5), 988–1012.

Lorecchio, C. (2021). Pursuading crowds. *unpublished manuscript*.

Luenberger, D. G. (1997). *Optimization by vector space methods.* John Wiley & Sons.

Lyu, C. (2023). Information design for social learning on a recommendation platform. *Available at SSRN 4275515*.

Martin, K., Ryan, C. T., & Stern, M. (2016). The slater conundrum: duality and pricing in infinite-dimensional optimization. *SIAM Journal on Optimization*, *26*(1), 111–138.

Myerson, R. B. (1986). Multistage games with communication. *Econometrica: Journal of the Econometric Society*, 323–358.

Orlov, D., Skrzypacz, A., & Zryumov, P. (2020). Persuading the principal to wait. *Journal of Political Economy*, *128*(7), 2542–2578.

Rayo, L., & Segal, I. (2010). Optimal information disclosure. *Journal of Political Economy*, *118*(5), 949–987.

Renault, J., Solan, E., & Vieille, N. (2017). Optimal dynamic information provision. *Games and Economic Behavior*, *104*, 329–349.

Schäl, M. (1975). On dynamic programming: compactness of the space of policies. *Stochastic Processes and their Applications*, *3*(4), 345–364.

Smolin, A. (2021). Dynamic evaluation design. *American Economic Journal: Microeconomics*, *13*(4), 300–331.

Sundaram, R. K. (1996). *A first course in optimization theory.* Cambridge university press.