STAGGERED ROLLOUT FOR INNOVATION ADOPTION[*]

Ricardo Fonseca[†]

July 28, 2023

Preliminary and incomplete: please do not circulate

Abstract

Consider a government trying to reach mass immunity through vaccination before new virus mutations kick in, a matchmaker (or service) platform wanting to reach a critical mass of adopters on one side of the market, or a planner coordinating the take-up of a new pest control product in a region. When the adoption of innovation generates information on its value for others, take-up occurs with delay, which is undesirable for a principal who wants to reach a target adoption rate as soon as possible. We study how a principal can use supply availability to hamper strategic delays and discuss how optimal rollout plans should change depending on the distribution of payoffs in the population. We also characterize when take-up contingent commitment power can be valuable to the principal.

[*]I thank Bobak Pakzad-Hurson for his supervision and kind support. I also thank Jack Fanning, Teddy Mekonnen, Kareen Rozen, Roberto Serrano, and Brown University seminar audiences for constructive comments, discussions, .

[†]Brown University, 8 Fones Alley, Providence, RI 02906. Email: ricardo_fonseca@brown.edu

I Introduction

In various situations, a product or service supplier wants to reach an adoption target as soon as possible. Examples of such settings include a government trying to reach mass immunity before new virus mutations come by, a matchmaker (or service) platform wanting to reach a critical mass of adopters as soon as possible on one side of the market (Evans and Schmalensee, 2016) or a new product for pest control by farmers (Reeves, Ohtsuki, and Fukui, 2017). In general, new innovative products in different sectors must often show their profitability to managers and company shareholders as soon as possible.

One of the most consistent regularities found in cases of innovation diffusion is the S-shaped format of adoption curves. That is, take-up at first increases convexly, at an increasing rate, and after a certain point, the growth rate decreases, and the curve becomes concave. An example from a highly influential work by Ryan and Gross (1943) is the adoption of hybrid corn by Iowa farms in the 1920s and 1930s:

Figure 1: Example of S-shaped adoption curve

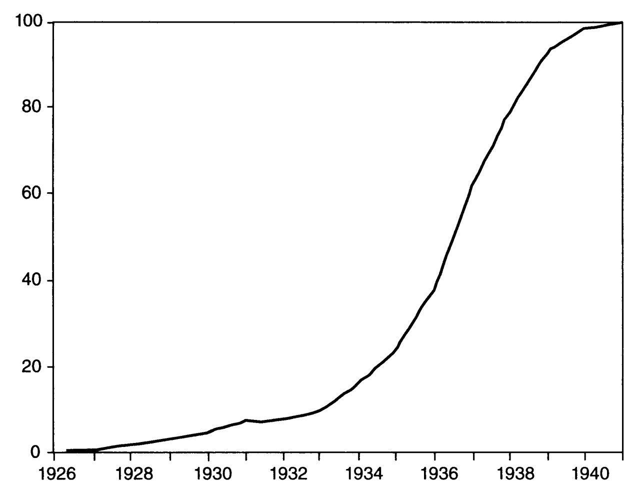

Percent	0.0	0.4	1.4	3.0	4.6	6.9	7.3	9.6	15.8	24.0	38.0	62.0	79.0	93.0	98.5	99.7
Year	1926	1927	1928	1929	1930	1931	1932	1933	1934	1935	1936	1937	1938	1939	1940	1941

FIGURE 5. PERCENT OF ADOPTERS OF HYBRID CORN IN TWO IOWA COMMUNITIES, 1926–1941
(*From Ryan and Gross 1943, Figure 4*)

This adoption pattern is widespread in diffusion studies. In his influential book, Rogers Everett (2003a) goes as far as to say that:

"The S-curve of diffusion is so ubiquitous that students of diffusion often expect every innovation to be adopted over time in an S-shaped pattern".

One important reason we observe lags in adopting a new product of unknown value is social learning: the adoption of a good by others generates information on quality that other agents can gather, but only over time. The literature empirically documenting social learning, especially for health and agricultural innovations, is large.[1] Ryan and Gross (1943), cited above, present survey evidence that the most influential factor in farmers' decision on whether to implement hybrid corn was the previous adoption of the technique by their neighbors.

Adoption delays are not desirable for a supplier who wants to reach a target adoption rate quickly. This paper studies how far a planner can reach a target rate earlier by committing to product supply plans.

To understand how controlling availability for the good might lead to faster take-up, consider first a supplier willing to serve any individual willing to adopt the innovation at any point in time irreversibly. Knowing that some people might want to take up the good before her creates incentives for her to strictly prefer to *strategically delay* adoption and learn more about it before adopting.

Consider now a setting where the product is only available in batches, with any unclaimed units subject to take-up as long as others did not claim them. If the future batch is set to come only too late, a potential adopter would rather take up earlier, given that she has to compete with others for the scarce units available.

This paper studies how a principal can use competition for scarcely available units of a good (or service) to reach an adoption target faster. Importantly, we characterize the optimal supply strategy for a principal with commitment power. We also discuss when stronger commitment power for the principal, allowing for changes contingent on take-up, is advantageous.

We consider a model with a continuum of agents who can irreversibly adopt a good of unknown binary quality. They learn about the quality of the good through a perfect bad news Poisson process. In particular, if no news arrives, they become more confident that the product is high quality.

To understand how powerful a tool competition can really be, suppose that all players homo-

[1] For developing countries, see Besley and Case (1993); Besley, Case et al. (1994); Foster and Rosenzweig (1995); Kremer and Miguel (2007); Conley and Udry (2010); Dupas (2014). For evidence of forward-looking behavior, see Munshi (2004); Bandiera and Rasul (2006).

geneously have homogeneous valuations for the good whenever the state is good. Suppose that slightly (arbitrarily so) less than enough mass to serve all agents is made available to all potential adopters. Suppose that there is an equilibrium in which some individual decides to take up the product at a later time than time 0. If more people than available units decide to adopt at this final moment, some will not get it. But they could apply just one moment before and get the good for sure, which is a profitable deviation. If, instead, we see adoption happening up to some point in time with a final positive mass of adopters, given the supply restriction, some agent who does not get the good can apply earlier to get it and will be able to. Therefore, *all* adoption happens at time 0.

If a principal's objective is to hit an adoption target as soon as possible, as when wanting to reach mass immunity from vaccines before new mutations kick in, or take-up of a new pest control product in a region, or adoption of a new matching platform wanting to reach a critical mass of adopters, supply becomes a relevant tool, especially given that the logic described in the previous paragraph does not depend on the amount of scarcity: any level not enough to serve all agents leads to immediate adoption by strategic agents.

Consider now the tension coming from introducing two types of agents, each with different valuations for the product in the good state: competition generates earlier and more "learning" for agents less willing to adopt the good. Note, however, that this can only work as long as we do not have incentives for higher types to also free-ride and adopt later on. The main point is that given the choice between taking up at the initial moment and at the final one, with the less willing agents, to have scarcity play a role in the benefit of the principal, we must have high types preferring the former.

Note that "simple" supply plans, with only the release of a couple of batches, are optimal for the principal when there are only two types. The reasoning is simple: if any other plan were to be strictly better than any in this class, it must lead to earlier adoption for the more skeptical types. But then this would mean that no high-type agents would be willing to adopt from the start. The same logic holds true when more types are considered.

Adding a new type leads to new insights: the optimal plan consists of, essentially two steps: i) Checking whether adding a mid-batch is strictly better for the principal (which is done by checking for the mid-type preference over adoption times), but if so, ii) try to please the mid-type as much as possible. The intuition for this second step comes from the fact that restricting supply for the highest type agents to have earlier mid-type adoption. This comes from the fact that earlier adoption from higher types is *more valuable* for the principal than later adoption of mid-level valuation types. In turn, this comes from the fact that adoption is easier to be achieved for more willing types (those with a higher valuation) and preferred because the whole stock of adoption is used for changing beliefs, and therefore a positive mass of earlier adopters will lead to a higher area over time then later adopters.

Finally, we discuss some extensions of the model for different principal objective functions, and welfare, and discuss what happens when the number of types goes to infinity and we are close to a uniform distribution of types. In particular, given a few restrictions, we show that the S-shaped pattern document in adoption studies can be accounted for when considering the uniform distribution of types. This means that there is no need to assume stochastic adoption opportunities, or myopic behavior, as done by the literature so far.

Generalizing things to a model with an infinite number of types, it is natural to think of valuations as uniformly distributed. The S-shaped adoption cited above can result from this heterogeneity, without relying on random opportunities to adopt or myopia from agents.

Related Literature

This paper relates to several branches of the social learning literature, which has used various assumptions in various models exploring different situations.

Most closely related to our setting are Frick and Ishii (2020) and Laiho and Salmi (2018), that we briefly describe below:

Frick and Ishii (2020) consider a continuum population of agents with homogeneous preferences who can irreversibly take up a good of uncertain quality and see a perfect Poisson public signal depending on the mass of adopters. They find that the unique equilibrium adoption curve is increasing convexly up to a certain point for bad news signals. They also find that having more potential adopters of innovation can decrease welfare in some particular situations. We build upon their model, also considering a perfect bad news Poisson process. Instead, we focus on the problem of a principal with an objective different from the welfare maximization of agents and who can control product availability.

Laiho and Salmi (2018) also consider a continuum model of agents who can irreversibly take up a good of uncertain quality and studies how monopolists can use dynamic prices depending on information generated by a Brownian signal. Our paper focuses on a different type of tool when prices cannot be used: artificial scarcity. Besides that, the principal in our setting cares about reaching an adoption target as soon as possible, a fundamentally different objective from maximizing profit or revenue. In their setting, a principal can use the increasing optimism of agents to extract surplus from their higher willingness to pay. This channel is completely absent from our setting and should not be present in some of the situations we are referencing (e.g.the vaccination case).

Some papers also consider social learning when signals of quality depend on the adoption decisions of others. Young (2009) focuses on parsing out adoption from mimicking and social learning through the shape of the adoption curve, assuming that agents are myopic about their dynamic incentives to delay adoption. Perego and Yuksel (2016) analyze a model in which network-connected agents can learn from their own experience or that of others, aggregating information, and show that increasing connectivity may decrease information quality. Wolitzky (2018) studies a model with observable outcomes of a random sample of players but not their actions. It shows that inefficiencies can persist as the sample size goes to infinity and even increase with it.

Other papers have considered artificial scarcity as a tool a principal may use to achieve its objectives. In a recent paper, Parakhonyak and Vikander (2023) show that artificial scarcity may be used as a quality signal if players observe only overdemand and may infer high quality by assuming that others observed high signal realizations. Scarcity is a signaling tool, unlike in the present paper, where it is used to induce faster take-up through competition effects. DeGraba (1995); Nocke and Peitz (2007); Möller and Watanabe (2010) consider models in which there is no social learning from the number of adopters, as in this paper.

Two other papers of note are Bonatti (2011) and Che and Mierendorff (2019). They consider monopolists and social planners with one or many goods of different but unknown quality. Similarly to Laiho and Salmi (2018), Bonatti (2011) studies a monopoly firm of a durable good that uses dynamic prices to maximize revenue. Che and Mierendorff (2019) study recommender systems (think about this term, it sounds strange) of goods facing short-lived customers (who, therefore, cannot strategically delay adoption). A social planner recommends a good product to some. Still, not all agents, in the absence of perfectly good news, as committing to recommending it to all agents leads to a less informative recommendation (spam).

The papers above show a diverse set of reasons for a planner or monopolist to restrict access to goods to increase profits or welfare. We contribute to this literature by showing that competition effects may be used to reach adoption targets faster in settings where signals of quality may depend on aggregate adoption for reasons that are different from profit maximization, which is not a good representation of the objectives of suppliers of innovative goods in many cases.

Herding models (sounds like I've mentioned it before - change it with a sentence explaining what these are about) assume that individuals observe their actions but not signals of quality depending on their actions. It is common for this literature to be referred to as "social learning" literature. Still, our settings are substantially different, although related, in that the previous decisions of others matter for each agent. It is also the case that agents observe private signal realizations, and the literature here is vast (Scharfstein and Stein, 1990; Banerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992; Smith and Sørensen, 2000; Eyster and Rabin, 2014; Smith, Sørensen, and Tian, 2021).

There is a vast literature, involving many fields besides economics, on the diffusion of innovations. As mentioned above, a good review can be seen on Young (2009). Another famous reference is written by one of the founders of the diffusion field Rogers Everett (2003b). As mentioned before, almost all papers assume that agents myopically adopt the innovation.

Part of the diffusion literature focuses on social learning in networks. This includes (Jackson, 2010; Mossel, Sly, and Tamuz, 2015; Akbarpour and Jackson, 2018) and has been recently reviewed by Golub and Sadler (2017). The focus here is on network structure, so the forward-looking behavior of agents is generally not assumed. S-shaped adoption curves are also robustly observed.

This paper is structured as follows: section 2 presents the model setup; section 3 will then state and give intuition to the results; finally, section 4 discusses the results from the previous section and considers extensions.

II Model

II.A Basic Structure

Consider a population composed of a continuum of agents with mass normalized to 1. We will use a superscript i for general agents. There is a persistent state of the world $\omega \in \{b,g\}$ (also referred to henceforth as "bad" and "good" states, respectively) with a probability $p(\omega) \in (0,1)$ of a ω realization. At each moment $t \in [0,\infty)$, if there is available supply, each agent can apply to take up the good irreversibly. We will use the terms "take up the good" and "adopt the innovation", or "adopt" interchangeably.

All agents receive the same payoff from adopting if the state is b, which we normalize to -1. If the state is g, however, agents receive type-dependent payoffs. To ease notation, we will denote an agent type by her g-state payoff v^n, with $n \in \mathcal{N} \equiv \{1,2,...,N\}$, and $v^n > v^{n+1}$ for any $n < N$. We also have that $v^n > 0$ for each n. The entire vector of valuations will be denoted, by $v^{\mathcal{N}} \equiv (v^1,...,v^N)$. Agents get a flow payoff of 0 at each moment of time when they do not adopt. The mass of agents of type v^n for each $n \in \mathcal{N}$ is given by q^n, with $\sum q^n_{n \in \mathcal{N}} = 1$. The vector representing the mass of agents with each valuation is represented by $q^{\mathcal{N}} \equiv (q^1,...,q^N)$. Agents discount future payoffs exponentially with a common discount rate $r > 0$.

II.A.1 Learning Process

Agents learn about the state of the world through a public signal. A public signal process better represents the settings that motivate this paper, such as vaccination and pest control. We focus on **perfect bad news Poisson public signals**, which are frequently studied by the strategic experimentation literature (e.g. Keller and Rady (2015)). This means that no signal realization can ever happen if the state is good. If it is bad, however, there is a positive probability that a perfectly informative signal will reveal to all that the state is b. When a realization occurs, we say that a *breakdown* happened, borrowing the term from the literature.

Crucially, previous adoption decisions of others influence the arrival rate of news. This means that learning happens socially, but not only through the observation of adoption decisions of others, as in herding models[2]. Formally, the arrival rate of bad news at time t depends on the mass of adopters up to that point in time, M_t, absent realizations up to that point. The arrival rate at time t is given by $\beta M_t dt$, with $\beta > 0$. We denote by $M^{\emptyset}$ the adoption path such that $M_t = 0$ $\forall t \geq 0$.

Representing by $\mu_0 = p(g) \in (0,1)$ the common prior that $\omega = g$, the posterior, *absent a breakdown*, μ_t, by using Bayes' Rule, is given by:

$$\mu_t = \frac{\mu_0}{\mu_0 + (1-\mu_0)e^{-\int_0^t \beta M_\tau d\tau}}$$

[2]Note that another difference between usual herding models and our setting is that the former has a queue for adoption, with each agent being able to see only the decisions of the others who came before her

As described above, if a realization of the signal happens at a time t, μ_t discontinuously goes to 0, and all agents learn that the state is bad. Given the payoffs described above, no agent ever adopts the innovation from this point forward.

Note that in this setting, as long as a positive mass of agents decides to adopt at any moment in time, $\{\mu_t\}_t$ is strictly increasing in t and $\lim_{t\to\infty}\mu_t = 1$. This is the case as if no bad news arrives, agents become strictly more optimistic about the state of the world, arbitrarily so over time. If no mass adopts the good from time 0 to a time t, then no news can arrive, and we end up with $\mu_t = \mu_0$. Note that even if a positive mass adopts exactly at time t, we still have $\mu_t = \mu_0$, as beliefs can only change over time, so that for any $t' > t$, we have that $\mu_{t'} > \mu_t$.

It is worth pointing out that the above model, with heterogeneous payoffs in the good state and a common prior, can also be modified to represent settings in which the value of the innovation is homogeneous but priors are heterogeneous. Essentially, both capture that individuals have a "threshold" belief above which agents decide to adopt the innovation.[3]

II.A.2 Myopic and Equilibrium Mass of Adopters

We define the mass of adopters absent breakdowns, henceforth also called an *adoption path*, by M_t, as a non-decreasing right-continuous function from $[0,\infty)$ to $[0,1]$.

We will often consider two types of behavior by the agents: 1) they can **myopically** adopt the innovation as soon as they perceive that it is profitable for them to do so when compared to never adopting; 2) they can **strategically** choose the best moment to adopt, given that no breakdowns happened and have knowledge of the payoff distribution of agents.

Denote the right-continuous function $a^i : [t,\infty) \to \{0,1\}$ to be an adoption plan for agent i that is equal to 1 at the time of take-up, depending on adoption up to that point in time. Each collection of plans leads to a *consistent* adoption path $M_t = \int\int_0^t a^i_\tau d\tau d\eta$.

Each adoption path that we will consider induces a time $t^{a^i} \geq 0$ in which the agent adopts absent a breakdown up to this point in time for any path in which there is a positive mass of takers at some past moment. This is without loss of generality from the point above that $\mu_t \to 1$ in that case and the structure of payoffs (0 from not adopting, bounded value v if adopting when the state is bad). We define a value of waiting for agent i, at time t, for strategy a^i, by $V_t^{a^i}$, and any path $\{M_t\} \neq M^{\emptyset}$ so that:

[3] Thresholds beliefs for adoption are widely used in adoption studies. For example, the network adoption literature cited in the Related Literature I section often considers models with this feature.

$$V_t^{a^i} = e^{-r(t^{a^i}-t)}\left(\mu_t v^n - (1-\mu_t)e^{-\int_t^{t^{a^i}} \beta M_\tau d\tau}\right)$$

If agents adopt as soon as they consider it to be better than never taking it up, we say that they are taking up the good myopically.

Definition 1. *An agent i adopts the innovation* ***myopically*** *if she does so at* $t_i^M \equiv \inf\{t | V_t^{a^i} \geq 0\}$

Note that if $V_t^{a^i} < 0$ for every t, agent i would never adopt the good myopically.

A type v^n agent i would therefore myopically adopt at time t_i^M satisfying:

$$\mu_{t_i^M} v^n - (1-\mu_{t_i^M}) = 0$$

When considering agents behaving strategically, we focus on equilibrium outcomes. We will assume that, given a *deterministic* adoption path M_t^E, each agent i must optimally choose an adoption time. The space of all possible a_t^i functions is given by Ω_t^i. Note that given the continuum of agents, to condition behavior on deterministic adoption paths is without loss of generality. Formally, the value for a v^n type agent to waiting at time t is given, then, by $V_t^n = \sup_{a_t^i \in \Omega_t} V_t^{a_t^i}$, and agents adopt as soon as this value is greatest.

The first equilibrium that we are going to consider is the **free supply equilibrium**, in which the principal has no active role in restricting good availability and all players can take up whenever they apply to get the good.

Definition 2. *A* ***free-supply equilibrium*** *is a set of strategy profiles* a_t^i *such that, for every agent i, :*

1. *If* $V_t^n > \mu_t v^n - (1-\mu_t)$, $a_t^i = 0$ *for any i with type* v^n
2. M_t^E *is consistent with* a^i *for each i, so* $M_t^E = \int_0^t \eta(a_t^i) d\eta$.

Each equilibrium is consistent with a unique mass of cumulative adopters M_t^E for an adoption strategy profile.

Lemma 1 from Frick and Ishii (2020), referred to as the "Quasi-single Crossing Property of Equilibrium Incentives" show that the value of waiting can only be higher than the value of adopting now once.

Definition 3. *We say that an economy ξ has the **quasi-single crossing property** if for any free-supply equilibrium adoption path $\{M_t^E\}_t$ and any type v^n, if there is a t such that $V_t^n < \mu_t v^n - (1-\mu_t)$, then $V_{t'}^n < \mu_{t'} v^n - (1-\mu_{t'})$ for any $t' > t$. Similarly, if there is a t such that $V_t^n \leq \mu_t v^n - (1-\mu_t)$, then $V_{t'}^n \leq \mu_{t'} v^n - (1-\mu_{t'})$ for any $t' > t$.*

Proposition 1. *Any economy ξ satisifies the quasi-single property.*

We prove that this holds in our setting in A.1 .

II.B Principal

A principal, who also does not know the state of the world and shares the same prior μ_0 with the agents, wants to reach a target $\bar{M} < 1$ adoption rate as soon as possible. Formally, her flow payoff is given by 1 if $\bar{M}$ is reached, 0 otherwise. If $q_N < 1 - \bar{M}$, the principal can ignore these agents with lower type v^N and reach her target. We can assume, then, without loss of generality, that $\bar{M} > 1 - \min_{i \in I} q_i$, and therefore the principal wants to see agents of all types adopting. Like the agents, the principal is forward-looking and discounts future payoffs by the rate $r > 0$. We will henceforth denote the fundamental variables of our model $\xi = (\mu_0, \beta, r, v^{\mathcal{N}}, q^{\mathcal{N}}, \bar{M})$, which we will call an *economy*.

A few points about the principal's payoff must be emphasized at this point: 1) it does not depend on the state of the world or agents' payoffs, only aggregate adoption 2) it depends on time independently of the probability of reaching the adoption target. This first point means that our principal is **not** a social planner, although we will discuss the welfare implications of optimal supply plans in our coming Discussion section.

Note that a principal can never expect to do better than when all players adopt myopically, which leads to the $\{M_t^O\}_t$ adoption path. This is the case because no agent would adopt before it is myopically profitable to do so in equilibrium and delays by any positive mass of agents can only lead to further delays from others, as then μ_t will increase relatively more slowly.

II.C Homogeneous Valuation

Suppose that agents can apply and get the good at any point in time. This section will consider the case with a single valuation type $v > 0$, so that $N = 1$. This setting is quite similar to the one considered by Frick and Ishii (2020), with the two significant differences being the normalization of the mass of potential adopters and the lack of stochastic opportunities to adopt.

The $\mu_{t^*}v-(1-\mu_{t^*})=0$ condition on myopic take-up leads to adoption either happening for all types at $t=0$ or never, depending on whether v is greater or lower than $v_0^M\equiv(1-\mu_0)/\mu_0$ which can be seen as a threshold value for when agents want to take-up myopically. If $v>v_0^M$, all apply as soon as possible; if $v<v_0^M$, no agent is willing to adopt. If no positive mass of agents adopts, though, beliefs stay the same as the prior, and the game is just the same over time. Therefore, we have that:

$$M_t^O=\begin{cases}1\ \forall t \text{ if } v\geq v_0^M\\ 0\ \forall t \text{ if } v<v_0^M\end{cases}$$

What about strategic take-up? Note first that if $v<v_0^M$, no player wants to take up at $t=0$, beliefs do not change, and $M_t^E=0$. This is also the case in the general model with many times, by the same logic, and therefore if the highest type is such that $v^1<v_0^M$, then $M_t^O=M_t^E=0$ always. One can also see that this is the case by noting that $0\leq M_t^E\leq M_t^O$ always and the above description of M_t^O when $v<v_0^M$.

Take M_t^E as given and assume that $v\geq v_0^M$. One can take a first-order condition on the value of waiting at time t and expecting to pick up at time t^*, given by $V_{t,t^*}^n=e^{-r(t^*-t)}\left(\mu_t v-(1-\mu_t)e^{-\int_t^{t^*}\beta M_\tau d\tau}\right)$ and find that t^* is such that:

$$v=v_0^M e^{-\int_0^t\beta M_\tau^E d\tau}\left(\frac{\beta}{r}M_t^E+1\right)$$

As long as $v<v_0^M\left(\frac{\beta}{r}+1\right)\equiv v_0^E$. We denote the value v_t^E as an agent's minimum threshold payoff to be indifferent between take-up at t or one instant later. If $v\geq v_0^E$, all agents prefer to take up at $t=0$. They are sufficiently optimistic to adopt from the starting time of 0. Note then that the above equation uniquely defines M_t^E. It goes from a positive value up to 1, reached at a time T^E.

One crucial point in our later discussion is that M_t^E increases **convexly** when $v\in(v_0^M,v_0^E)$. This is a consequence of the fact that as μ_t increases over time, only the perspective of greater and greater gains can lead to agents being willing to wait.

Remark 1. *If $v\geq v_0^M$, there is a unique free-supply equilibrium adoption path $\{M_t^E\}_t$, which is i) strictly increasing and ii)convex up to $T^E\equiv\min\{t|M_t^E=1\}$.*

The proof can be found in A.2.

The graphical representation is the following:

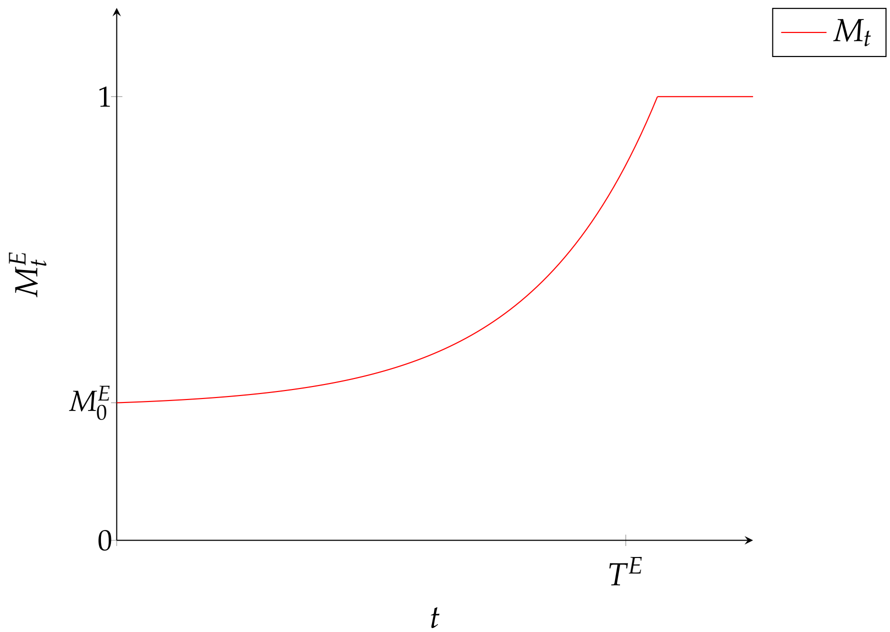

III Results

III.A Adoption-contingent Supply Plans

Suppose that the planner can commit to a supply plan contingent on the take-up that has happened up to each moment of time. Formally, take any weakly increasing function $M:[0,\infty)\rightarrow[0,1]$, and denote the space of all such functions by $\mathcal{M}$. Denote the set of all realized take-up that may have happened up to time t by $\mathcal{M}^t$. At time 0, a planner can choose any function $S_t^A:\mathcal{M}\rightarrow[0,1]$. This function determines how much available supply agents can have at each moment in time. As an example, the principal can set $S_t^A=1$ for any t such that $M_t<\bar{M}$ and $S_t^A=0$ for any $t>t'\equiv\inf\{t|M_t=\bar{M}\}$, so that no more units of the good are available anytime strictly after t'. Agents can adopt, so $a_t^i=1$, only if $S_t^A>0$. If the mass of applicants at a certain point in time is greater than the mass of available supply, the mass of takers is chosen uniformly at random. A final requirement is that we do not want $\{S_t^A\}_t$ to be such that

Definition 4. *A **supply-path** $\{S_t^A\}$ is such that*

1. *For every t, $S_t^A\geq 0$*

2. *Take any convergent sequence $S_{t^k}^A$ such that there is an $\epsilon>0$ satisfying $S_{t^k}^A>\epsilon>0$ for every k, converging to S_*^A. We must have $S_*^A>0$.*

The second condition above means that we are ruling out the possibility of having available supply at every point before or after a time t, but not at time t itself.

For contingent-supply plans, we must redefine an equilibrium for the game. We will name such an equilibrium a **supply-restricted equilibrium**:

Definition 5. *For any supply path $\{S_t^A\}_t$, a **supply-restricted equilibrium** is a set of strategy profiles a_t^i such that, for every agent i, :*

1. *If $V_t^n > \mu_t v^n - (1-\mu_t)$, $a_t^i = 0$ for any i with type v^n*
2. *M_t^E is consistent with a^i for each i, so $M_t^E = \int \int_0^t \eta(a^i(M_\tau)) d\tau d\eta$.*
3. *$M_t \leq S_t^A$ for every $t \geq 0$.*

The difference when comparing it to the free-supply equilibrium comes exclusively from the last bullet point: take-up can happen only if there are units available.

Note that any deviation by a zero mass of agents does not change the adoption path, and therefore the principal cannot change supply based on these types of deviations. If we focus, though, on equilibria in which only positive masses can apply to take up the innovation, the planner can achieve the myopic adoption path by threatening to withhold supply forever whenever the adoption path of agents is at any point different from the myopic adoption path M_t^O. We will refer to this strategy as the **Gim-trigger supply plan**.

Definition 6. *The **Grim-trigger** supply plan sets supply $S_t^A = 0$ for any $t > t^*$ and $S_t^A = 1$ for any $t \leq t^*$, where $t^* \equiv \inf\{t \geq 0 | M_t^E < M_t^O\}$.*

A unique induced equilibrium adoption path exists for any given supply plan $S^A(.)$. We say that a plan **induces** an adoption path M_t if this is the unique equilibrium adoption path given it. Formally:

Definition 7. *A supply plan $\{S_t^A\}_t$ induces an equilibrium take-up path $\{M_t^E\}_t$ if there is a **supply-restricted equilibrium** associated with that path.*

Proposition 2. *Any supply plan $\{S_t^A\}_t$ induces a unique supply-restricted equilibrium.*

The proof can be found in A.3.

We also have that the principal can induce $\{M_t^O\}_t$ by using contingent supply plans, proved in A.4:

Remark 2. *For any economy ξ, the principal can induce M_t^O with the Grim-trigger supply plan.*

This means the principal can reach her adoption target at her first-best time when using contingent plans as above. It is worth noting that this ability does not depend on the number of types in v^I, unlike some results that we will present later on.

III.B Non-contingent Supply Plans

Given our remark above on the power of take-up contingent supply plans, we will focus on our primary setting: when the principal can commit to realizing extra supply at particular times but not dependent on take-up. This means that she cannot add supply when demand exceeds available mass for the good or when the take-up rate is low.

Formally, instead of supply plans being functions of observed adoption paths, they will be functions of times. Define the incremental mass of supply at time t by $S_t:[0,\infty)\to[0,1]$. Therefore the available mass of goods is given by $\{S^A\}_t$ with $S_t^A=\int_0^t S_\tau d\tau - M_t$, the integral of all increments after accounting for all mass already claimed.

Suppose that the principal was to set $S_0=\bar{M}<1$. What would be the induced equilibrium take up path M_t^E?

First, note that we must have a positive mass of applicants in finite time so that $M_t>0$ for some $t<\infty$. Otherwise, any agent would want to apply and get the good at time $t=0$.

Suppose now that $t^*=\inf\{t|M_t^E>0\}>0$. Then any agent i that applies at t^* has the same belief $\mu_{t^*}=\mu_0$ but, given discounting, a lower payoff. Therefore any agent would rather deviate and apply at $t=0$ instead.

The argument above establishes that any equilibrium take-up path must be such that $M_0^E>0$. This means that we must have $T^E=\inf\{t|M_t^E=\bar{M}\}<\infty$ so that all will eventually get the good.

Consider now the mass of applicants at T^E. If $A_{T^E}>S_{T^E}^A$, there is a lottery to choose when to get the good. But then any such agent would instead apply and get the good $\epsilon>0$ earlier. If, instead, we have smooth take-up up to T^E, given that $\bar{M}<1$, there is an agent i who would instead get it at time T^E but is unable to. Therefore we cannot have $T^E>0$ and all apply to get the good at $t=0$.

Theorem 1. *For any $\bar{M}<1$. If $v>v_0^M$, $T^E(\bar{M})=0$*

The proof can be found in A.5.

This means that *arbitrary* amounts of scarcity can lead to myopic take-up behavior.

III.C Two Types

Suppose now instead that there are two types $v^1 > v^2$, so that $I = \{1,2\}$.

Firstly, let's rule out some uninteresting cases:

1. As noted in the previous section, if $v^1 < v_0^M$ the only equilibrium take-up path is $M_t^E = 0$ for any $t \geq 0$.
2. If $v_2 > v_0^M$, both types of agents are willing to adopt at time 0 myopically, and the principal can use arbitrary scarcity to have all apply at $t = 0$, as in the one-type case.

Suppose instead, then, that $v^1 > v_0^M > v^2$. This means high-types are willing to take up the good myopically, and the low-types are not. Remember that we assume that $\bar{M} > 1 - \min_{i \in I}$ so the principal does not want to ignore a type to reach her target adoption rate.

Define by $T_1^D(m_1)$ the time a v^1 type is indifferent between adopting at this time with a mass m_1 of agents adopting at time 0 and no additional mass adopted between 0 and $T_1^D(m_1)$. Define also $T_2^M(m_1)$ to be the time in which any v^2 types is willing to adopt myopically when a mas m_1 of agents adopting at time 0 and no additional mass adopted between 0 and $T_2^M(m_1)$.

Suppose first that $T_1^D(q_1) < T_2^M(q_1)$. Then the principal can use arbitrary scarcity at time 0 and set $S_0 = q_1 - \epsilon$ for low ϵ. By the argument above, all v^1 types will apply at $t = 0$. But if $S_{T_2^M(q_1)} = q_2 - \epsilon_2$, all v^2 types will also apply to get the good myopically.

The plan above is arbitrarily close to myopic take-up, so it is optimal for the principal in that the adoption target is hit as soon as possible.

Suppose now instead that $T_1^D(q_1) > T_2^M(q_1)$. The planner is no longer able to reach the adoption at time $T_2^M(q_1)$, as then any v^1 type would rather apply to get the good at this time instead of $T_2^M(q_1)$ (for high enough $\bar{M}$. Note though that $T_1^D(m_1)$ is increasing in m_1 (more learning is available creates incentives to delay take-up longer), but $T_2^M(m_1)$ is *decreasing* in m_1, as less learning has happened to induce the lower types to take-up myopically.

Note that $T_1^D(0)=0$, so there is a unique m_1^* such that $T_1^D(m_1^*)=T_2^M(m_1^*)$. The principal can then set $S_0=m_1^*-\epsilon$ and then serve the remainder (up to $\bar{M}$) at $T_2^M(m_1^*)$.

Note that this plan is optimal; any optimal plan must have v^2 types taking up myopically. Given that, we need to have v^1 types preferring to take up at $t=0$ then with the second batch.

Note also that we must have $M_0^E>0$.

The plan above is simple, as it only has two batches of supply made available. One can wonder whether other plans (potentially with more batches or even continuous supply) can improve upon it. The answer is negative. To get the intuition, suppose that some plan is strictly better, with $T^{E'}<T^E$. As v^2 types take up the good myopically, the amount of learning happening (governed by the integral of M_t^E up to that point) is the same. But then all v^1 types would prefer taking up at this moment then at $t=0$, which contradicts what was said above.

The following theorem summarizes the discussion in the previous paragraph (proved in A.6:

Theorem 2. *The optimal supply plan when $I=\{1,2\}$ and $v^1>v_0^M>v^2$ consists of using two batches at times 0 and t_2, $m_2=\bar{M}-m_1$ and $m_1=q_1-\epsilon$ and if $T_1^D(q_1)<T_2^M(q_1)$ or $m_1=m_1^*-\epsilon$ for $T_1^D(m_1^*)=T_2^M(m_1^*)$.*

Graphical representation:

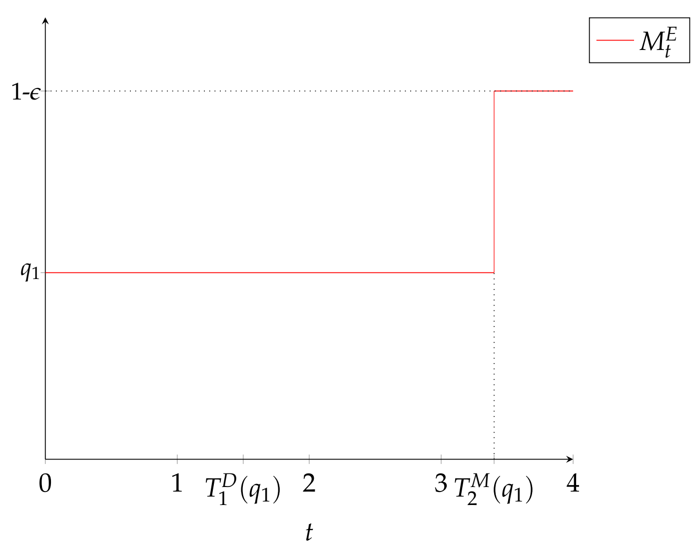

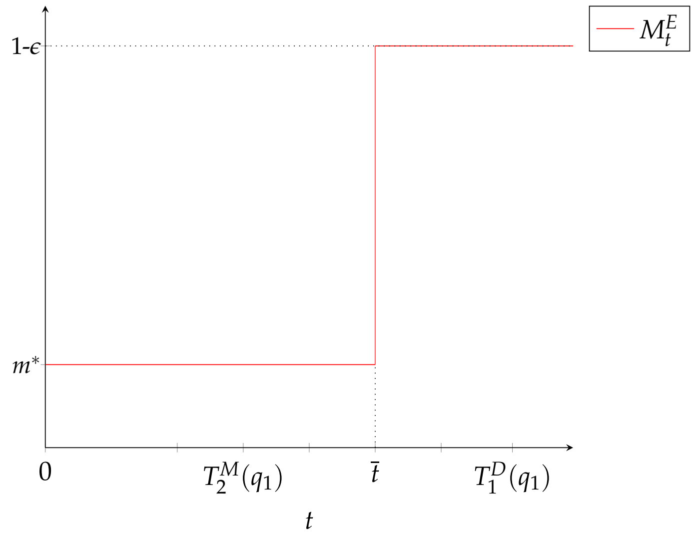

III.D Three Types

Suppose now that there are three possible types $v^1 > v^2 > v^3$. As before, if $v^1 < v_0^M$, no agent would ever adopt, and if $v^3 > v_0^M$, the principal can simply use an arbitrary amount of scarcity to reach her target at time 0. This leaves us with the following cases:

1. $v^1 > v_0^M > v^2 > v^3$
2. $v^1 > v^2 > v_0^M$

We will break down the optimal supply algorithm for the first case. It is not without loss of generality to do so, but the steps for the second one are very much analogous.

The first point in our analysis is to point out that the optimal supply plan must have at most 3 batches:

Lemma 1. *If $v^{\mathcal{N}} = (v^1, v^2, v^3)$, there is an optimal supply plan with at most 3 points $B \equiv (t_1, t_2, t_3)$ with $S_t > 0$ for $t \in B$.*

The proof can be found in A.7, and the intuition goes as follows: a positive mass of agents must adopt at time $t = 0$, or else nothing will change. The only agents willing to do that (for any plan) are v^1 types. Any optimal plan must have all v^3 type agents adopting myopically (in a last batch with less supply than willing adopters). The three batches are to be released at three points in time, one at 0, one at a final moment for myopic take-up of the lowest types, and finally, one in which any v^2 type first adopts. Can any batch in between each of these types strictly help the principal? No, because if it could, it would lead to higher beliefs at an earlier time (whether for v^1 if between t_1 or t_2 or for v^2 if between t_2 or t_3) and therefore all types taking-up the good earlier would rather do so later on.

Before going through the algorithm itself, let's first define some important variables:

1. $T_1^D(m_1)$ is the time in which v^1 is indifferent between 0 and this time when m_1 take it at first. So $0 \sim^{v^1} T_1^D(m_1)$, or

$$v^1 - v_0^M = e^{-rT_1^D(m_1)}(v^1 - v_0^M e^{-\beta m_1 T_1^D(m_1)})$$

2. $T_2^E(m_1)$ is the preferred time for a v^2-type agent to adopt up given that a m_1 mass has done so at $t=0$ and no new mass has done it after that.

3. $T_3^M(m_1, m_2, t_2)$ is the type in which the v^3 type is indifferent between adopting now or never again, given that a mass m_1 decided to do so at time 0 and an extra mass m_2 at time t_2.

There is a substantial difference between allowing for an arbitrarily small amount of scarcity and not. We will consider first a limiting setting in which we have $\epsilon \to 0$ scarcity but people behave as if there was positive scarcity, solve for the model in this world and then show that we can get arbitrarily close to it with positive ϵ amount of total scarcity.

Definition 8. *A limiting economy ε^L has zero-mass of agents competing excluded from take-up from the good*

We can now go to the full three types algorithm for a limiting economy:

1. Compare $T_1^D(q_1)$ and $T_3^M(q_1, 0, 0)$. If the former is greater than the latter, batch up to the point m_1^* in which the v^1-types are indifferent. Otherwise, continue.

2. Check if $T_2^E(q_1) > T_3^M(q_1, 0, 0)$. If so, release batches at times 0 and $T_3^M(q_1, 0, 0)$, with $m_1 = q_1$ and $m_2 = \bar{M} - q_1$. Otherwise, continue.

3. Check if $T_1^D(q_1) \prec^{v^2} T_3^M(0, q_1)$. If so, no batching for the v^2 types is profitable. Otherwise, proceed.

4. If $T_3^M(q_2, q_1) \succ^{v^2} T_1^D(q_1) \succ^{v^2} T_3^M(0, q_1)$, there are two options: If $T_2^E(q_1) > T_1^D(q_1)$, go for the former. Otherwise, go for m_2^* at T_1^D making v^2 indifferent.

5. If $T_1^D(q_1) \succ^{v^2} T_3^M(q_2, q_1)$, then pick m_1^* such that $T_1^D(m_1^*) \sim^{v^2} T_M^3(q_2, m_1^*)$. Compare this to $T_2^E(q_1)$. One of these two is optimal.

Here is the intuitive reasoning for each point:

1. Suppose that adding a mid-batch is better for the principal. Then by definition, it will lead to a higher payoff to types v^1 than before, as it implies the same amount of learning happening earlier. But then no v^1 type would want to take up at time 0.

2. If the condition is met, then $T_3^M(q_2,q_1) \succ^{v^2} T_1^D(q_1) \succ^{v^2} T_3^M(0,q_1)$. Adding a new batch with positive mass take-up would lead to an even more relatively desirable $T_3^M(q_2,q_1)$, though—contradiction with take-up at the new batch.

3. Suppose that any batch for v^2 can make things better. This means the final batch will come even earlier, with the same amount of learning. But then v^2 must strictly prefer taking up at this last batch, a contradiction.

4. A second batch with q_2 at $T_1^D(q_1)$ is not feasible (they would all prefer to take up with v^3). Bringing more v^2 types later rather than sooner is always better. However, one should never decrease m_1 to get more v^2 earlier.

5. This comes from the concavity of the objective function being minimized between these two points and the fact that the minimum cannot be after $T_2^E(q_1)$ or reversing the preference relation.

We note first that the restriction on three batches is not vacuous, and two batches are not enough in some cases, as in our example, that can be found on section B of the Appendix. We name the algorithm above the **three-type optimal algorithm** (TTOA) and the supply plan induced by it $\{S^{TTOA}\}$. With these two concepts, we state our main theorem (proved in A.8:

Theorem 3. *If $v^{\mathcal{N}} = (v^1, v^2, v^3)$ and $v^1 > v_0^M > v^2 > v^3$, then $\{S^{TTOA}\}$ induces the optimal supply plan for any limiting economy ε^L.*

Define by T^L the time $M_t^E = \bar{M}$ for the limiting economy. Then, for any $\epsilon > 0$, the principal can reach its target at $T^L + \epsilon$ by perturbing the induced supply path by TTOA. This is stated in the following proposition and proven in the Appendix.

Definition 9. *Take a supply plan $\{S_t\}$. A γ-perturbation $\{S_t^\gamma\}$ has $S_t' > S_t - \gamma$ for every $t \geq 0$.*

Proposition 3. *For any $\epsilon > 0$, the principal can induce $T^* = T^L + \epsilon$ by using a γ-perturbation of $\{S^{TTOA}\}$ for a sufficiently small γ.*

The intuition for this result comes from the fact that the arbitrarily small scarcity leads to early take-up and behavior that is very similar to the one from the limiting economy ε^L.

IV Discussion and Extensions

IV.A Discussion

The results from the previous section give us insights into how can a planner achieve a target adoption rate as soon as possible. As discussed, this objective is common for a diverse array of problems, such as reaching immunization rates, dealing with pests, or reaching a critical mass of adopters for a matching platform. Given the distribution of valuations, the principal knows when to add new batches of the good and how many units to use.

As discussed in our literature review, there are other reasons why a principal may want to restrict supply (notably for persuasion purposes), but the question we pose here might add to the arsenal of tools available for public policymakers that need to deal with informational free-riding.

IV.B Extensions

In this section, we discuss some extensions of particular interest.

IV.B.1 Principal's Preference for Adoption Paths

If the objective of the planner is not to reach a target $\bar{M}$ as soon as possible but rather to get a preferred adoption path M_t, with strict preferences for paths that are weakly above at each point in time, then by definition scarcity, by either supplying fewer units or later, can never be desirable. Note that the myopic take-up path M_t^O is strictly preferred to any strategic one M_t^E. Define an *optimal supply path* as one that is not strictly preferred by any other. The hands-free supply path is optimal, then.

IV.B.2 Social Planner

There is a clear gap between the interests of the principal and that of agents in our model. In particular, the former wants take-up to happen no matter what, and a social planner would rather have no adoption if $\omega = b$.

Note that in the model, welfare does not increase with faster-take-up: players take up myopically and might be more rapidly adopting a bad product.

IV.B.3 Uniformly Distributed Values

One question of interest is to analyze the limits of supply restrictions when the number of types increases. In particular, should one batch a finite number of types in this case?

We will focus here on the case with $v \sim U[\underline{v}, \bar{v}]$. Myopic behavior is then governed by the following:

$$\dot{v}_t^M = -\beta \frac{\bar{v} - v_t^M}{\bar{v} - \underline{v}} v_t^M$$

With $v_0^M = \frac{1-\mu_0}{\mu_0}$.

One can solve the differential equation and find that the solution is of the form:

$$v_t^M = \frac{k_2 \bar{v}}{k_2 + e^{t\bar{v}k_1}}$$

where k_1 and k_2 are constants:

$$k_1 = \frac{\beta}{\bar{v} - \underline{v}}$$

and k_2 is given by the initial condition:

$$k_2 = \frac{1-\mu_0}{\bar{v}\mu_0 - (1-\mu_0)} \tag{1}$$

1. v_t^M is decreasing
2. If $v_t^M \geq \bar{v}/2$, v_t^M is concave decreasing, and if $v_t^M \leq \bar{v}/2$, v_t^M is convex decreasing. To see that, just note that

   $$\dot{v}_t^M = -k_1(\bar{v} - v_t^M) v_t^M$$

   and therefore:

   $$\ddot{v}_t^M = -k_1 \dot{v}_t^M (\bar{v} - 2v_t^M)$$

   As $k_1 \dot{v}_t^M = -k_1^2(\bar{v} - v_t^M) v_t^M < 0$, we have our result.
3. The inequalities from the previous point imply the opposite for the M_t^O: if v_0^M is high enough, it starts out convex increasing, but eventually it becomes concave increasing.
4. The inflection point of M_t^O is given by t^I such that $v_{t^I}^M = \bar{v}/2$. Note that it might not be reached if v_0^M starts below $\bar{v}/2$ (people already start out too optimistic), or if $\underline{v} > \bar{v}/2$, in which case take-up always happens convexly.

We conclude that M_t^O, the myopic adoption path, has an S-shaped adoption form for optimistic enough agents.

Strategic behavior is governed by the following differential equation:
If take-up happens in equilibrium, we must have:

$$v_t^E = \frac{1-\mu_t}{\mu_t}\left(\frac{\beta}{r}M_t^E + 1\right)$$

For our uniform example, this implies:

$$v_t^E = \frac{1-\mu_t}{\mu_t}\left(\frac{\beta}{r}M_t^E + 1\right)$$

This leads to:

$$v_t^E = \frac{1-\mu_0}{\mu_0}e^{-\beta\int_0^t M_\tau^E d\tau}\left(\frac{\beta}{r}M_t^E + 1\right)$$

Using that $v_0^M = (1-\mu_0)/\mu_0$, we have:

$$v_t^E = v_0^M e^{-\beta\int_0^t M_\tau^E d\tau}\left(\frac{\beta}{r}M_t^E + 1\right)$$

Differentiating both sides with respect to t, we get:

$$\dot{v}_t^E = -v_0^M \beta M_t^E e^{-\beta\int_0^t M_\tau^E d\tau}\left(\frac{\beta}{r}M_t^E + 1\right) + v_0^M e^{-\beta\int_0^t M_\tau^E d\tau}\frac{\beta}{r}\dot{M}_t^E$$

so that:

$$\dot{v}_t^E = -\beta M_t^E v_t^E + v_0^M e^{-\beta\int_0^t M_\tau^E d\tau}\frac{\beta}{r}\dot{M}_t^E$$

which leads to:

$$\dot{v}_t^E = -\beta M_t^E v_t^E + \frac{v_t^E}{\left(\frac{\beta}{r}M_t^E + 1\right)}\frac{\beta}{r}\dot{M}_t^E$$

As $M_t^E = \frac{\bar{v}-v_t^E}{\bar{v}-\underline{v}}$,

$$\dot{M}_t^E = -\frac{\dot{v}_t^E}{\bar{v}-\underline{v}}$$

we have that:

$$\dot{v}_t^E = -\beta M_t^E v_t^E - \frac{v_t^E}{\left(\frac{\beta}{r} M_t^E + 1\right)} \frac{\beta}{r} \frac{\dot{v}_t^E}{\bar{v} - \underline{v}}$$

Denoting:

$$k = \frac{\beta}{r(\bar{v} - \underline{v})}$$

we have:

$$\dot{v}_t^E = -\beta M_t^E v_t^E - \frac{v_t^E}{\left(\frac{\beta}{r} M_t^E + 1\right)} k \dot{v}_t^E$$

So that:

$$\dot{v}_t^E = -\frac{\left(\frac{\beta}{r} M_t^E + 1\right) \beta M_t^E v_t^E}{v_t^E k + 1} < 0$$

Unlike the myopic adoption path, one cannot find a closed-form solution to the equilibrium adoption path. However, we would like to remark that for high enough μ_0, M_t^E must start increasing convexly.

The following figure plots the adoption paths M_t^O (green) and M_t^E (blue) for particular values $\mu_0 = 0.5$, $\beta = r = 1$, and $v \sim U[0.2, 1.2]$:

Figure 4: S-shaped adoption curves for uniformly distributed valuations

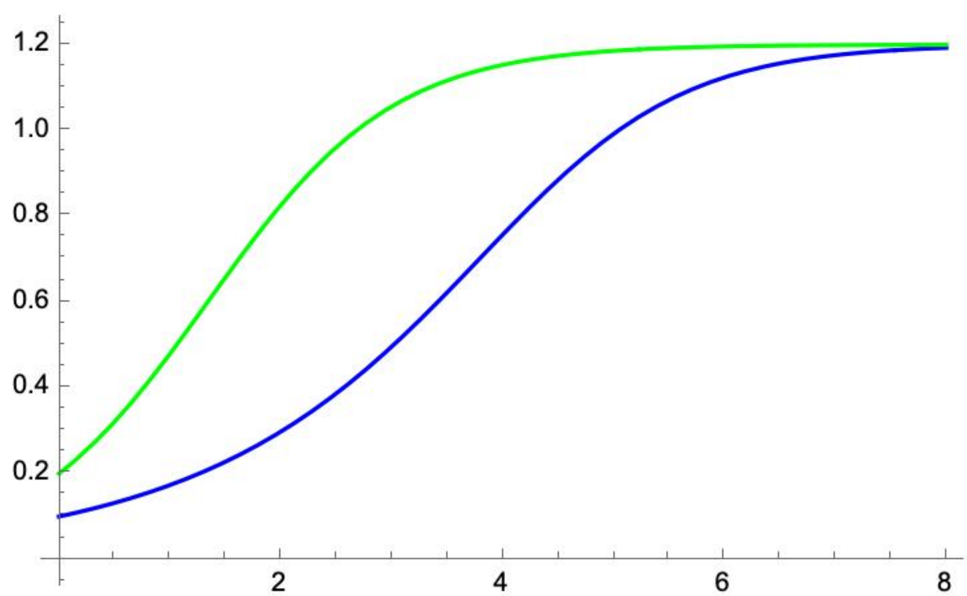

IV.C Different Regions

We finally briefly discuss what would happen if two regions with different arrival rates of bad news ($\beta_1 > \beta_2$), but with such breakdown visible for both, and the same payoff distributions and other parameters. The principal need to have adoption from all types in both regions and can restrict supply in any particular one.

Despite the potential advantage of the region with the higher beta, the target adoption must be reached *simultaneously* in both regions. This is the case because we need to have the lowest types in both regions adopting myopically, and they share not only a prior but also beliefs.

This suggests that areas with less access to health care (such as rural areas), and in which the assessment of side effects of a vaccine could be slower to get, should not be left to have skeptics taking later on, if one needs to reach adoption targets fast, as is the case for any rapidly mutating viral disease, as COVID-19.

References

Akbarpour, Mohammad and Matthew O Jackson. 2018. "Diffusion in networks and the virtue of burstiness." *Proceedings of the National Academy of Sciences* 115 (30):E6996–E7004.

Bandiera, Oriana and Imran Rasul. 2006. "Social networks and technology adoption in northern Mozambique." *The economic journal* 116 (514):869–902.

Banerjee, Abhijit V. 1992. "A simple model of herd behavior." *The quarterly journal of economics* 107 (3):797–817.

Besley, Timothy and Anne Case. 1993. "Modeling technology adoption in developing countries." *The American economic review* 83 (2):396–402.

Besley, Timothy, Anne Case et al. 1994. "Diffusion as a learning process: Evidence from HYV cotton." Tech. rep.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of political Economy* 100 (5):992–1026.

Bonatti, Alessandro. 2011. "Menu pricing and learning." *American Economic Journal: Microeconomics* 3 (3):124–63.

Che, Yeon-Koo and Konrad Mierendorff. 2019. "Optimal dynamic allocation of attention." *American Economic Review* 109 (8):2993–3029.

Conley, Timothy G and Christopher R Udry. 2010. "Learning about a new technology: Pineapple in Ghana." *American economic review* 100 (1):35–69.

DeGraba, Patrick. 1995. "Buying frenzies and seller-induced excess demand." *The RAND Journal of Economics* :331–342.

Dupas, Pascaline. 2014. "Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment." *Econometrica* 82 (1):197–228.

Evans, David S and Richard Schmalensee. 2016. *Matchmakers: The new economics of multisided platforms*. Harvard Business Review Press.

Eyster, Erik and Matthew Rabin. 2014. "Extensive imitation is irrational and harmful." *The Quarterly Journal of Economics* 129 (4):1861–1898.

Foster, Andrew D and Mark R Rosenzweig. 1995. "Learning by doing and learning from others: Human capital and technical change in agriculture." *Journal of political Economy* 103 (6):1176–1209.

Frick, Mira and Yuhta Ishii. 2020. "Innovation adoption by forward-looking social learners." Tech. rep., Yale.

Golub, Benjamin and Evan Sadler. 2017. "Learning in social networks." *Available at SSRN 2919146* .

Jackson, Matthew O. 2010. *Social and economic networks*. Princeton university press.

Keller, Godfrey and Sven Rady. 2015. "Breakdowns." *Theoretical Economics* 10 (1):175–202.

Kremer, Michael and Edward Miguel. 2007. "The illusion of sustainability." *The Quarterly journal of economics* 122 (3):1007–1065.

Laiho, Tuomas and Julia Salmi. 2018. "Social Learning and Monopoly Pricing with Forward Looking Buyers." Tech. rep., Working paper.

Möller, Marc and Makoto Watanabe. 2010. "Advance purchase discounts versus clearance sales." *The Economic Journal* 120 (547):1125–1148.

Mossel, Elchanan, Allan Sly, and Omer Tamuz. 2015. "Strategic learning and the topology of social networks." *Econometrica* 83 (5):1755–1794.

Munshi, Kaivan. 2004. "Social learning in a heterogeneous population: technology diffusion in the Indian Green Revolution." *Journal of development Economics* 73 (1):185–213.

Nocke, Volker and Martin Peitz. 2007. "A theory of clearance sales." *The Economic Journal* 117 (522):964–990.

Parakhonyak, Alexei and Nick Vikander. 2023. "Information design through scarcity and social learning." *Journal of Economic Theory* 207:105586.

Perego, Jacopo and Sevgi Yuksel. 2016. "Searching for Information and the Diffusion of Knowledge." *Unpublished manuscript, New York University* .

Reeves, T, H Ohtsuki, and S Fukui. 2017. "Asymmetric public goods game cooperation through pest control." *Journal of Theoretical Biology* 435:238–247.

Rogers Everett, M. 2003a. *Diffusion of Innovations*. New York: Free Press.

———. 2003b. "Diffusion of innovations." *New York* 12.

Ryan, Bryce and Neal C Gross. 1943. "The diffusion of hybrid seed corn in two Iowa communities." *Rural sociology* 8 (1):15.

Scharfstein, David S and Jeremy C Stein. 1990. "Herd behavior and investment." *The American economic review* :465–479.

Smith, Lones and Peter Sørensen. 2000. "Pathological outcomes of observational learning." *Econometrica* 68 (2):371–398.

Smith, Lones, Peter Norman Sørensen, and Jianrong Tian. 2021. "Informational herding, optimal experimentation, and contrarianism." *The Review of Economic Studies* 88 (5):2527–2554.

Wolitzky, Alexander. 2018. "Learning from Others' Outcomes." *American Economic Review* 108 (10):2763–2801.

Young, H Peyton. 2009. "Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning." *American economic review* 99 (5):1899–1924.

Appendix: For Online Publication

Ricardo Fonseca

A Appendix

This document presents proofs omitted in the main text and additional theoretical results.

A.1 Proposition 1

Before going through the proof of the proposition itself, we mention some preliminary results:

(i) V_t^n is continuous in t for any equilibrium adoption path $\{M_t^E\}_t$.

(ii) If $\lim_{t\to\infty} V_t^n < lim_{t\to\infty}\{\mu_t v^n - (1-\mu_t)\}$ if $\{M_t^E\}_t \neq M^{\emptyset}$.

Proofs:

(i)
$$V_t^n = \int_t^{\infty} e^{-(r(s-t)}\frac{\mu_t}{\mu_s}\max\{\mu_t v^n - (1-\mu_t), V_s^n\} ds$$

Which is clearly continuous in t.

(ii) If $\{M_t^E\}_t \neq M^{\emptyset}$, then $\mu_t \to 1$, and therefore $lim_{t\to\infty}\{\mu_t v^n - (1-\mu_t)\} = v^n > \lim_{t\to\infty} V_t^n$. The last inequality comes from $v^n > V_t^n$ for every t and that as μ_t is arbitrarily close to 1, only the time discount matters.

We must show that any economy ξ satisfies the quasi-single crossing property, as laid out in Definition 3. We will proceed by contradiction:

1. Suppose first that for some time t and type v^n, $V_t^n < \mu_t v^n - (1-\mu_t)$ and there is a time $t' > t$ such that $V_{t'}^n \geq \mu_{t'} v^n - (1-\mu_{t'})$.
2. Suppose now that for some time t and type v^n, $V_t^n = \mu_t v^n - (1-\mu_t)$ and there is a time $t' > t$ such that $V_{t'}^n \geq \mu_{t'} v^n - (1-\mu_{t'})$.

To see why the first statement cannot hold, note that if M_t^E is continuous at t, V_t^n is differentiable at this point. Otherwise, it is still right-differentiable. Also, as V_t^n is continuous in t, there is an $\epsilon > 0$ such that $V_{t+\epsilon}^n < \mu_{t+\epsilon} v^n - (1-\mu_{t+\epsilon})$ We can write, for any $\tau \in (t, t+\epsilon)$:

$$V_\tau^n = \int_\tau^{t+\epsilon} e^{-r(s-\tau)}[\mu_\tau v^n - (1-\mu_\tau)e^{-\int_\tau^s \beta M_x^E dx}]ds + e^{-r(t+\epsilon-\tau)}[\mu_\tau v^n - (1-\mu_\tau)e^{-\int_\tau^{t+\epsilon}\beta M_x^E dx}]V_{t+\epsilon}^n$$

Using Ito's lemma, we find that the right derivative is given by:

$$\dot{V}_\tau^n = (r+(1-\mu_\tau)\beta M_\tau^E)V_\tau^n - (\mu_\tau v^n - (1-\mu_\tau))$$

The difference between $\mu_t v^n - (1-\mu_t)$ and V_t^n has a right derivative equal to $\beta M_t^E \mu_t(1-\mu_t) - (r + (1-\mu_\tau)\beta M_\tau^E)V_\tau^n - (\mu_\tau v^n - (1-\mu_\tau))$.

As we know $\lim_t V_t^n < \lim_t(\mu_t v^n - (1-\mu_t)$, therefore one can define moments $\underline{t} < t'$ such that $V_{\underline{t}}^n = \mu_{\underline{t}} v^n - (1-\mu_{\underline{t}})$ and analogously a $\bar{t} > t'$ such that $V_{\bar{t}}^n = \mu_{\bar{t}} v^n - (1-\mu_{\bar{t}})$.

Note that the right derivative of $\mu_t v^n - (1-\mu_t)$ at points $\underline{t}$ and $\bar{t}$ are given by $r(\mu_{\underline{t}} v^n - (1-\mu_{\underline{t}})) - \beta M_{\underline{t}}^E(1-\mu_{\underline{t}})$ and $r(\mu_{\bar{t}} v^n - (1-\mu_{\bar{t}})) - \beta M_{\bar{t}}^E(1-\mu_{\bar{t}})$. As the first must be non-negative and the second non-positive, together with $\mu_{\underline{t}} \leq \mu_{\bar{t}}$, we need to have $\underline{t} = \bar{t} \equiv t^*$.

Therefore, for $\tau < t^*$, we get that $r(\mu_t au v^n - (1-\mu_\tau)) \leq \beta M_\tau^E(1-\mu_\tau)$. But then, together with the fact that $V_\tau^n < \mu_\tau v^n - (1-\mu_t au)$, we get that

$$0 < -rV_\tau^n \mu_\tau + (\mu_\tau v^n - (1-\mu_\tau) - V_\tau^n)\beta M_\tau^E(1-\mu_\tau)$$

As the right-hand side is the right-derivative of $\mu_\tau v^n - (1-\mu_\tau) - V_\tau^n$, we have that this value is increasing, and therefore we cannot have $V_{t^*}^n = \mu_{t^*} v^n - (1-\mu_{t^*})$. A contradiction and conclusion to the proof.

The second statement can be proved analogously.

A.2 Remark 1

The proof that the equilibrium adoption path $\{M_t^E\}_t$ for an economy ξ has these four properties: existence, uniqueness, the fact that it is strictly increasing over time, and convex for a homogeneous valuation economy follows directly from Theorem 1 from Frick and Ishii (2020). The only difference between our setup and theirs is that we do not have stochastic adoption opportunities. However, the proof that the equilibrium path is convex, exists, and is unique goes exactly the same way.

To see that it is strictly increasing, suppose the path is constant in some interval $[t_1, t_2]$ with $t_2 > t_1$

and $M^E_{t_2} < 1$. Then, it must be the case that for some individuals waiting at a time $t \in (t_1, t_2)$ is at least as good as adopting at time t_2. By homogeneity of the value of adoption in the good state, though, and the fact that $M^E_{t_1} > 0$, we must have $V^n_{t_1} \leq \mu_{t_1} v^n - (1-\mu_{t_1})$. But then, by Proposition 1, we must have $V^n_t < \mu_t v^n - (1-\mu_t)$, contradicting our assumption that waiting at t is preferred.

A.3 Proposition 2

We will proceed in two parts: i) existence and ii) uniqueness.

i) Existence If $v^1 < v^M_0 \equiv \frac{1-\mu_0}{\mu_0}$, then $M^{\emptyset}$ is a supply-restricted equilibrium, with all types optimally choosing not to adopt, $a^i_t = 0$ every period. This is clearly consistent with the three requirements for a supply-restricted equilibrium.

If $v^1 \geq v^M_0$, then for any $t^F \equiv \min\{t | S^A_t > 0\}$ know that $M^E_{t^F} > 0$. Note that t^F is well-defined because of the second property of a supply path (see Definition 4). Given that, we know that $\lim_t \mu_t = 1$ and, therefore, eventually $a^n_t = 1$ is optimal for any type v^n. Given a supply plan $\{S^A_t\}_t$ and the structure of the game, there is an optimal time to adopt absent breakdown t^n_* for each type, and therefore a Nash equilibrium for the game.

ii) Uniqueness By the definition of a supply-restricted equilibrium, every point t is such that either M^E_t is continuous, and there is indifference for some type v^n, or there is a jump in M^E_t. With indifference, the unique adoption path must satisfy the following:

$$M^E_t = \frac{r(\mu_t v^n - (1-\mu_t))}{\beta(1-\mu_t)}$$

This means that a **unique adoption path** is compatible with indifference for a type v^n. During jumps, there is obviously only one adoption decision compatible with the equilibrium path. As these are the two possible cases, we conclude that any supply plan induces a unique adoption path.

A.4 Remark 2

We must show that a Grim-trigger supply plan will induce $\{M^O_t\}$. As there is exactly one supply-restricted equilibrium, we need only to prove M^O_t is an equilibrium adoption path.

Given $\{M^O_t\}_t$, suppose a profitable deviation exists for a positive mass of agents of type v^n at time t'. By commitment power, the planner can commit to $S^A_t = 0$ for any $t > t'$. The payoff for any such player is equal to 0 by this deviation. By following the strategy, though, the player receives $\mu_{t'} v^n - (1-\mu_{t'}) \geq 0$, where the last inequality comes from the definition of $\{M^O_t\}_t$, which contradicts the fact that it is a profitable deviation.

A.5 Theorem 1

We need only to show that $V_0 > \mu_0 v - (1-\mu_0)$. To see that, note first that $M^{\emptyset}$ is not an equilibrium, by the assumption that $v > v_0^M$. Therefore $\lim_t \mu_t = 1$ and there is a time T in which all agents strictly prefer to adopt $a_T = 1$. Take T to be the moment any agent last gets to take up the good. There are two options as to what happens at T if it is strictly greater than 0:

(i) $\eta(i|a_T^i = 1) > S_T^A$.

As described, a lottery will happen at T, and a fraction of agents will get the good. But then there is an $\epsilon > 0$, so it is better for these players to apply at $T - \epsilon$. To see that, not that the payoff at T is given by:

$$\frac{\eta(i|a_T^i=1)}{\bar{M} - M_T^E} e^{-rT} (\mu_T v - (1-\mu_T)) e^{-\int_0^T \beta M_\tau^E d\tau}$$

Which is strictly lower than the payoff from applying at $T - \epsilon$, by continuity of μ_t:

$$e^{-rT} (\mu_{T-\epsilon} v - (1-\mu_{T-\epsilon})) e^{-\int_0^{T-\epsilon} \beta M_\tau^E d\tau}$$

(ii) Otherwise, $\eta(i|a_T^i = 1) \leq S_T^A$. If the inequality is strict, by the definition of T, some agent decides never to apply for available units of the good, even though it is profitable to do so at time T. If it holds with equality, as $\bar{M} < 1$, some agents never get the good, and get a payoff of 0. However, by $v > v_0^M$, applying at time 0 is profitable: as $T > 0$, the good is available at time 0.

Given that these two cases contradict equilibrium behavior, we conclude that we must have $T = 0$, and all agents apply at time 0.

A.6 Theorem 2

The argument for why up to two batches are enough follows the proof on Lemma 1 and will therefore be omitted here.

We must show that the supply plan described is indeed optimal given the family of up to two batches.

Note first that the time the first batch is released must be 0. To see that, note that if this first t_1 is strictly greater than 0, the game is the same from 0 to t_1, and the principal is strictly worse-off.

There are, then, three variables to choose from:

1. How many units to release at time 0, m_1
2. When to release the second batch, t_2
3. How many units to release at the second batch, m_2

One can see that m_2 can easily be derived by $\bar{M} - m_1$. We need to determine, then, t_2 and m_1.

Note also that it is optimal for the principal to have the type v^2 agents to adopt *myopically*. This is the case because they cannot adopt before that, and anytime after is just decreasing the payoff for the principal. Therefore we establish that t_2 will be such that the payoff of v^2 is 0.

If $T_1^D(q_1) < T_2^M(q_1)$, we have that one should set the supply plan as stated. Otherwise, we must serve arbitrarily close to m_1^*, the point when type v^1 is indifferent between pickup at this time and 0.

This concludes the proof.

A.7 Lemma 1

Take the principal's problem with the added restriction of using up to three batches and denote the time when the associated equilibrium adoption $\{M_t^E\}_t$ path reaches the target $\bar{M}$ by T_*^{3B}. Suppose that adding a new batch induces a new equilibrium adoption path $\{M_t^{E'}\}_t$ that reaches the target (absent breakdowns) at a time T_* that is strictly lower than T_*^{3B}.

In any optimal supply plan with three types, we must have the induced equilibrium adoption path, with the lowest type adopting myopically. This means that we must have the following:

$$\int_0^{T_*} M_\tau^{E'} d\tau = \int_0^{T_*^{3B}} M_\tau^E d\tau$$

It is also clear that adoption must start at time 0 and must include the highest types v^1. They must, therefore, weakly prefer adoption at time 0 to adoption at time T^{3B}. The first moment in which a type v^2 adopts must also have this characteristic.

We have three possibilities for the equilibrium adoption path M_t^E:

- It might jump at some time t'.

- It might stay constant.
- It might increase convexly for an interval.

Note that, given the equality above, all types must strictly prefer to take up at T_* after $\{M_t^{E'}\}_t$ when compared to T^{3B} after $\{M_t^{E'}\}_t$. Denote the first time in which any type v^2 adopts by T_*^2. Denote also the preferences of a type v^n picking up with a certain adoption plan M_t in mind by $\succeq^{n|M_t}$, with indifference represented by $\sim^{n|M_t}$. For example, $t \succeq^{1|M_t} t'$ if type v^1 weakly prefers taking up at time t to adopting at time t'.

By the fact that M_t^E is an equilibrium adoption plan, we must have $0 \succeq^{1|M_t^E} T_*^2$, $0 \succeq^{1|M_t^E} T_*^{3B}$, $T_*^2 \succeq^{2|M_t^E} 0$, $T_*^2 \succeq^{2|M_t^E} T_*^{3B}$, $T_*^{3B} \succeq^{3|M_t^E} 0$ and $T_*^{3B} \succeq^{3|M_t^E} T_*^2$.

Suppose that $\{M_t^E\}$ is a step function with jumps at the three points 0,T_*^2 and T_*^{3B}. If the inequalities above are all strict, then the principal was able to reach her first-best $\{M_t^O\}_t$, we conclude that a new batch cannot help in this case, and our previous assumption is false.

If the inequalities are not strict, though, by the structure of the problem, either $0 \sim^{1|M_t^E} T_*^2$ or $T_*^2 \sim^{2|M_t^E} T_*^{3B}$. If the latter, as the types v^2 all prefer to take at $T_* < T_*^{3B}$, we must have that no v^2 type would want to pick up the good at the time of a second batch and we are back to the three types case. If the former, though, any fourth batch between 0 and $T_*^{2'}$, the time of the second batch in this alternative 4 batches plan would be preferred by all v^1 types and any between $T_*^{2'}$ and T_* would be preferred by all v^2 types, and therefore no positive mass would want to take up before.

We conclude, then, that three batches cannot be improved upon by a new batch.

A.8 Theorem 3

The proof of the theorem will be done through a series of steps, following the intuitive discussion done in the main text:

STEP 1: If $T_1^D(q_1) > T_3^M(q_1,0,0)$, the optimal supply-plan has two batches. One at time 0 serving up to m_1* and the second at $T_3^M(m_1^*,0,0)$, serving $\bar{M} - m_1^*$.

Suppose that another plan $\{S_t'\}_t$ is strictly preferred by the principal. By definition, then, we need to have its induced adoption path $\{M_t'\}$ being such that $M_{T'}' = \bar{M}$ at some $T' < T_3^M(m_1^*,0,0)$. As, by definition, agents of type v^3 must prefer adopting at time T' instead of never, they we must have $\int_0^{T'}$

B Example of optimal supply plan with 3 batches

Suppose that we have valuations for the three types given by $(v^1, v^2, v^3) = (1.2, 0.8, 0.5)$, a mass of each type given by $(q_1, q_2, q_3) = (0.5, 0.3, 0.2)$, $\beta = r = 1$ and $\mu_0 = 0.5$, so that $v_0^M = 1$.

Then $T_3^M(q_1, 0, 0) = 2ln(2) \sim 1.39$, $T_1^D(q_1) = 1.16587$, so we have that $T_1^D(q_1) < T^M(q_1, 0, 0)$.

What about $T_2^E(q_1)$? Then one is maximizing $e^{-t}(0.8 - e^{-q_1 t})$, so that $T_2^E(q_1) = 1.2572 < T_3^M(q_1, 0, 0)$, so we are also good here.

Finally, we also have that $T_1^D(q_1) \succ^{v^2} T_3^M(0, q_1)$, from the equations determining payoffs.

Together, these equations mean that mid-value agents are willing to take up at a point in time in which high-types no longer want to and can, therefore, only speed up learning for low-types. The $T_1^D(q_1) \succ^{v^2} T_3^M(0, q_1)$ condition guarantees that learning is not such that the mid-level agent would rather wait and take-up with the lowest types.