Providing Certainty*

Andrew B. Choi† Christoph Schlom‡ Chengyang Zhu§

April 2025

Preliminary and incomplete; please do not circulate.

Abstract

We study the endogenous provision of certainty. A principal chooses a policy at a future date, and wishes to match the policy to an uncertain state. An agent chooses when to make an irreversible investment, and wishes to invest only if he expects the policy will be favorable. Information about the state is publicly and gradually revealed over time. To incentivize the agent to invest early, the principal provides policy certainty. This is inefficient – the agent's benefit from certainty is always outweighed by the principal's cost from reduced policy flexibility. We apply our results to environmental subsidies and procurement of vaccines.

Keywords: policy uncertainty, moral hazard, stopping time

1 Introduction

It has long been recognized that economic agents are more willing to make long-term investments when there is less uncertainty about future policies. Progress has been made in understanding how certainty about fiscal, trade, or environmental policies encourage investment.[1] However, an important question remains unanswered: how should a policymaker endogenously *provide* certainty about future policies, taking into account not only the benefit but also the cost of certainty?

Suppose a policymaker reduces an agent's uncertainty about tomorrow's policy by announcing today her policy for tomorrow. This ties the policymaker's hands and prevents her from adapting the policy to states of the world that are uncertain today but will be

*We thank Tilman Borgers, Yeon-Koo Che, Emir Kamenica, Amanda Friedenberg, David Miller, Doron Ravid, and Mu Zhang for helpful comments and suggestions.

†University of Michigan, choiandy@umich.edu

‡UC Davis, cschlom@ucdavis.edu

§Boston University, zhuc@bu.edu

[1]For example, see Fernández-Villaverde et al. (2015), Gulen and Ion (2016), Handley and Limão (2017), Gowrisankaran et al. (2024), and Chen (2024).

realized tomorrow. That is, to provide certainty to the agent, the policymaker must herself become more vulnerable to the underlying uncertainty of the world.

At the heart of this trade-off is a tension between a principal who wishes to match her future decision to an uncertain future state, and an agent who wishes to match her current action to the principal's future decision. We study how the resolution of this conflict is shaped by the strategic behavior of the players. We uncover a novel form of moral hazard – because the agent has control over his investment decision, the principal provides too much certainty, and the agent delays investment.

In our model, a principal chooses a policy at some final period, and an agent chooses when, if ever, to make an irreversible investment. The principal wishes to match the policy to an uncertain binary state of the world – when the state is good, she prefers a higher policy. The principal always wants the agent to invest, while the agent wants to invest only if he expects the principal's policy to be sufficiently high. Information about the state is gradually revealed over time via a signal process that is publicly observed by both the principal and the agent. Because the signal is public, the agent has incentives to delay investing and wait for more information about the state, using the information to predict the principal's policy. We assume that the signal has a "breakdown" structure – in each period, either the players become certain that the state is bad, or the good state becomes more likely.

The principal commits ex ante to a policy rule, which specifies the policy to be implemented as a function of the terminal history of signal realizations. We assume that the agent's investment decision is not observable by the principal, so the policy rule cannot depend on the timing of the investment. We solve for the optimal policy rule that maximizes the principal's expected payoff subject to the agent's incentive constraints.

We first show that the optimal policy rule *provides certainty* to the agent by making the policies under the good and bad states less dispersed, relative to a Pareto-efficient benchmark (Theorem 1(a)). The reason is that the principal wants the agent to invest early before there is chance for bad news to be revealed, and to induce the agent to do so, the principal must lower the agent's option value of waiting for information. However, the agent's value of waiting comes from being able to avoid investing when the policy is low. Hence, to induce early investment, the principal must ensure that when the policy is low, it is not too low. This leads the principal to raise the policy in the bad state and lower it in the good state, reducing their spread. As we will explain, this can be viewed as a novel implication of the "bad news principle" (Bernanke, 1983) in a strategic environment.

Providing certainty is costly for the principal because she loses the flexibility to tailor policy to the state. In fact, we show that the level of certainty provided is always inefficiently high (Theorem 1(a) and Proposition 1). That is, both players would be better if they could agree that the principal will provide less certainty and the agent will not delay investment. Because the agent's investment decision cannot be contracted on, such

a Pareto improvement cannot be achieved; providing less certainty would increase the agent's value of waiting for more information and make him delay investing. Perhaps surprisingly, the result that the principal provides too much certainty remains true regardless of how much the agent values certainty. Even if the agent places a very high value on certainty, so that it would be efficient for the principal to provide a high level of certainty, the optimal policy rule requires that the principal provide even more of it.

When there are more than two periods, the principal may *front-load* certainty – conditional on the bad state, the policy is higher if the signal "broke down" in an earlier period (Theorem 1(b)). Again, this serves to reduce the agent's value of information and dissuade him from delaying investment.

We apply our results to study the design of environmental subsidies (Section 5). Here, because the agent's payoff is additively separable in his investment timing and the principal's subsidy, the optimal subsidy rule does not front-load certainty. That is, the subsidy is independent of the history of signal realizations, conditional on the state. Our results also shed light on how a buyer should procure a good from a seller, when the seller must first succeed in R&D to be able to produce the good.

We first analyze a simple two-period example (Section 2), previewing many of our main results. We then set up and analyze a multi-period model with a richer payoff structure (Sections 3 and 4), which allows us to both generalize the results from the example and obtain new insights. We then use our results to study applications (Section 5).

Related Literature. There is a long-standing literature that studies irreversible investment under uncertainty (Bernanke, 1983; Dixit et al., 1994; Friedman, 1968; Rodrik, 1991; Fernández-Villaverde et al., 2015; Gulen and Ion, 2016; Handley and Limão, 2017; Gowrisankaran et al., 2024; Chen, 2024).[2] These papers mostly take uncertainty as exogenous and study the benefit of policy certainty in encouraging investment. In contrast, we study how a principal to endogenously provides certainty, trading off the benefits against the costs, to incentivize a strategic agent.

The literature on dynamic commitment under uncertainty (Athey et al., 2005; Halac and Yared, 2014, 2022; Choi, 2024) has focused on settings where future information arrives privately to the principal. Choi (2024) is a close antecedent to this paper; in our language, it considers a principal who delays an agent's stopping decision by promising today to provide certainty tomorrow. In these papers, the focus is on how the principal incentivizes her future self to be truthful after receiving private information. In contrast, we study an environment where the principal and the agent observe the same public signal. This is highly relevant in many settings, such as in designing environmental or macroeconomic policies, but have received little attention. We show that the agent's ability to wait for future public information is precisely what causes the principal to

[2]Relatedly, Baker et al. (2016) introduces a way to measure policy uncertainty.

provide too much certainty. Enriching the agent's moral hazard constraint while dropping the principal's incentive constraints allows us to tractably derive sharp insights about the optimal dynamic policy rule.

Orlov et al. (2020) studies a sender who discloses information to a receiver about an exogenous state of the world, in an environment with a publicly observed stochastic process. Koh et al. (2024) provides a method for solving such information design problems. Grenadier and Wang (2005), Board (2007), and Kruse and Strack (2015) study principal-agent models where the agent chooses a stopping time and has private information. The application of our model to bilateral trade relates to advance market commitments for vaccines (Kremer and Glennerster, 2004; Kremer et al., 2020).

2 Example – Renewing Subsidies

Suppose a principal (she) wants an agent (he) to invest in a project. There are two periods; the agent can invest in period 0, invest in period 1, or never invest. The investment has an upfront cost of $I > 0$. Once the agent invests, in each period, he produces one unit of a good at zero marginal cost. That is, if the agent invests in period 0, he produces in each of periods 0 and 1, while investing in period 1 means that he produces only in period 1. An interpretation is that a government wishes to encourage a developer to build a wind farm which generates clean electricity once it is built.

The agent's benefit from production comes solely from per-unit subsidies provided by the principal. We assume that the subsidy in period 0 is exogenously fixed at $s > 0$, while the subsidy in period 2, denoted $y \geq 0$, is the principal's choice variable. An interpretation is that s has been determined by past legislation, which expires after period 0. The principal's cost of paying y is $(2-\theta)y$, where θ is an uncertain state of the world that takes a value of $\theta = 1$ (good) or $\theta = 0$ (bad). For example, $\theta = 0$ may represent a negative fiscal shock that increases the government's opportunity cost of subsidies. In period 0, the parties share a common prior $p = \Pr(\theta = 1)$. At the beginning of period 1, before the agent makes the investment decision, θ is publicly revealed.

The game is summarized by the following timeline:

Period 0: Agent chooses $a_0 \in \{\texttt{Invest}, \texttt{No}\}$.

If $a_0 = \texttt{Invest}$, agent receives $s - I$.

Period 1: State $\theta \in \{0, 1\}$ is publicly realized.

If $a_0 = \texttt{No}$, agent chooses $a_1 \in \{\texttt{Invest}, \texttt{No}\}$.

If $a_1 = \texttt{Invest}$, agent receives $-I$.

If $a_0 = \texttt{Invest}$ or $a_1 = \texttt{Invest}$, agent receives y, and principal receives $-(2-\theta)y$.

Letting $\delta \in (0,1]$ denote the discount factor, the agent's final payoff is equal to $s+\delta y-I$ if he invests in period 0, equal to $\delta y - \delta I$ if he invests in period 1, and equal to 0 if he never invests. We assume $s \in ((1-\delta)I, (1-\delta p)I)$. The lower bound means that the benefit s obtained by investing early is greater than the benefit from discounting the investment cost, $(1-\delta)I$, which is obtained by investing late. As will become clear, the upper bound means that s is not so large as to dominate the entire benefit from waiting, which includes not only $(1-\delta)I$ but also the informational benefit of observing θ. For simplicity, we assume that the principal does not incur a cost from paying the exogenous subsidy s.[3]

We will postpone specifying the principal's benefit from inducing the agent to invest. For now, suppose the principal's objective is to minimize the principal's expected subsidy cost while inducing the agent to invest in period 0. Throughout, we endow the principal with ex ante commitment power and assume the public state θ is court-verifiable, so that the principal can commit to a level of y that is contingent on the realization of θ. We are mainly interested in the case where the timing of the agent's investment is not observable by the principal.[4] However, we will start by considering a benchmark case with observable (and verifiable) investment.

2.1 Benchmark: Verifiable Investment

When the principal can verify the exact period in which the agent invested, it is without loss for the principal to commit to provide no subsidies unless the agent invests in period 0. Then, the agent has no reason to wait until period 1 – he should either invest in period 0 or never invest. Thus the principal's problem is to choose a subsidy rule $y(\theta)$, which specifies the subsidy to be paid if the agent invests in period 0 and the realized state is θ, in order to minimize expected cost subject to the *individual rationality* constraint that the agent prefers investing in period 0 to never investing:

$$\begin{aligned} &\min_{y(0),y(1)} py(1) + 2(1-p)y(0) \\ &\text{subject to } \; s + \delta(py(1) + (1-p)y(0)) - I \geq 0, \end{aligned} \tag{IR}$$

where the LHS is the agent's expected payoff from investing in period 0. It is easy to see that $(\hat{y}(0), \hat{y}(1)) = (0, \frac{I-s}{p\delta})$ solves the above problem with the IR constraint binding. Intuitively, since the IR constraint only depends on the expected policy, it is optimal for

[3]This may be because the government operates on a lose-it-or-use-it budget. Our analysis remains unchanged even if we assume that s is costly for the principal.

[4]Since the principal subsidizes production, she must be able to observe it. However, it may nevertheless be difficult for the principal to observe the exact timing of the intermediate investments that led to output. Even when investment can be observed, concerns about distortion of input choice may lead policymakers to subsidize output (Parish and McLaren, 1982). In the US, the leading federal subsidy for wind farm developers has been a tax credit that is proportional to the amount of electricity generated (Aldy et al., 2023).

the principal to only provide subsidies when doing so is cheap.

2.2 Unverifiable Investment

Now, consider the principal's problem when the timing of the agent's investment cannot be verified. In this case, the benchmark outcome cannot be implemented. Indeed, suppose the principal commits to pay $\hat{y}(\theta)$ in state θ for period-1 production, regardless of investment timing. Then, the agent has a profitable deviation: he can wait until period 1, observe the state, and then invest if and only if $\theta = 1$. This gives the agent an expected payoff of $(1-\delta p)I - s$, which is strictly positive by assumption and thus greater than the expected payoff of 0 that the agent gets from investing in period 0.

To prevent such a deviation, a feasible subsidy rule must additionally satisfy the *no-waiting constraint*:

$$s + \delta(py(1) + (1-p)y(0)) - I \geq \delta p(y(1) - I), \tag{NW}$$

where the LHS is once again the agent's expected payoff from investing in period 0, and the RHS is the expected payoff from waiting and then investing in period 0 only in the good state. In principle, the agent may also consider two other deviations: wait and invest only in the bad state, and wait and invest in both states. We ignore the former deviation for now, and will shortly verify that this is not a profitable deviation. Due to our assumption that $s > (1-\delta)I$, the latter deviation is always dominated by investing in period 0; intuitively, if the agent plans to invest in period 1 no matter what, he should just invest in period 0.

Therefore, the principal's problem under unverifiable investment is

$$\min_{y(0),y(1)} py(1) + 2(1-p)y(0)$$
$$\text{subject to (IR), (NW).}$$

It is straightforward to see that both constraints should bind at the optimum, and that the optimal solution is $(y^*(0), y^*(1)) = \left(\frac{(1-\delta p)I - s}{\delta(1-p)}, I\right)$. Since $y^*(0) < y^*(1)$, we have verified that the agent will not wish to deviate by waiting and then investing only in the bad state.

Comparing y^* to $\hat{y}$, we see that $y^*(0) > \hat{y}(0)$ and $y^*(1) < \hat{y}(1)$. That is, investment being unverifiable leads the principal to *provide certainty* to the agent by reducing the difference between subsidies under the two states.[5] To see why this must be optimal, observe that in (NW), the terms on each side containing $y(1)$ cancel out, so the principal must raise $y(0)$ in order to satisfy (NW). Intuitively, waiting for information is valuable

[5] Note that, in this example, the expected policy is the same in the two scenarios, $\mathbb{E}[y^*(\theta)] = \mathbb{E}[\hat{y}(\theta)] = \frac{I-s}{\delta}$.

for the agent because, and only because, it allows him to avoid investing in the bad state. So, while raising $y(1)$ equally raises he agent's expected payoff from investing in period 0 and not waiting and thus has no impact on the agent's decision to wait, raising $y(0)$ reduces the agent's value of information and thus induces him not to wait. Since a higher $y(0)$ relaxes the IR constraint, the principal can lower $y(1)$.

The bad news principle (Bernanke, 1983) states that, for the decision whether to invest early, only the severity of future bad news matters. Our example uncovers a *strategic implication* of the bad news principle. For the agent's decision whether to invest in period 0 or wait, what matters is that when the subsidy is low, it is not too low. Therefore, to prevent the agent from waiting, the principal must provide more subsidies precisely when subsidies are more costly for her.

The provision of certainty leads to a loss of efficiency. The subsidy rule y^* is clearly more costly than $\hat{y}$ for the principal, as she provides the same level of subsidies in expectation but pays relatively more in the bad state. On the other hand, conditional on investing in period 0, the agent has zero value for certainty because her payoff is linear in the subsidy; the agent's expected payoff is 0 under both y^* and $\hat{y}$. In other words, under y^*, the principal provides certainty which is costly for herself but worthless to the agent. Clearly, y^* is Pareto-dominated by $\hat{y}$.

We view this as a novel form of moral hazard. In a standard moral hazard problem (Holmström, 1979), a principal conditions wage on endogenous output in order to incentivize an agent to exert effort; this exposes the agent to risk, which is inefficient if the agent is risk-averse. In our model, a principal conditions her policy on an exogenous state of the world. The principal promises to enact a favorable policy in the bad state, not because she wants to insure a risk-averse agent, but because this destroys the agent's option value of waiting for information and induces him to invest early.

One may wonder if the loss of efficiency in the current example relies on the agent being risk neutral. That is, if the agent values certainty, might the principal's provision of certainty actually improve welfare? Our general model will show that this is never the case. Regardless of how much the agent values certainty, the principal will always provide too much of it.

2.3 Optimal Investment Timing

So far, we have considered the cost-minimizing subsidy rule that induces the agent to invest in period 0. Given our results, should the principal induce the agent to invest early or induce him to wait for information? Before answering this question, we must now fully specify the principal's payoffs. Let us assume that the principal receives a benefit $b > 0$ if the agent invests, and that this benefit does not depend on when the agent invests.[6]

[6]For example, the principal may only care that the investment occurs before the next election.

Then, the principal's expected payoff from inducing the agent to invest in period 0 is

$$\begin{aligned} & b - [py^*(1) + 2(1-p)y^*(0)] \\ =& b - \frac{2}{\delta}(I - s) + pI. \end{aligned} \tag{1}$$

The principal's expected payoff from inducing the agent to wait until period 1 and invest only in state $\theta = 1$ is

$$p(b - I). \tag{2}$$

The principal induces the agent to invest in period 0 whenever (1) is positive and is greater than (2). This is the case if the investment cost I is low enough.[7]

Remark 1. The principal has no intrinsic preference for *early* investment, as she obtains the same benefit b regardless of when the agent invests. Nevertheless, the principal provides costly certainty in order to induce the agent to invest early. This is because the agent obtains an exogenous benefit $s - (1 - \delta)I$ from investing early. Although this exogenous benefit is not by itself large enough to make the agent invest early, it becomes enough once it is supplemented by the certainty provided by the principal. In turn, this allows the principal to induce investment at a lower cost to herself.

Finally, we consider how the agent's moral hazard (i.e. non-verifiability of investment) affects the timing of investment. When investment is verifiable, the principal's expected payoff from inducing the agent to invest in period 0 is $b - p\hat{y}(1) = b - \frac{I-s}{\delta}$, which is strictly lower than (1). On the other hand, her expected payoff from inducing the agent to wait and invest only in the good state is $p(b - I)$, which is the same as (2). Therefore, the agent's moral hazard always weakly delays investment. If the principal induces late investment without moral hazard, she also induces late investment when there is moral hazard; however, even if the principal induces early investment without moral hazard, she may induce late investment when there is moral hazard.

There are two main takeaways from the example described in this section. First, the principal provides certainty in order to induce the agent to invest, and this is inefficient. Second, moral hazard delays investment. In subsequent sections, we explore how these conclusions generalize in a model with a richer payoff structure and multiple periods.

3 Setup

Preliminaries. A principal (she) and an agent (he) interact over a finite number of periods $t \in \{0, 1 \ldots, T\}$. We will abuse notation and sometimes let T denote the set of

[7] For example, this happens if $I = 2$, $s = 0.4$, $p = 0.5$, $\delta = 0.9$, and $b = 4$.

all periods. In period T, the principal chooses a policy $y \in Y = [0, \overline{y}] \subset \mathbb{R}_+$. The agent chooses when, if ever, to make a single irreversible investment; that is, the agent solves a stopping time problem. We say that the agent stops at time $t = \infty$ to mean that the agent has not stopped by the end of period T. There is a payoff-relevant state of the world θ which can take values 1 (good) or 0 (bad), with a common prior $\Pr(\theta = 1) \in (0, 1)$.

To introduce assumptions on the players' payoffs, we start with a definition.

Definition 1. A function $f(\theta, y, t)$ satisfies **aggregate positive monotonicity (APM)** if, whenever

$$\sum_i \alpha_i f(\theta_i, y_i, t) \geq 0$$

for $t \in T$ and a finite sequence of $\alpha_i > 0$, θ_i and y_i, then, for all $t' < t$, we have:

$$\sum_i \alpha_i f(\theta_i, y_i, t') \geq \sum_i \alpha_i f(\theta_i, y_i, t). \tag{3}$$

We will interpret this definition shortly. If the agent invests at time t, the state is θ, and the principal's policy is y, the agent's payoff is $u(\theta, y, t)$. We impose the following assumptions on u.

Assumption 1. *The agent's payoff function $u(\theta, y, t)$ satisfies the following conditions:*

(a) u is differentiable in y, strictly increasing in y, and weakly concave in y.

(b) $u(1, y, t) \geq u(0, y, t)$ for all y, t.

(c) $u(\theta, y, \infty) = 0$ for all θ, y.

(d) $u(0, 0, t) < 0$ for all $t \in T$.

(e) $u_y(\theta, y, t) \leq u_y(\theta, y, t')$ for $t > t'$, where u_y denotes the partial derivative of u with respect to y.

(f) For $y > y'$ and $t > t'$, we have $u_y(\theta, y, t) - u_y(\theta, y', t) \geq u_y(\theta, y, t') - u_y(\theta, y', t')$.

(g) u satisfies APM.

(h) $u_y(0, y, t) \leq u_y(1, y, t)$, for all y, t.

(b) means that the agent prefers the good state to the bad state. (c) normalizes to 0 the agent's payoff from never investing, and (d) says that the agent prefers not to invest if the state is bad the the principal provides the lowest possible policy of 0. (e) means that the earlier the agent invests, the stronger his preference for higher policy; (f) means that investing earlier increases the curvature of u with respect to y, that is, $u_y(\theta, y, t)$ is supermodular fixing θ. APM (g) is a weak discounting notion. It says that, provided the agent is receiving a "good" lottery of rewards by stopping tomorrow, he would prefer to receive that same lottery and stop today. (h) says the agent's marginal value for policy is greater when the state is good.

If the agent invests at time t, the state is θ, and the principal's policy is y, the principal's payoff is $v(\theta, y, t)$. We impose the following assumptions on v.

Assumption 2. *The principal's payoff function $v(\theta, y, t)$ satisfies the following conditions:*

(a) v is differentiable in y and strictly concave in y.

(b) $v(1, y, t) \geq v(0, y, t)$ for all y, t.

(c) $v_y(\theta, y, \infty) \leq 0$ for all θ, y.

(d) $v(\theta, 0, t)$ is non-negative and is non-increasing in t.

(e) v satisfies APM.

(f) $v_y(0, y, t) \leq v_y(1, y, t)$, for all y, t.

(b) means that the principal prefers the good state to the bad state. (c) says that if the agent never invests, the principal does not benefit from setting a higher policy. (d) means that the principal can guarantee herself a non-negative payoff by providing the lowest possible policy of 0; and when she does, she weakly prefers earlier investment. (f) says that the principal's marginal value of policy (which may be negative) is greater in the good state.

Our assumptions on the payoff functions u and v capture a wide range of applications, such as those provided in Section 5.

Markov Signals. In each period, the principal and the agent symmetrically update information about θ by observing a public Markov signal x_t. Each x_t takes a value of 1 (good) or 0 (bad), and $x_T = \theta$ so that the players fully learn the state in period T. Signal 0 is absorbing, while $\Pr(x_{t+1} = 1 | x_t = 1) = p_t \in (0, 1)$. We say that a breakdown occurred in period t to mean that the signal transitioned from $x_{t-1} = 1$ to $x_t = 0$. The initial signal is assumed to be $x_0 = 1$.

Timeline. Within each period $t \in \{0, 1, \ldots, T\}$, the timeline of events is as follows:

1. Signal x_t is publicly realized.
2. If the agent has not already invested, he decides whether to invest or not.
3. If $t = T$, the principal implements policy y.

Policy Rule. Define two types of terminal histories at period T:

$$\mathbf{a} := (\overbrace{1, \ldots, 1}^{T \text{ entries}}), \qquad \mathbf{b}_t := (\overbrace{1, \ldots, 1}^{t-1 \text{ entries}}, \overbrace{0, \ldots, 0}^{T-t+1 \text{ entries}}).$$

The history $\mathbf{a}$ indicates that there has been no breakdown. The history $\mathbf{b}_t$ indicates that there was a breakdown at period t. Let $\mathcal{H} := \{\mathbf{a}, \mathbf{b}_1, \ldots, \mathbf{b}_T\}$ denote the set of possible terminal histories, and $\mathbf{h}$ a generic element of $\mathcal{H}$.

The principal commits ex ante to a **policy rule** $y : \mathcal{H} \to Y$. We are thereby making three assumptions regarding the principal's commitment power. First, the principal can condition her choice of policy on both the state and the signals. It is common for policies to be contingent on uncertain future states. For example, the Clean Electricity Production Credit in the US depends on annual US greenhouse gas emissions, which may be viewed as an exogenous state from the point of an individual firm.[8] Second, we assume that the policy rule does not depend on the agent's choice of whether and when to invest. This captures the fact that it is often difficult for a policymaker to observe the investment decisions of individuals, let alone their exact timings.[9] Finally, we restrict attention to deterministic policy rules.

Agent's Problem. The agent's choice of when to irreversibly invest can be modelled as the choice of a stopping time. Formally, let $\mathcal{T}$ denote the set of stopping times that are adapted to the natural filtration with respect to $\{x_t\}_{t\in\mathcal{T}}$. Taking as given the policy rule $y(\cdot)$ announced by the principal, the agent chooses $\tau \in \mathcal{T}$ to maximize his expected payoff:

$$\max_{\tau\in\mathcal{T}} \mathbb{E}[u(\theta, y(\mathbf{h}), \tau)],$$

where the expectation is taken over the terminal state θ, the terminal history $\mathbf{h}$, and the stopping time τ.

Principal's Problem. The principal's problem is to choose a policy rule $y(\cdot)$ and a stopping time τ to maximize her expected payoff, subject to the agent's incentive constraint:

$$[\mathcal{P}_0] \qquad \max_{y(\cdot),\,\tau\in\mathcal{T}} \mathbb{E}[v(\theta, y(\mathbf{h}), \tau)]$$
$$\text{subject to } \mathbb{E}[u(\theta, y, \tau)] \geq \mathbb{E}[u(\theta, y, \tau')] \qquad \forall \tau' \in \mathcal{T}.$$

[8] As our analysis will show, the optimal policy rule depends not only on the payoff-relevant state, but also on the history of signal realizations, which is not payoff-relevant conditional on the state. We view this as a normative result.

[9] One interpretation of our model is that there exists a continuum of small homogeneous agents. Then, even if the principal could condition her policy on the aggregate investment of the economy, this would not affect the incentives of the agents because each agent takes the policy as given.

4 Optimal Policy Rule

To analyze the principal's problem $[\mathcal{P}_0]$, we first study a simpler problem in which the agent is constrained to only choose among a smaller, one-dimensional class of stopping times (Section 4.1). We then argue that the solution to this simpler problem also solves the principal's original problem $[\mathcal{P}_0]$ (Theorem 2).

4.1 Simple Stopping Times

For $t \in \{0, 1, \ldots, T, \infty\}$, the **simple stopping time** $\tau_t \in \mathcal{T}$ is the random variable defined by

$$\tau_t = \begin{cases} t & \text{if } x_t = 1 \\ \infty & \text{if } x_t = 0, \end{cases}$$
$$\tau_\infty = \infty.$$

That is, an agent adopting τ_t always waits until period t; if the period-t signal is good, he stops immediately, and otherwise he never stops.[10] Given a policy rule $y(\cdot)$, the agent's expected payoff from the simple stopping time τ_t is

$$U_t(y) = \begin{cases} p_1 p_2 \cdots p_t \mathbb{E}[u(\theta, y(\mathbf{h}), t) \mid x_t = 1] & \text{if } t < \infty \\ 0 & \text{if } t = \infty, \end{cases}$$

where the expectation is taken with respect to the terminal state θ_T and the terminal history $\mathbf{h} \in \mathcal{H}$. In contexts where the policy rule is unambiguous, we will omit the argument and simply write U_t.

The principal's **auxiliary problem** is

$$[\mathcal{P}] \qquad \max_{y(\cdot),\, \tau_t} \mathbb{E}[v(\theta, y(\mathbf{h}), \tau_t)]$$
$$\text{subject to } U_t \geq U_{t'} \qquad \forall t' \in \{0, 1, \ldots, T, \infty\}.$$

The auxiliary problem $[\mathcal{P}]$ differs from $[\mathcal{P}_0]$ in two aspects – the principal cannot induce the agent to adopt a non-simple stopping time, and the agent cannot deviate to a non-simple stopping time. We will show in Section 4.2 that the former is without loss of optimality, while the latter is without "gain" of optimality.

[10] Non-simple stopping times are those in which the agent stops after observing a breakdown signal $x_t = 0$. Although such behavior is intuitively suboptimal, we cannot a priori rule it out because the policy is endogenous.

It will be useful to also define the **inner auxiliary problem** as follows:

$$[\mathcal{P}(t)] \qquad \max_{y(\cdot)} \mathbb{E}[v(\theta, y(\mathbf{h}), \tau_t)]$$
$$\text{subject to } U_t \geq U_{t'} \qquad \forall t' \geq t.$$

The problem $[\mathcal{P}(t)]$ differs from $[\mathcal{P}]$ in two aspects. First $[\mathcal{P}(t)]$ takes as given the stopping time τ_t. Second, $[\mathcal{P}(t)]$ assumes that the agent cannot deviate to any "earlier" simple stopping time $\tau_{t'}$ with $t' < t$. That we can ignore such deviations is established by the following lemma.

Lemma 1. *Let $V(t)$ denote the value of $[\mathcal{P}(t)]$. The following statements are true:*

(a) For each $t \in T$, there exists a solution to $[\mathcal{P}(t)]$ with $y(\mathbf{b}_s) = 0$ for all $s \leq t$.

(b) There exists a solution (y, τ_{t^*}) to $[\mathcal{P}]$ such that y solves $[\mathcal{P}(t^*)]$ and t^* solves $\arg\max_t V(t)$.

(c) If (y, τ_{t^*}) solves $[\mathcal{P}]$, then y solves $[\mathcal{P}(t^*)]$ and t^* solves $\arg\max_t V(t)$.

Proof. See Appendix A.3. □

The intuition for Lemma 1 is that, if y solves $[\mathcal{P}(t)]$, and the agent wishes to deviate to an earlier stopping time under y, then it is better for the principal to induce the earlier stopping time. Lemma 1(b) indicates that, in order to solve $[\mathcal{P}]$, it is enough to solve the problem $[\mathcal{P}(t)]$ for each t and then maximize over t. Lemma 1(c) allows us derive necessary properties of the solution to $[\mathcal{P}]$, which are stated in Theorem 1 below.

In stating Theorem 1, we shall say that a constraint is *slack* to mean that the value of the maximization program $[\mathcal{P}]$ remains unchanged even if one were to remove the constraint.[11]

Theorem 1. *Suppose (y, τ_t) solves $[\mathcal{P}]$ for some $t \in T$, and suppose $y(\mathbf{h})$ is interior for $\mathbf{h} \in \{\mathbf{b}_{t+1}, \ldots, \mathbf{b}_T, \mathbf{a}\}$. Then the following statements hold:*

(a) It must be that

$$\frac{v_y(0, y(\mathbf{b}_T), t)}{u_y(0, y(\mathbf{b}_T), t)} \leq \frac{v_y(1, y(\mathbf{a}), t)}{u_y(1, y(\mathbf{a}), t)}, \tag{4}$$

where (4) holds with equality if and only if the constraint $U_t \geq U_T$ is slack.

(b) It must be that

$$y(\mathbf{b}_{t+1}) \geq y(\mathbf{b}_{t+2}) \geq \cdots \geq y(\mathbf{b}_T). \tag{5}$$

[11] That is, a constraint is slack if and only if the corresponding Lagrangian multiplier may be set to be 0. Note that this is weaker than saying that the constraint holds with a strict equality at the solution, since it may be that the multiplier is zero, yet the unconstrained solution happens to satisfy the constraint.

For each $s = 1, 2, \ldots, T-t-1$, we have $y(\mathbf{b}_{t+s}) = y(\mathbf{b}_{t+s+1})$ if and only if the constraint $U_t \geq U_{t+s}$ is slack.

Proof. See Appendix A.5. □

When the inequality (4) is strict, the policy in the good state is inefficiently low, and policy in the bad state is inefficiently high. This can be viewed as a violation of the Borch rule. Both the principal and the agent would be better off if they could agree that the agent would use the stopping time τ_t, and that the principal chooses policy more flexibly in response to the state θ. However, such a Pareto improvement is not feasible because it would lead the agent to wait for more information instead of stopping at time t after seeing $s_t = 1$.

One can show that the optimal policy always satisfies $y(a) > y(b_T)$. As a result, lowering $y(a)$ and raising $y(b_T)$ decreases the difference between the two policies. Therefore, we interpret (4) to mean that *the principal provides too much certainty.*

The inequalities (5) mean that *certainty is front-loaded* – the principal provides more certainty when the breakdown occurs earlier. To discourage the agent from waiting, the principal makes earlier signals less informative about her choice of policy.

Benchmark: Verifiable Investment. For comparison, consider a benchmark case where the timing of investment is verifiable. Then, to induce the agent to adopt the stopping time τ_t, the principal just needs to ensure that the agent prefers τ_t to never stopping (i.e. adopting τ_∞). That is, the principal's auxiliary problem is

$$\max_{y(\cdot),\, \tau_t} \mathbb{E}[v(\theta, y(\mathbf{h}), \tau_t)]$$
$$\text{subject to } U_t \geq 0.$$

In this case, the policy only depends on the final payoff-relevant state, and not the entire history.

Proposition 1. *Suppose (y, τ_t) solves the principal's problem with verifiable investment, and $y(\mathbf{h})$ is interior for $\mathbf{h} \in \{\mathbf{b}_{t+1}, \ldots, \mathbf{b}_T, \mathbf{a}\}$. Then it must be that*

$$y(\mathbf{b}_{t+1}) = y(\mathbf{b}_{t+2}) = \cdots = y(\mathbf{b}_T) =: y(\mathbf{b})$$

and

$$\frac{v_y\left(0, y(\mathbf{b}), t\right)}{u_y\left(0, y(\mathbf{b}), t\right)} = \frac{v_y\left(1, y(\mathbf{a}), t\right)}{u_y\left(1, y(\mathbf{a}), t\right)}.$$

Given τ_t, the policy rule $y(\cdot)$ is Pareto efficient.

Proof. See Appendix A.4. □

4.2 Optimal Policy Rule

The solution to the auxiliary problem $[\mathcal{P}]$, which we described in Theorem 1, in fact solves the original problem $[\mathcal{P}_0]$.

Theorem 2. *If* (y, τ_t) *solves* $[\mathcal{P}]$*, and* $y(\mathbf{h})$ *is interior for* $\mathbf{h} \in \{\mathbf{b}_{t+1}, \ldots, \mathbf{b}_T, \mathbf{a}\}$*, then* (y, τ_t) *solves* $[\mathcal{P}_0]$*.*

The theorem contains two results. First, the original program $[\mathcal{P}_0]$ allows the principal to induce the agent to choose a non-simple stopping rule, but the principal never benefits from doing so. Second, under the policy y that solves $[\mathcal{P}]$, agent has no profitable deviation to a non-simple stopping rule. This is where the "guess-and-verify" approach proves useful. The properties of the solution to the auxiliary problem (Theorem 1) allows us to argue that for any non-simple stopping rule, there exists a simple stopping rule that is better for the agent.

5 Applications

5.1 Renewing Subsidies

Consider the extension of the two-period example in Section 2 to T periods. Once the agent invests, he receives the exogenous benefit s in every period. That is, the agent's payoff from investing in period t is

$$u(\theta, y, t) = (\delta^t + \cdots + \delta^{T-1})s + \delta^T y - \delta^t I.$$

It can be verified that this satisfies the conditions in Assumption 1 (see Appendix A.2). The transition probability in each period is constant at $p \in (0, 1)$. We assume that the principal's cost of subsidy is quadratic:

$$v(\theta, y, t) = b - \frac{1}{2}y^2.$$

Using Theorems 1 and 2, we obtain the following result:

Proposition 2 (No Front-Loading)**.** *Suppose* (y, τ_t) *for some* $t \in T$ *solves the subsidy renewal problem. Then, exactly one of the following three conditions holds:*

(a) $U_t > \max\{U_{t+1}, \ldots, U_T, 0\}$

(b) $U_t > \max\{U_{t+1}, \ldots, U_T\}$ *and* $U_t = 0$

(c) $U_t > \max\{U_{t+1}, \ldots, U_{T-1}\}$ *and* $U_t = U_T = 0$.

Proof. See Appendix A.6. □

The proposition says that says that only the last two constraints, $U_t \geq U_T$ and $U_t \geq 0$, may bind at the optimum. Strikingly, this means that the optimal policy rule does not involve any front-loading of uncertainty, as all $y(b_t)$'s are equal. One may be concerned about a policymaker's ability to commit to condition her policy on past signal realizations that are payoff-irrelevant. Proposition 2 says that, in the context of environmental subsidies, it is without loss of optimality for the policy to only depend on the payoff-relevant state.

5.2 Bilateral Trade

The principal (a buyer) wishes to procure from the agent (a seller) a good at time T. The seller may or may not be able to produce the good at time T. For example, the good may be a vaccine for a new virus that takes a stochastic amount of time to develop. To produce the good, the seller must first invest in development, and then the development must succeed. If the seller invests in period t, success arrives with probability q in each period $t, t+1, \ldots, T$. Hence, if the seller invests in period t, the probability that she succeeds by period T is $1-(1-q)^{T-t+1}$.

If the seller does not invest by period T, both players obtain 0 payoffs. If the seller invests but does not succeed in development by period T, the buyer's payoff is 0, and the seller's payoff is $-\delta I$, where $I > 0$ is the cost of investment and $\delta \in (0,1)$ is the discount rate. If the seller invests and succeeds by period T, the players have a chance to trade the good. If y units of the good are traded, the seller's payoff is[12]

$$y - I,$$

and the buyer's payoff is

$$\theta y^{\frac{1}{2}} - y,$$

where $\theta \in \{0,1\}$ represents the buyer's uncertain future value for the good. For example, the value of a vaccine may depend on how quickly a new virus is spreading. To incentivize the seller to invest early, the buyer can make the following form of *advance market commitment* (Kremer and Glennerster, 2004; Kremer et al., 2020): if the seller's research succeeds, and if the terminal state is θ_T, the buyer will purchase $y(\theta_T)$ units of the good.

The seller's expected payoff from investing in period t, conditional on policy y and terminal state θ_T, is

$$(1-(1-q)^{T-t+1})y - \delta^t I, \tag{6}$$

Although (6) is a (conditional) expected payoff because the coefficient of y is a probability,

[12] We assume in this example that the per-unit price of the good is fixed at 1.

we may nevertheless define the seller's payoff function $v(\theta, y, t)$ to be (6) to incorporate this application under our setup.

Notably, this is an example where the investment timing t interacts with y. The earlier the agent invests, the more exposed he is to the principal's policy. This is in contrast to the subsidies application, where the agent's payoff was additively separable in t and y.

References

Aldy, J. E., T. D. Gerarden, and R. L. Sweeney (2023): "Investment versus output subsidies: Implications of alternative incentives for wind energy," *Journal of the Association of Environmental and Resource Economists*, 10(4), 981–1018. 5

Athey, S., A. Atkeson, and P. J. Kehoe (2005): "The optimal degree of discretion in monetary policy," *Econometrica*, 73(5), 1431–1475. 3

Baker, S. R., N. Bloom, and S. J. Davis (2016): "Measuring economic policy uncertainty," *The quarterly journal of economics*, 131(4), 1593–1636. 3

Bernanke, B. S. (1983): "Irreversibility, uncertainty, and cyclical investment," *The quarterly journal of economics*, 98(1), 85–106. 2, 3, 7

Board, S. (2007): "Selling options," *Journal of Economic Theory*, 136(1), 324–340. 4

Chen, L. (2024): "The Dynamic Efficiency of Policy Uncertainty: Evidence from the Wind Industry," *Working Paper*. 1, 3

Choi, A. B. (2024): "I'll Tell You Tomorrow: Committing to Future Commitments," *Working Paper*. 3

Dixit, R. K., A. K. Dixit, and R. S. Pindyck (1994): *Investment under uncertainty*. Princeton university press. 3

Fernández-Villaverde, J., P. Guerrón-Quintana, K. Kuester, and J. Rubio-Ramírez (2015): "Fiscal volatility shocks and economic activity," *American Economic Review*, 105(11), 3352–3384. 1, 3

Friedman, M. (1968): "The role of monetary policy," *American economic review*, 58(1), 1–17. 3

Gowrisankaran, G., A. Langer, and W. Zhang (2024): "Policy uncertainty in the market for coal electricity: The case of air toxics standards," Discussion paper, Working Paper. 1, 3

Grenadier, S. R., and N. Wang (2005): "Investment timing, agency, and information," *Journal of Financial Economics*, 75(3), 493–533. 4

Gulen, H., and M. Ion (2016): "Policy uncertainty and corporate investment," *The Review of Financial Studies*, 29(3), 523–564. 1, 3

Halac, M., and P. Yared (2014): "Fiscal rules and discretion under persistent shocks," *Econometrica*, 82(5), 1557–1614. 3

——— (2022): "Instrument-based versus target-based rules," *The Review of Economic Studies*, 89(1), 312–345. 3

Handley, K., and N. Limão (2017): "Policy uncertainty, trade, and welfare: Theory and evidence for China and the United States," *American Economic Review*, 107(9), 2731–2783. 1, 3

Holmström, B. (1979): "Moral hazard and observability," *The Bell journal of economics*, pp. 74–91. 7

Koh, A., S. Sanguanmoo, and W. Zhong (2024): "Persuasion and optimal stopping," *arXiv preprint arXiv:2406.12278*. 4

Kremer, M., and R. Glennerster (2004): *Strong medicine: creating incentives for pharmaceutical research on neglected diseases*. Princeton University Press. 4, 16

Kremer, M., J. Levin, and C. M. Snyder (2020): "Advance market commitments: insights from theory and experience," in *AEA Papers and Proceedings*, vol. 110, pp. 269–273. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203. 4, 16

Kruse, T., and P. Strack (2015): "Optimal stopping with private information," *Journal of Economic Theory*, 159, 702–727. 4

Orlov, D., A. Skrzypacz, and P. Zryumov (2020): "Persuading the principal to wait," *Journal of Political Economy*, 128(7), 2542–2578. 4

Parish, R. M., and K. R. McLaren (1982): "Relative cost-effectiveness of input and output subsidies," *Australian Journal of Agricultural Economics*, 26(1), 1–13. 5

Rodrik, D. (1991): "Policy uncertainty and private investment in developing countries," *Journal of Development Economics*, 36(2), 229–242. 3

A Proofs

A.1 Proof of Theorem 2 – No stopping after a breakdown

Additional notation. A **history**, $\mathbf{g}$, is a sequence of states. For any $t \in T$, and any $s \leq t$, define:

$$\mathbf{a}_t := (\overbrace{1, \ldots, 1}^{t \text{ entries}}), \qquad \mathbf{b}_{s,t} := (\overbrace{1, \ldots, 1}^{s-1 \text{ entries}}, \overbrace{0, \ldots, 0}^{t-s+1 \text{ entries}}).$$

The history $\mathbf{a}_t$ indicates that it is period t and there has been no breakdown. The history $\mathbf{b}_{s,t}$ indicates that it is period t and there was a breakdown at period s. Additionally, define the history $\mathbf{a}_0$, which indicates that it is period 0. Let

$$\mathcal{G} := \mathbf{a}_0 \cup (\cup_t \mathbf{a}_t) \cup (\cup_{t,s \leq t} \mathbf{b}_{s,t})$$

denote the set of possible histories. Notice that a terminal history is simply a history of length T, and that, accordingly, $\mathcal{H} \subset \mathcal{G}$.

Say that a history, $\mathbf{g}'$, of length t, is a **sub-history** of a strictly longer history, $\mathbf{g}$, if the first t entries of $\mathbf{g}$ are given by $\mathbf{g}'$. The history $\mathbf{a}_0$ is defined to be a sub-history of any other history. We will indicate the sub-history relation using the shorthand $\mathbf{g}' \prec_{sh} \mathbf{g}$.

A **stopping rule**, d, is a map from $\mathcal{G}$ to the set $\{\texttt{wait}, \texttt{stop}, \varnothing\}$, with the property that:

$$d(\mathbf{g}) = \varnothing \iff \exists\, \mathbf{g}' \prec_{sh} \mathbf{g} \text{ with } d(\mathbf{g}') = \texttt{stop}.$$

Given some stopping rule, let $\tau : \mathcal{H} \to T \cup \infty$ denote the **stopping time function**, which maps a terminal history to the time the agent would stop under that rule (or to ∞ if the agent would never stop). There is a one-to-one relationship between stopping rules, d, and stopping time functions, τ; we will use τ_d to represent the stopping time function induced by stopping rule d, and d_τ to represent the stopping rule induced by stopping time function τ. Notice that this is a notation change from the main text: d and τ are now both functions of histories, rather than random variables, and τ represents the stopping time.

Define the **simple stopping rules**, d_t, for $t \in T$, by:

$$\begin{aligned} d_t(\mathbf{a}_t) &= \texttt{stop} \\ d_t(\mathbf{b}_{i,j}) = d_t(\mathbf{a}_s) &= \texttt{wait}, \text{ for all } j, i \leq j \text{ and all } s < t \\ d_t(\mathbf{a}_x) &= \varnothing, \text{ for all } x > t. \end{aligned}$$

The rule d_t describes an agent who stops iff it is period t and a breakdown has not occurred. Further define $d_\infty(\mathbf{g}) = \texttt{wait}$, for all $\mathbf{g} \in \mathcal{G}$, which describes an agent who never stops.

The following notations related to agent utility will be useful. First, define:

$$p(\mathbf{g}) := \begin{cases} \left(\prod_{i=1}^{s-1} p_i\right)(1-p_s) & \mathbf{g} = \mathbf{b}_{s,t} \\ \prod_{i=1}^{t} p_i & \mathbf{g} = \mathbf{a}_t, \end{cases}$$

to be the probability that the history $\mathbf{g}$ will be reached at some point. (Notice that p is also defined for terminal histories, via the relationships $\mathbf{a} = \mathbf{a}_T$ and $\mathbf{b}_t = \mathbf{b}_{t,T}$.) Then, define:

$$u\langle \mathbf{h}, \tau | y \rangle := \begin{cases} u(\theta = 0, y(\mathbf{b}_t), \tau(\mathbf{b}_t)) & \mathbf{h} = \mathbf{b}_t \\ u(\theta = 1, y(\mathbf{a}), \tau(\mathbf{a})) & \mathbf{h} = \mathbf{a}, \end{cases}$$

to be the agent's (discounted) utility given terminal history h, stopping time function τ and policy y. Similarly, define:

$$u\langle \mathbf{g} | y \rangle := \begin{cases} u(\theta = 0, y(\mathbf{b}_s), t) & \mathbf{g} = \mathbf{b}_{s,t} \\ (\prod_{i=t+1}^{T} p_i) u(\theta = 1, y(\mathbf{a}), t) + & \\ \quad \sum_{s=t+1}^{T} (\prod_{i=t+1}^{s-1} p_i)(1-p_s) u(\theta = 0, y(\mathbf{b}_s), t) & \mathbf{g} = \mathbf{a}_t, \end{cases}$$

to be the agent's (discounted) expected utility from stopping at history $\mathbf{g}$, given policy y.

Denote the agent's (discounted) expected utility from using stopping rule d, given policy y, by $u\langle d|y\rangle$. It will be useful to compute $u\langle d|y\rangle$ in two ways. First, conditioning on terminal histories, we have:

$$u\langle d|y\rangle = p(\mathbf{a}) u\langle \mathbf{a}, \tau_d | y \rangle + \sum_t p(\mathbf{b}_t) u\langle \mathbf{b}_t, \tau_d | y \rangle. \tag{7}$$

Second, conditioning on stopping histories, we have:

$$u\langle d|y\rangle = \sum_{\{\mathbf{g} \in \mathcal{G} : d(\mathbf{g}) = \texttt{stop}\}} p(\mathbf{g}) u\langle \mathbf{g} | y \rangle. \tag{8}$$

Next, the following simple consequences of aggregate positive monotonicity will be useful. The first simply says that u is monotonic when it is non-negative.

Lemma 2 (Positive monotonicity)**.** *For any $\theta \in \{0,1\}$, $y \in \mathbb{R}$ and $t \in T$, such that $u(\theta, y, t) \geq 0$, and any $s < t$, we have:*

$$u(\theta, y, s) \geq u(\theta, y, t).$$

Proof. Use the single-element sequence $(1, \theta, y)$ in the definition of (APM). □

The second says that, if, no matter tomorrow's news, the agent would derive non-negative utility from stopping, then he is better off stopping today than tomorrow – in

other words, he does not want pointless delay. We call this the early sure-thing principle (ESTP).

Lemma 3 (Early sure-thing principle)**.** *Fix a policy rule, $y : \mathcal{H} \to \mathbb{R}$.*

(a) For any $t \in T$, and $s < t$, such that $u\langle \mathbf{b}_{s,t}|y\rangle \geq 0$, we have:

$$u\langle \mathbf{b}_{s,s}|y\rangle \geq u\langle \mathbf{b}_{s,t}|y\rangle.$$

(b) For any $t \in T$, such that $u\langle d_{t+1}|y\rangle \geq 0$, $u\langle \mathbf{b}_{t+1,t+1}|y\rangle > 0$, we have:

$$u\langle d_t|y\rangle > u\langle d_{t+1}|y\rangle.$$

Proof. (a) Set $\theta = 0$ and $y = y(\mathbf{b}_s)$ in Lemma 2.

(b) Using equation (7), we have:

$$\begin{aligned} u\langle d_t|y\rangle - u\langle d_{t+1}|y\rangle =& p(\mathbf{b}_{t+1})u(0, y(\mathbf{b}_{t+1}), t) \\ &+ p(\mathbf{a})u(1, y(\mathbf{a}), t) + \sum_{i=t+2}^{T} p(\mathbf{b}_i)u(0, y(\mathbf{b}_i), t) \\ &- p(\mathbf{a})u(1, y(\mathbf{a}), t+1) - \sum_{i=t+2}^{T} p(\mathbf{b}_i)u(0, y(\mathbf{b}_i), t+1), \end{aligned}$$

where the third line is $u\langle d_{t+1}|y\rangle$. Then, (APM) implies that the sum of the second and third lines is non-negative, while (a) implies that the first line is strictly positive. Thus, $u\langle d_t|y\rangle > u\langle d_{t+1}|y\rangle$. □

Let us now define another optimization problem for the principal, which we call the principal's inner problem.

$$[\mathcal{P}_0(t)] \qquad \max_{y(\cdot)} \mathbb{E}[v(\theta, y(\mathbf{h}), \tau_t)]$$
$$\text{subject to } \mathbb{E}[u(\theta, y, \tau_t)] \geq \mathbb{E}[u(\theta, y, \tau')] \qquad \forall \tau' \in \mathcal{T}.$$

In $[P_0(t)]$, as in $[P(t)]$, the principal seeks to induce the agent to use the simple stopping rule d_t; however, in $[P_0(t)]$, the principal must also guard against deviations to non-simple stopping rules. Thus, $[P(t)]$ is a relaxation of $[P_0(t)]$.

Our proof of Theorem 2 proceeds in two parts. Our first result, Lemma 4, implies that any solution to $[P(t)]$ is feasible for (and therefore solves) $[P_0(t)]$. Our second result, **??**, shows that any solution to $[P_0]$ is a solution to $[P_0(t)]$, for some t – that is, the principal would optimally induce a simple stopping time. Together, these results establish the equivalence between $[P_0]$ and $[P]$.

Lemma 4 (Theorem 2, part a)**.** *Fix $t^* \in T$, and let y be a solution to $[\mathcal{P}(t^*)]$. Then, we must have $u\langle \mathbf{b}_{s,t}|y\rangle \leq 0$, for all $t \in T, s \leq t$.*

Proof. Suppose not – i.e., there exists some $\hat{s}, \hat{t}$ such that $u\langle \mathbf{b}_{\hat{s},\hat{t}}|y\rangle > 0$. Then, by ESTP, it must be that $u\langle \mathbf{b}_{\hat{s},\hat{s}}|y\rangle > 0$. Let s be the largest value of $\hat{s}$ satisfying $u\langle \mathbf{b}_{\hat{s},\hat{s}}|y\rangle > 0$. (Notice that $s > t^*$ by Lemma 1.)

Claim 1. For any $s' \leq s$, we have $u\langle \mathbf{b}_{s',s'}|y\rangle > 0$.

Proof. By Theorem 1, we have that $y(\mathbf{b}_{s'}) \geq y(\mathbf{b}_s)$. Then, we have:

$$\begin{aligned} u\langle \mathbf{b}_{s',s'}|y\rangle &= u(\theta = 0, y(\mathbf{b}'_s), s') \\ &\geq u(\theta = 0, y(\mathbf{b}_s), s') && \text{[Monotonicity of } u\text{]} \\ &\geq u(\theta = 0, y(\mathbf{b}_s), s) && \text{[ESTP]} \\ &= u\langle \mathbf{b}_{s,s}|y\rangle > 0. \end{aligned}$$

Next, let $t \in T \cup \infty$ be the first period such that $t > t^*$ and $u\langle d_t|y\rangle = u\langle d_{t^*}|y\rangle$. Such a t must exist, because if not, we can improve the principal's payoff by slightly lowering $y(\mathbf{b}_s)$; recall that $y(\mathbf{b}_s) > 0$ because $u(\theta = 0, 0, s) < 0$. However, we will argue case-wise that such a t can not exist, producing a contradiction.

Case 1: $t \leq s$.

Observe that $u\langle d_t|y\rangle \geq 0$, (since $u\langle d_t|y\rangle = u\langle d_{t^*}|y\rangle \geq u\langle d_\infty|y\rangle = 0$) and $u\langle \mathbf{b}_{t,t}|y\rangle > 0$, (by *Claim 1*). Thus, by ESTP, $u\langle d_{t-1}|y\rangle > u\langle d_t|y\rangle$. But, since y is feasible for $[\mathcal{P}(t^*)]$, we must have $u\langle d_{t^*}|y\rangle \geq u\langle d_{t-1}|y\rangle$, so that $u\langle d_t|y\rangle = u\langle d_{t^*}|y\rangle$ cannot hold.

Case 2: $s < t < \infty$.

Let r be the first period satisfying $s \leq r < t$, and $y(\mathbf{b}_{r+1}) < y(\mathbf{b}_s)$. If there is no such r, let $r = t$. By the complementary slackness condition in Theorem 1, we have $u\langle d_r|y\rangle = u\langle d_{t^*}|y\rangle$. Now, let $m := u(\theta = 0, y(\mathbf{b}_s), t^*)$. By positive monotonicity, we have $m \geq u\langle \mathbf{b}_{s,s}|y\rangle > 0$. By construction, if the agent stops at t^*, then he is guaranteed at least m from any breakdown occurring between t^* and r. We therefore have:

$$u\langle d_{t^*}|y\rangle \geq \prod_{i=1}^{t^*} p_i \left(\sum_{i=t^*+1}^{r} (1 - p_i) \prod_{j=t^*+1}^{i-1} p_j \right) m + u\langle d_r|y\rangle,$$

contradicting the assertion that $u\langle d_r|y\rangle = u\langle d_{t^*}|y\rangle$, since $m > 0$.

Case 3: $t = \infty$.

Observe that this can only be the case if we have:

$$y(\mathbf{b}_s) = \cdots = y(\mathbf{b}_T) \tag{9}$$

and

$$\frac{v_y(0, y(\mathbf{b}_T), t^*)}{u_y(0, y(\mathbf{b}_T), t^*)} = \frac{v_y(1, y(\mathbf{a}), t^*)}{u_y(1, y(\mathbf{a}), t^*)}, \tag{10}$$

since otherwise one of the complementary slackness results in Theorem 1 would imply $u\langle d_{t'}|y\rangle = u\langle d_{t^*}|y\rangle$ for some $t' < \infty$.

Using (10), Assumption 1 (h) and Assumption 2 (f), we conclude that $y(\mathbf{b}_T) \leq y(\mathbf{a})$. To see this, observe from the FOCs for $y(\mathbf{a})$ and $y(\mathbf{b}_T)$, (19) and (20), that $v_y(0, y(\mathbf{a}), t^*) < 0$ and $v_y(1, y(\mathbf{b}_T), t^*) < 0$. Now, suppose towards contradiction that $y(\mathbf{a}) < y(\mathbf{b}_T)$. By concavity of u and v, we have:

$$\frac{v_y(1, y(\mathbf{a}), t^*)}{u_y(1, y(\mathbf{a}), t^*)} > \frac{v_y(1, y(\mathbf{b}_T), t^*)}{u_y(1, y(\mathbf{b}_T), t^*)}$$

and, by assumptions 1 (h) and 1 (f), we have:

$$\frac{v_y(1, y(\mathbf{b}_T), t^*)}{u_y(1, y(\mathbf{b}_T), t^*)} \geq \frac{v_y(0, y(\mathbf{b}_T), t^*)}{u_y(0, y(\mathbf{b}_T), t^*)};$$

together, these give:

$$\frac{v_y(1, y(\mathbf{a}), t^*)}{u_y(1, y(\mathbf{a}), t^*)} > \frac{v_y(0, y(\mathbf{b}_T), t^*)}{u_y(0, y(\mathbf{b}_T), t^*)},$$

contradicting (10).

Now, let $m := u(\theta = 0, y(\mathbf{b}_s), s-1)$, and observe, by positive monotonicity, that $m > 0$. But then, using (9) the fact that $y(\mathbf{a}) \geq y(\mathbf{b}_T)$ and Assumption 1 (b), we have that an agent who stops at $s-1$ must receive at least m. We therefore have:

$$u\langle d_{s-1}|y\rangle \geq p(\mathbf{a}_{s-1})m,$$

which is greater than $u\langle d_\infty|y\rangle = 0$, since $m > 0$. Since $u\langle d_{t^*}|y\rangle \geq u\langle d_{s-1}|y\rangle$, we have contradicted $u\langle d_{t^*}|y\rangle = u\langle d_\infty|y\rangle$.

□

An immediate consequence of Lemma 4 is that any solution to $[P(t)]$ also satisfies $[P_0(t)]$.

Corollary 1. *Fix $t^* \in T$, and consider any policy, y, which solves $[P(t^*)]$. Then, y also solves $[P_0(t^*)]$.*

Proof. Lemma 4 implies that y is feasible for $[P_0(t^*)]$. Indeed, consider any non-simple stopping rule, d. Let t be the (unique, if it exists) period at at which $d(\mathbf{a}_t) = \texttt{stop}$, or $t = \infty$ if there is no such period. Let $\{\mathbf{b}_{r_i,s_i}\}_i$ denote the set of other histories at which d stops. Then, from equation (8), we have:

$$u\langle d|y\rangle = u\langle d_t|y\rangle + \sum_i p(\mathbf{b}_{r_i,s_i})u\langle \mathbf{b}_{r_i,s_i}|y\rangle;$$

by Lemma 4 we have that the second term is weakly negative, and therefore

$$u\langle d|y\rangle \leq u\langle d_t|y\rangle. \tag{11}$$

Then, since y is feasible for $[P(t^*)]$, we have $u\langle d_t|y\rangle \leq u\langle d_{t^*}|y\rangle$; combined with (11), we conclude that $u\langle d|y\rangle \leq u\langle d_{t^*}|y\rangle$. Since this holds for any d, we have that y is feasible for $[P_0(t^*)]$. Then, since $[P(t^*)]$ is a relaxation of $[P_0(t^*)]$, it must be that y solves $[P_0(t^*)]$. □

Finally, we would like to argue that the principal should optimally induce a simple stopping rule – i.e., that there exists some t^*, such that any solution to $[P_0(t^*)]$ also solves $[P_0]$. This proof is not included in this draft, but we sketch it below. Suppose the principal uses a policy, y, which induces the agent to use a non-simple stopping rule, in which he stops following a breakdown at some time, s, and following no breakdown at time $t^* > s$. The principal can profitably deviate to the policy $\tilde{y}$, which is equal to y, except that at all $t \in (s, t^*]$ it equals $y(s)$. Policy $\tilde{y}$ induces the agent to stop at $s-1$, and is preferable to y for the principal, since it induces the agent to stop earlier, and, conditional on breakdown in $(s, t^*]$, produces a better outcome for the principal than breakdown at s did under y.

A.2 A sufficient condition for APM

We record the following useful class of utility functions (which we call $\beta - \gamma - \delta$ utility functions), and which, as we show below, automatically satisfy APM.

Definition 2. An agent utility function, $u(\theta, y, t)$ is $\beta - \gamma - \delta$ if it can be written:

$$u(\theta, y, t) = \gamma(t)[\beta(t)(w(\theta, y) + \delta(t)) - 1],$$

for some weakly decreasing $\gamma, \beta, \delta : T \to \mathbb{R}_{>0}$, and some $w : \{0, 1\} \times \mathbb{R}_+ \to \mathbb{R}$.

The $\beta - \gamma - \delta$ class is quite broad, and includes the utility functions in all of our applications. It describes an agent who temporally discounts a state-and-policy-dependent "reward" for investment, w, and a cost of investment, normalized to 1, in a way that is consistent with the early sure-thing principle.

Lemma 5 ($\beta-\gamma-\delta$ implies APM)**.** *Suppose that u is $\beta-\gamma-\delta$. Then, u satisfies (APM).*

Proof. Fix some sequence $(\alpha_i, \theta_i, y_i)_i$, with $\alpha_i > 0$ for all i, and suppose that

$$\sum_i \alpha_i u(\theta_i, y_i, t) \geq 0, \tag{12}$$

for some $t \in T$. We would like to show that

$$\sum_i \alpha_i u(\theta_i, y_i, s) \geq \sum_i \alpha_i u(\theta_i, y_i, t),$$

for any $s < t$.

Rewriting (12), we find:

$$\sum_i \alpha_i(\gamma(t)[\beta(t)(w(\theta_i, y_i) + \delta(t)) - 1]) = A\gamma(t)[\beta(t)(X + \delta(t)) - 1] \geq 0, \tag{13}$$

where $A := \sum_i \alpha_i$ and $X := \frac{1}{A}\sum_i \alpha_i w(\theta_i, y_i)$. Then, for any $s < t$, we have:

$$\begin{aligned}
A\gamma(t)[\beta(t)(X + \delta(t)) - 1] &\geq 0 && && (14)\\
X + \delta(t) &\geq 0 && [A, \gamma(t), \beta(t) > 0] && (15)\\
\beta(s)(X + \delta(s)) &\geq \beta(t)(X + \delta(t)) && [\beta, \delta \text{ decr.}] && (16)\\
\beta(s)(X + \delta(s)) - 1 &\geq \beta(t)(X + \delta(t)) - 1 \geq 0 && [(14), (16)] && (17)\\
A\gamma(s)[\beta(s)(X + \delta(s)) - 1] &\geq A\gamma(t)[\beta(t)(X + \delta(t)) - 1] && [\gamma \text{ decr.}] && (18)
\end{aligned}$$

and so:

$$\sum_i \alpha_i u(\theta_i, y_i, s) \geq \sum_i \alpha_i u(\theta_i, y_i, t),$$

as desired. □

A.3 Proof of Lemma 1

Proof of (a). The existence of a solution follows from the Weierstrass Theorem. Suppose that a policy rule y with $y(\mathbf{b}_s) > 0$ for some $s < t$ solves $[\mathcal{P}(t)]$. Consider the policy rule $\tilde{y}$ that satisfies $\tilde{y}(\mathbf{b}_s) = 0$ for all $s < t$ and is otherwise identical to y. By Assumption 2(c), the principal's ex ante expected payoff must be weakly higher under $\tilde{y}$ than y. Furthermore, note that U_t does not depend on the past policies $y(\mathbf{b}_s)$ for $s < t$. Hence, switching the policy rule from y to $\tilde{y}$ does not cause the agent to stop later than t.

Proof of (b). Suppose $\hat{t} \in \arg\max_t V(t)$, and let $\hat{y}$ be a solution to $[\mathcal{P}(\hat{t})]$ that satisfies the condition in (a). We will argue that, whenever $(\hat{y}, \hat{t})$ does not solve $[\mathcal{P}]$, one can find $s < \hat{t}$ such that $(\hat{y}, s)$ solves $[\mathcal{P}]$, $\hat{y}$ solves $[\mathcal{P}(s)]$, and $s \in \arg\max_t V(t)$.

Suppose $(\hat{y}, \hat{t})$ does not solve $[\mathcal{P}]$. This means that $(\hat{y}, \hat{t})$ must be infeasible under $[\mathcal{P}]$. That is, it must that $U_{t'}(\hat{y}) > U_t(\hat{y})$ for some $t' < \hat{t}$. Let $s < \hat{t}$ be such that τ_s is the agent's optimal choice of simple stopping time given y.

Fixing the policy rule at $\hat{y}$, let us compare the principal's expected payoff when the agent plays τ_s, to $V(\hat{t})$, which is the principal's expected payoff when the agent plays $\tau_{\hat{t}}$. We make two observations. First, if the agent plays $\tau_{\hat{t}}$, conditional on a breakdown occurring on or after period $\hat{t} + 1$ (or never at all), the principal's conditional expected payoff must be non-negative; otherwise, the principal would be better off committing to implement the policy $y = 0$ regardless of the state. Then, since the principal's payoff

function satisfies APM, the principal's conditional expected payoff is weakly higher if the agent stops earlier at $s < \hat{t}$. Second, conditional on a breakdown occurring between period s and $\hat{t}$, by Assumption 2(d), the principal weakly prefers that the agent stops. Given these two observations, the principal's payoff is weakly higher when the agent plays τ_s than when he plays $\tau_{\hat{t}}$. However, since we assumed $\hat{t} \in \arg\max_t V(t)$, it must be that the principal's payoff is in fact the same whether the agent plays τ_s or $\tau_{\hat{t}}$. Therefore, $\hat{y}$ solves $[\mathcal{P}(s)]$, and $s \in \arg\max_t V(t)$. Since $\hat{y}$ induces τ_s, it must be that $(\hat{y}, s)$ solves $[\mathcal{P}]$.

Proof of (c). By (b), it must be that the value of $[\mathcal{P}]$ is equal to $\max_t V(t)$. The statement follows.

A.4 Proof of Proposition 1

Under the verifiable investment case, fixing a time t, following the same approach as the proof of Theorem 1, the first order condition with respect to any $y(\mathbf{h})$, where λ is the Lagrange multiplier on the individual rationality constraint, is

$$v_y\left(\theta, y(\mathbf{h}), t\right) + \lambda u_y\left(\theta, y(\mathbf{h}), t\right) = 0.$$

This says that the first order conditions with respect to $y(\mathbf{b}_{t+s})$ for all s are the same, thus define $y(\mathbf{b})$ to be the optimal policy if the final state is 0. Rearranging the first order conditions to substitute out the multiplier to get the equation required.

A.5 Proof of Theorem 1

By Lemma 1(c), it is enough to prove that the statements hold for any y that solves $[\mathcal{P}(t)]$. Recall that this problem is

$$\begin{aligned}
&\max_{y(\cdot)} \mathbb{E}_{\theta,\mathbf{h}}[v(\theta, y(\mathbf{h}), \tau_t)] \\
=&\max_{y}(1-p_{t+1})v(0, y(\mathbf{b}_{t+1}), t) + p_{t+1}(1-p_{t+2})v(0, y(\mathbf{b}_{t+2}), t) + \ldots + \\
&\qquad + p_{t+1}p_{t+2}\cdots p_T v(1, y(\mathbf{a}), t)
\end{aligned}$$

subject to $U_t \geq U_{t+s}$ for all $s = 1, 2, \ldots, T-t$ and $U_t \geq U_\infty = 0$. Let λ_s be the Lagrange multiplier corresponding to the constraint $U_t \geq U_{t+s}$. Then the first order condition with respect to $y(\mathbf{a})$ is

$$v_y(1, y(\mathbf{a}), t) + \sum_{r=1}^{T-t} \lambda_r \left[u_y\left(1, y(\mathbf{a}), t\right) - u_y\left(1, y(\mathbf{a}), t+r\right)\right] + \lambda_\infty u_y\left(1, y(\mathbf{a}), t\right) = 0 \qquad (19)$$

and the first order condition with respect to each $y(\mathbf{b}_{t+s})$ for $s = 1, 2, \ldots, T-t$, is

$$v_y(0, y(\mathbf{b}_{t+s}), t) + \sum_{r=1}^{s-1} \lambda_r \left[u_y\left(0, y(\mathbf{b}_{t+s}), t\right) - u_y\left(0, y(\mathbf{b}_{t+s}), t+r\right)\right] + \\ + \sum_{r=s}^{T-t} \lambda_r u_y\left(0, y(\mathbf{b}_{t+s}), t\right) + \lambda_\infty u_y\left(0, y(\mathbf{b}_{t+s}), t\right) = 0 \quad (20)$$

Next we show that $y(\mathbf{b}_{t+s}) \geq y(\mathbf{b}_{t+s+1})$ for all $s = 1, 2, \ldots, T-t$. Compare the first order conditions with respect to $y(\mathbf{b}_{t+1})$ and $y(\mathbf{b}_{t+2})$.

$$\begin{cases} v_y(0, y(\mathbf{b}_{t+1}), t) + \left(\sum_{r=1}^{T-t} \lambda_r + \lambda_\infty\right) u_y(0, y(\mathbf{b}_{t+1}), t) = 0 \\ v_y(0, y(\mathbf{b}_{t+2}), t) + \left(\left(1 - \frac{u_y(0, y(\mathbf{b}_{t+2}), t+1)}{u_y(0, y(\mathbf{b}_{t+2}), t)}\right) \lambda_1 + \sum_{r=2}^{T-t} \lambda_r + \lambda_\infty\right) u_y(0, y(\mathbf{b}_{t+2}), t) = 0. \end{cases}$$

Suppose towards a contradiction that $y(\mathbf{b}_{t+1}) < y(\mathbf{b}_{t+2})$. Then $v_y(0, y(\mathbf{b}_{t+1}), t) > v_y(0, y(\mathbf{b}_{t+2}), t)$ from v_y decreasing in y, thus

$$\left(\sum_{r=1}^{T-t} \lambda_r + \lambda_\infty\right) u_y(0, y(\mathbf{b}_{t+1}), t) \\ < \left(\left(1 - \frac{u_y(0, y(\mathbf{b}_{t+2}), t+1)}{u_y(0, y(\mathbf{b}_{t+2}), t)}\right) \lambda_1 + \sum_{r=2}^{T-t} \lambda_r + \lambda_\infty\right) u_y(0, y(\mathbf{b}_{t+2}), t),$$

but since u_y is decreasing in y, we have $u_y(0, y(\mathbf{b}_{t+1}), t) \geq u_y(0, y(\mathbf{b}_{t+2}), t)$, further, our assumption says that $\frac{u_y(0, y(\mathbf{b}_{t+2}), t+1)}{u_y(0, y(\mathbf{b}_{t+2}), t)} \leq 1$, which means that

$$\left(\sum_{r=1}^{T-t} \lambda_r + \lambda_\infty\right) \geq \left(\left(1 - \frac{\partial_y u(0, y(\mathbf{b}_{t+2}), t+1)}{\partial_y u(0, y(\mathbf{b}_{t+2}), t)}\right) \lambda_1 + \sum_{r=2}^{T-t} \lambda_r + \lambda_\infty\right),$$

contradicting the inequality above.

The argument is the same for all other $y(\mathbf{b}_{t+s}) \geq y(\mathbf{b}_{t+s+1})$, but the inequality $y(\mathbf{b}_T) \geq y(\mathbf{a})$ does not necessarily hold because θ changes from 0 to 1 between the histories $\mathbf{b}_T$ and $\mathbf{a}$, thus θ is no longer held fixed in v_y and $\partial_y u$.

For the second part of the theorem, suppose $y(\mathbf{b}_{t+s}) = y(\mathbf{b}_{t+s+1})$. By comparing the first order conditions with respect to $y(\mathbf{b}_{t+s})$ and $y(\mathbf{b}_{t+s+1})$, we get

$$v_y(0, y(\mathbf{b}_{t+s}), t) + \sum_{r=1}^{s-1} \lambda_r \left(u_y(0, y(\mathbf{b}_{t+s}), t) - u_y(0, y(\mathbf{b}_{t+s}), t+r)\right) + \\ + \sum_{r=s}^{T-t} \lambda_r u_y(0, y(\mathbf{b}_{t+s}), t) + \lambda_\infty u_y(0, y(\mathbf{b}_{t+s}), t) = 0$$

$$v_y(0, y(\mathbf{b}_{t+s+1}), t) + \sum_{r=1}^{s} \lambda_r \left(u_y(0, y(\mathbf{b}_{t+s+1}), t) - u_y(0, y(\mathbf{b}_{t+s+1}), t+r)\right) + \\ + \sum_{r=s+1}^{T-t} \lambda_r u_y(0, y(\mathbf{b}_{t+s+1}), t) + \lambda_\infty u_y(0, y(\mathbf{b}_{t+s+1}), t) = 0$$

Since $y(\mathbf{b}_{t+s}) = y(\mathbf{b}_{t+s+1})$, all the terms in the two FOCs are the same, except for the coefficient of λ_s. Hence, using $u_y > 0$,

$$\lambda_s u_y(0, y(\mathbf{b}_{t+s}), t) = \lambda_s \left(u_y(0, y(\mathbf{b}_{t+s}), t) - u_y(0, y(\mathbf{b}_{t+s}), t+s)\right),$$

we conclude $\lambda_s = 0$, that is, the constraint $U_t \geq U_{t+s}$ is slack.

Conversely, suppose $y(\mathbf{b}_{t+s}) > y(\mathbf{b}_{t+s+1})$, then $v_y(0, y(\mathbf{b}_{t+s}), t) < v_y(0, y(\mathbf{b}_{t+s+1}), t)$. The first order conditions imply

$$\sum_{r=1}^{s-1} \lambda_r \left(u_y(0, y(\mathbf{b}_{t+s}), t) - u_y(0, y(\mathbf{b}_{t+s}), t+r)\right) + \\ + \sum_{r=s}^{T-t} \lambda_r u_y(0, y(\mathbf{b}_{t+s}), t) + \lambda_\infty u_y(0, y(\mathbf{b}_{t+s}), t) \\ > \\ \sum_{r=1}^{s} \lambda_r \left(u_y(0, y(\mathbf{b}_{t+s+1}), t) - u_y(0, y(\mathbf{b}_{t+s+1}), t+r)\right) + \\ + \sum_{r=s+1}^{T-t} \lambda_r u_y(0, y(\mathbf{b}_{t+s+1}), t) + \lambda_\infty u_y(0, y(\mathbf{b}_{t+s+1}), t) \quad (21)$$

Note that $u_y(0, y(\mathbf{b}_{t+s}), t) \leq u_y(0, y(\mathbf{b}_{t+s+1}), t)$ by concavity of u, so by comparing the second and the third term of both sides we have

$$\sum_{r=s}^{T-t} \lambda_r u_y(0, y(\mathbf{b}_{t+s}), t) + \lambda_\infty u_y(0, y(\mathbf{b}_{t+s}), t) \leq \\ \leq \sum_{r=s+1}^{T-t} \lambda_r u_y(0, y(\mathbf{b}_{t+s+1}), t) + \lambda_\infty u_y(0, y(\mathbf{b}_{t+s+1}), t).$$

Further, by supermodularity of u_y, we have that for all $r \leq s-1$,

$$u_y(0, y(\mathbf{b}_{t+s}), t) - u_y(0, y(\mathbf{b}_{t+s}), t+r) \leq u_y(0, y(\mathbf{b}_{t+s+1}), t) - u_y(0, y(\mathbf{b}_{t+s+1}), t+r).$$

Thus the only way we could have inequality (21) is if

$$\lambda_s u_y(0, y(\mathbf{b}_{t+s}), t) > \lambda_s \left[u_y(0, y(\mathbf{b}_{t+s+1}), t) - u_y(0, y(\mathbf{b}_{t+s+1}), t+s)\right],$$

and because of our assumption that u_y is decreasing in time all else equal, the right hand side must be positive. This says that $\lambda_s > 0$, and the constraint $U_t \geq U_{t+s}$ binds.

A.6 Proof of Proposition 2

We first show that the constraints $U_t \geq U_{t+r}$ for $r \in \{1, 2, \ldots, T-t-1\}$ do not bind. Suppose otherwise, and $U_t = U_{t+r}$ for some r. Then from $U_t(y) \geq \max\{U_{t+1}(y), \ldots, U_T(y), 0\}$ we have $U_{t+r-1} \leq U_{t+r}$. This simplifies to

$$(1-p)y(\mathbf{b}_{t+r}) \leq \frac{pg(t+r) - g(t+r-1)}{\delta^T},$$

where $g(t) := \delta^t \left(\frac{s}{1-\delta} - I\right) - \frac{s\delta^T}{1-\delta}$. Consider the right-hand side of the inequality. The mapping $t \mapsto \frac{pg(t+1)-g(t)}{\delta^T}$ is equal to

$$\frac{p\delta - 1}{\delta^{T-t}}\left(\frac{s}{1-\delta} - I\right) + (1-p)\frac{s}{1-\delta},$$

and note that this is strictly increasing in t if we have $\delta < 1$. This means that

$$\frac{pg(t+r) - g(t+r-1)}{\delta^T} < \frac{pg(t+r+1) - g(t+r)}{\delta^T}.$$

Further, by Theorem 1, we have $y(\mathbf{b}_{t+r}) \geq y(\mathbf{b}_{t+r+1})$, leading to

$$(1-p)y(\mathbf{b}_{t+r+1}) < \frac{pg(t+r+1) - g(t+r)}{\delta^T}.$$

Which is what the inequality $U_{t+r} < U_{t+r+1}$ simplifies to, and this says $U_t < U_{t+r+1}$, contradiction.

Next we show that the constraint $U_t \geq U_T$ can not be the only one that binds. Towards a contradiction, suppose $U_T = U_t > 0$. Then

$$U_T = p_1 \cdots p_T u(1, y(\mathbf{a}), T) = p_1 \cdots p_T \delta^T (y(\mathbf{a}) - I) > 0.$$

To solve for the optimal $y(\mathbf{a})$, following the approach in Theorem 1, the first order condition with respect to $y(\mathbf{a})$, with the principal having a quadratic cost function, is $y(\mathbf{a}) = \lambda_\infty$, where λ_∞ is the Lagrange multiplier on the constraint $U_t \geq 0$. Using $U_t > 0$ we conclude that this constraint is slack, and thus $\lambda_\infty = 0$. This is a contradiction because $y(\mathbf{a}) - I = -I < 0$.