Incentivizing Exploration and Stopping*

Christopher Almacen†

October 22, 2026

Job Market Paper

(Latest Version Here)

Abstract

We study a dynamic principal-agent problem where a principal seeks to acquire information about a project with uncertain returns and contracts with an agent to conduct the experiments. While the principal (she) commits to flexibly designing the experiments and rewards, the agent (he) can stop at any time and is constrained by his technology, which limits how much information he can generate. We show that there exists a *simple* optimal contract: the principal acquires information via breakthroughs, offers a deterministic and nondecreasing transfer schedule over time, and recommends stopping after any breakthrough realization. This result is driven by the players' exponential discounting and differing risk-attitudes over money. Moreover, it provides a theoretical foundation for restricting attention to a simple class of contracts.

*I am grateful to my advisor Hector Chade and to my committee members Mark Whitmeyer and Ilia Krasikov. I would like to thank Ludovic Renou, Eddie Schlee, Théo Durandard, Kun Zhang, Aniruddha Ghosh, Martín Grunwald, and the seminar participants at Arizona State University, the 37th Stony Brook International Conference on Game Theory, University of Arizona, and Cal Poly San Luis Obispo for the helpful comments and suggestions. This project has benefited from large language models. All errors are my own.

†Arizona State University, Department of Economics (email: calmacen@asu.edu)

1 Introduction

Incentivizing experimentation and innovation is central in many economic settings. For example, the National Institutes of Health (NIH) and National Science Foundation (NSF) allocate grants to principal investigators. The U.S. Food and Drug Administration (FDA) designs clinical trials to incentivize pharmaceutical firms to develop new drugs. Lastly, owners delegate to managers on how to invest in risky assets.

All of the above examples share a common theme: these are risk-sharing environments where a principal contracts with an agent to investigate a risky opportunity. An agent (he) controls whether and when to undertake a risky project, while a principal (she) wants to learn more about its value. A key and realistic feature is that when the principal writes a contract, she may have considerable flexibility in designing both the experiments and rewards. For instance, the NIH and NSF oversee both research progress and funding for researchers. The FDA specifies clinical-trial requirements and even allocates subsidies.[1] Lastly, owners specify business strategies and compensation for managers. To elaborate, when designing experiments, the principal usually has a great degree of freedom in specifying what signals to acquire over time, which statistical tests to conduct, and the overall sequential sampling process. Moreover, she can reward the agent in multiple ways. She can fund the agent with fixed payments or offer prizes depending on signal realizations.

However, experimentation is inherently dynamic, costly, and capacity constrained. It takes time to learn about a project with uncertain returns. This may be due to the agent's available technology, which limits how quickly he can generate information. Furthermore, the agent may be unable to commit in these long-term relationships. In particular, if continuing to experiment is too costly, he may opt to walk away. Hence, when the principal designs her contract, she must also recommend how and when the agent should stop experimenting. These frictions lead us to our research question: *How would the principal flexibly design an optimal contract subject to the agent's dynamic participation and capacity-constrained technology?*

To answer this question, we build a stylized continuous-time principal-agent model with two projects: one risky and one safe. The risky project can be in a good or bad state, and the agent has the decision right to accept or reject it. The safe project yields a normalized payoff. For the risky project, the principal cares about the agent's choice: both players prefer acceptance in the good state and rejection in the bad state, but they value these outcomes differently. Consequently, the common posterior belief at the time

[1]The FDA may allocate grants for orphan products to encourage the development of medical products for rare diseases or conditions (https://www.fda.gov/industry/orphan-products-grants-program/clinical-trials-grants-program). For more examples, see Yoder (2022).

of stopping determines the agent's decision and both players' terminal payoffs. Before experimentation begins, the principal designs and commits to a dynamic contract which is made up of three ingredients: an exploring rule, a transfer rule, and a stopping rule. The exploring rule represents dynamic experiments that the agent must conduct. We study a symmetric-information environment in which experiments produce hard information: the principal and the agent observe both the experiments and the signal realizations. The agent's technology limits the rate at which information can be generated. The transfer rule specifies the monetary incentives provided by the principal, and the stopping rule specifies when she recommends that the agent stop experimenting. Because the agent lacks commitment, he may leave the contract at any time.

We emphasize that the principal's contract space is rich and complicated. For instance, she can choose lotteries over rewards and history-dependent experiments and stopping times. And yet, our main result, Theorem 3.1, states that there exists a simple optimal contract.[2] There exists an optimal solution such that the principal chooses a piecewise-deterministic jump process as the exploring rule and a deterministic transfer rule that is weakly increasing over time. The stopping rule is such that if the players observe a breakthrough, then the principal recommends that the agent to stop immediately.

This result is significant because it provides a theoretical foundation for both the strategic experimentation and innovation literatures. With respect to the former literature, the majority of theoretical models take the information structure of experiments as *exogenously* given. In this paper, the principal *endogenously* designs the experiments. Thus, an important consequence of Theorem 3.1 is that it gives one reason why experiments would draw signals in a breakthrough fashion. The latter literature asked both positive and normative questions for why certain funding structures exist. For instance, should research be funded with grants, prizes or patents?[3] This paper contributes to this. To the best of my knowledge, this is the first paper to study a contracting problem where the principal jointly and flexibly designs the dynamic experiments and reward structure. In particular, Theorem 3.1 shows that funding with a deterministic payment schedule during experimentation can indeed be an optimal choice.

Next, we describe what model assumptions drive Theorem 3.1. The intuition for the exploring rule comes from the convexity of the players' exponential discounting. That is, both players have *risk-loving* time preferences. Because the exponential discount factor is convex in delay, both players love lotteries over time. The jump process is the continuous-

[2]When we say optimal, we mean the principal's expected utility maximizing contract.

[3]See Bryan and Williams (2021) and Azoulay and Li (2020) for surveys on innovation and scientific funding.

time analog; the principal designs experiments so as to wait for an infrequent but precise breakthrough. The key distinction is that the jump process also delivers enough promise value to the agent to continue experimenting. Meanwhile, the deterministic transfer rule reflects the players' differing attitudes over money. That is, although the principal can feasibly choose history-dependent stochastic lotteries, she can optimally reward the agent with a deterministic flow of money. Intuitively, since the principal is risk-neutral whereas the agent is risk-averse, if the principal were to offer a stochastic lottery, she can weakly improve by offering the certainty equivalent. The only incentive constraint she faces is the agent's dynamic participation; hence, the agent's continuation value also weakly increases with the certainty equivalent. This intuition describes why the principal weakly improves with deterministic transfers. However, Theorem 3.1 states that this improvement is also feasible. Lastly, the transfer rule is weakly increasing transfer because of the agent's dynamic participation constraint. If this constraint is binding, then the principal backloads rewards and thereby encourages the agent to continue experimenting.

Section 4 derives properties of the optimal contract and studies an extension of the model. The main result (Theorem 3.1) only describes the existence of a certain class of optimal contracts. In particular, it only specifies that designing experiment via breakthroughs are optimal. Proposition 4.1 states how the exploring rules are shaped; the principal faces a precision-frequency trade-off. Dynamic experiments via breakthroughs are shaped by two features: the precision of the signal and its frequency rate. The principal chooses the precision of the signal, which influences the players' posterior belief about the risky project. On the other hand, she must also choose the intensity rate at which she observes the signal. Limited by the agent's technology, she is forced to trade off these quantities. We also study a modified problem with an infinite horizon and perfectly patient players. Under these assumptions, an optimal contract again features an exploring rule that is a jump process and a constant transfer rule (Proposition 4.2). Relative to the main model, discounting plays a key role in predicting a transfer rule which backloads rewards.

Lastly, we outline the proof of Theorem 3.1 in Section 5. We do not use dynamic programming or Hamilton–Jacobi–Bellman techniques to solve the principal's problem. Instead, inspired by the methodologies of Sannikov and Zhong (2026) and Koh, Sanguanmoo, and Zhong (2026), we relax the principal's problem and optimize over the probability law of stopped belief-time pairs and deterministic transfers. After deriving necessary constraints for both capacity and the agent's dynamic participation, we show that our problem is amenable to strong duality and obtain necessary first-order conditions. We then use these conditions to establish the optimality of breakthrough experiments and

explicitly construct an optimal contract.

The paper proceeds as follows. Subsection 1.1 surveys the related literature. Section 2 presents the model and its assumptions. Section 3 states our main result. Section 4 examines properties of the optimal contract and a modified problem. Lastly, Section 5 outlines the proof for our main result. Detailed arguments and proofs appear in the Appendix.

1.1 Related Literature

This paper builds upon the classical sequential sampling model as initiated by the seminal works of Wald (1945) and Arrow, Blackwell, and Girshick (1949). Two important literatures extend this model: strategic experimentation and dynamic information acquisition. This paper contributes to both.

First, this paper contributes to the vast literature of both strategic information acquisition and experimentation.[4] Many of these papers study a dynamic principal-agent framework with moral hazard (e.g., Chade and Kovrijnykh (2016)), hidden information (e.g., Guo (2016), Grenadier, A. Malenko, and N. Malenko (2016), and McClellan (2022)), or both (e.g., Halac, Kartik, and Liu (2016)). Other papers study a strategic setting where the players interact noncooperatively (e.g., Keller, Rady, and Cripps (2005), Keller and Rady (2010), Keller and Rady (2015)), have differing costs in experimentation (e.g., Manso (2011)), have differing control rights (e.g., Henry and Ottaviani (2019)), or may covertly hide signals (e.g., Herresthal (2022)). In this paper, although the players have symmetric information, the main distinction is that the principal can design the dynamic experiments. That is, all of the aforementioned papers assume that the signal structure *exogenously* draws its signal realizations. Only the agent's technology constrains the rate of information acquisition. By allowing such flexibility, our main result provides a benchmark for these strategic games. We mention that Yoder (2022) studies a static mechanism design setting where the designer can choose both the experiments and transfers. In contrast, we study a dynamic setting and the principal chooses a dynamic Blackwell experiment.

Second, this paper also contributes to the dynamic information acquisition literature. Zhong (2022) studies a decision-maker's problem who can flexibly acquire information and choose the stopping time subject to a flow cost. Hébert and Woodford (2023) and Sannikov and Zhong (2026) study a similar problem except that the decision maker faces a flow constraint. Lastly Georgiadis-Harris (2026) studies the decision maker's problem

[4]Early references on strategic information acquisition include Lambert (1986) and Iossa and Stroffolini (2005).

where who can flexibly choose the dynamic information structure but takes the stopping time as exogenous. It is important to note that these papers also predict a jump process as an optimal solution. In section 3, we elaborate on how our main result, Theorem 3.1, compares with the results of these papers. Che and Mierendorff (2019) studies a related problem where the decision maker must allocate time to different information sources. Among the closest papers to ours is Hébert and Zhong (2026); they study a strategic interaction where the principal designs a dynamic experiment to keep the agent engaged. Model-wise, our paper has two key distinctions. First, the principal's preferences are time-dependent and risk-loving over time. Second, the principal can use transfers to incentivize the agent. Together, these features yield novel predictions for the joint design of experiments and transfers.

Lastly, this paper is related to the dynamic Bayesian persuasion and information design literature. Many of these papers study a strategic interaction where the sender can flexibly acquire information. Ball (2023) and Zhao (2025) study a strategic interaction where the sender can flexibly disclose messages to the receiver and where the unknown state evolves stochastically over time. Zhao et al. (2024) and Koh, Sanguanmoo, and Zhong (2026) study a dynamic principal-agent problem where the principal only controls information to influence the agent. In our model, the principal also controls information flexibly, but the agent's technology constrains its rate of generation; moreover, the principal can use transfers to provide incentives.

2 Setup

We study a dynamic principal-agent problem with two projects: a risky one and a safe one. The principal (she) designs a contract for the agent (he) to flexibly and dynamically acquire information about the risky project. The principal can use three choice variables to incentivize the agent: an *exploring rule*, a *stopping rule*, and a *transfer rule*. While she can fully commit to the contract, the agent can opt to prematurely end the process.

First, we establish some notation. We denote stochastic processes by angle brackets while deterministic paths and histories are denoted by parentheses. Concretely, given filtered probability space $(\Omega, \mathcal{F}, \langle \mathcal{F}_t \rangle, \mathcal{P})$ over index set I and subset $E \subseteq \mathbb{R}^n$, we let $\langle X_t \rangle$ denote an E-valued stochastic process. Given $\omega \in \Omega$, we let $(X_s)_{s \leq t} \equiv X_t(\omega)$ denote the realized history of X_t up to time $t \in I$.

Environment. Time is continuous with a finite horizon $T \in (0, +\infty)$. The safe project has deterministic value. On the other hand, the risky project has uncertain returns and is described by two states of the world: good state G or bad state B. We model the players'

belief about this with the pair $\mu = (\mu_G, \mu_B) \in \Delta^1$ where $\Delta^1 \equiv \{(\mu_G, \mu_B) \in \mathbb{R}^2_+ : \mu_G + \mu_B = 1\}$ denotes the (closed) unit simplex. We fix $\mu_0 \in \text{ri } \Delta^1$ to be the common prior belief of the risky project's quality.[5] In the same way, let $\mu_t \in \Delta^1$ denote the common posterior belief at time $t \in [0, T]$. The agent has the decision right to accept or reject the risky project. Let $X = \{0, 1\}$ denote the set of the agent's feasible actions. The action $x = 1$ means accept while $x = 0$ means reject. The principal can incentivize the agent with rewards; let $\psi_t \in \mathbb{R}_+$ denote a flow of transfer from the principal to the agent at time t.

Preferences. The principal and agent may have misaligned preferences in three different ways: stopping utility, flow costs, and imperfectly transferable utilities. We start by describing how the players value the projects. Both players value the safe project deterministically; we normalize this to zero. For the risky project, they both prefer accepting in the good state and rejecting in the bad state. However, they may differ in the relative value, and hence, they may disagree at which belief to accept the risky project. Formally, conditional on accepting, let $v_G, u_G > 0$ denote the principal and agent's values in the good state (respectively) and $v_B, u_B < 0$ denote their values in the bad state (respectively). For belief $\mu = (\mu_G, \mu_B) \in \Delta^1$, the indirect utilities for the agent and principal are

$$\left.\begin{aligned} U_A(\mu) &\equiv \max_{x \in X} \left\{x \cdot \left(\mu_G \cdot u_G + \mu_B \cdot u_B\right) + (1 - x) \cdot 0\right\} \\ x^*(\mu) &\in \arg\max_{x \in X} \left\{x \cdot \left(\mu_G \cdot u_G + \mu_B \cdot u_B\right) + (1 - x) \cdot 0\right\} \end{aligned}\right\} \quad \text{(Agent)}$$

$$U_P(\mu) \equiv x^*(\mu) \cdot (\mu_G \cdot v_G + \mu_B \cdot v_B) + (1 - x^*(\mu)) \cdot 0 \quad \text{(Principal)}$$

If the players are at a belief where the agent is indifferent between accepting and rejecting, we assume the agent takes the principal's preferred action. In terms of our model, we assume the principal's indirect utility $U_P(\cdot)$ is upper-semicontinuous. We formally record this assumption as follows.

Assumption 1. The principal's indirect utility $U_P : \Delta^1 \to \mathbb{R}$ is upper-semicontinuous.

Both players exponentially discount time and share a common discount factor $r \in (0, +\infty)$. To model the conflicting cost in experimentation, we assume the agent, and only the agent, incurs flow costs $c > 0$. Lastly, the players' utilities are imperfectly transferable because the principal and agent differ in their risk attitudes over money. We assume that the principal is risk-neutral over money while the agent is strictly risk-averse. Let utility function $u : \mathbb{R}_+ \to \mathbb{R}$ represent the agent's preferences over transfers. We formally impose the following structure:

[5] For subset $Z \subseteq \mathbb{R}^n$, we let ri Z denote the relative interior of Z.

Assumption 2. Utility function $u : [0, +\infty) \to \mathbb{R}$ satisfies the following:

(i) *Smoothness:* Twice continuously differentiable on $[0, +\infty)$;

(ii) *Monotonicity:* $u'(\psi) \geq 0$ for all $\psi \in [0, +\infty)$ and $u'(\psi) > 0$ if $\psi > 0$;

(iii) *Risk Aversion:* $u''(\psi) < 0$ for all $\psi \in [0, +\infty)$, and;

(iv) *Surjectivity:* $u(0) = 0$ and $u(\psi) \to +\infty$ as $\psi \to +\infty$.

Given belief $\mu \in \Delta^1$, time $t \in [0, T]$ and path $(\psi_t)_{s \leq t}$, the principal and agent's expected utilities are

$$e^{-r\tau} U_P(\mu) - \int_0^\tau e^{-rs} \psi_s ds \qquad \text{and} \qquad e^{-r\tau} U_A(\mu) + \int_0^\tau e^{-rs}(u(\psi_s) - c) ds.$$

Contracts. An important feature of this model is that we allow the principal a great degree of freedom in designing experiments. Specifically, she has the flexibility to choose both the sample space and information structure. In the language of the Information Acquisition literature, she chooses a Blackwell experiment. She can also use lotteries over transfers. Given the dynamic context, this means that a contract is a complete contingent plan which maps every history realization of signals and transfers to a Blackwell experiment, a flow of transfer (potentially randomized), and a binary stopping decision. However, notice that the players' expected utilities depend only on the realization of the posterior belief. Following the Bayesian persuasion literature, we do not explicitly model the signal realizations nor information structures. Instead, we track the evolution of the posterior beliefs.

We dive into the formal details of a contract. Let Ω be the set of càdlàg paths mapping from $[0, T]$ into $\Delta^1 \times \mathbb{R}_+$. Both players are Bayesian, and hence, update their beliefs via Bayes rule. The belief paths start from prior belief $\mu_0 \in \Delta^1$. A contract is a filtered probability space $(\Omega, \mathcal{F}, \langle \mathcal{F}_t \rangle, \mathcal{P})$ (over the time horizon $[0, T]$) along with a tuple $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$ where: (i) the first coordinate $\langle \mu_t \rangle$ is a martingale — a Δ^1-valued adapted stochastic process satisfying $\mu_s = \mathbb{E}_{\mathcal{P}}[\mu_t \mid \mathcal{F}_s]$ almost surely for $s \leq t$; (ii) the second coordinate $\langle \psi_t \rangle$ is a nonnegative real-valued progressively measurable stochastic process;[6] and, (iii) the third coordinate τ is an optional-stopping time with respect to filtration $\langle \mathcal{F}_t \rangle$. Whenever the context is clear, we will refer to the tuple $(\langle \mu_t \rangle, \langle \psi \rangle, \tau)$ as the contract, omitting the underlying filtered probability space. Given contract $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$, the principal and agent's expected utilities are, respectively, given by

$$\mathbb{E}_{\mathcal{P}} \left[e^{-r\tau} U_P(\mu_\tau) - \int_0^\tau e^{-rs} \psi_s ds \,\middle|\, \mathcal{F}_0 \right] \qquad \text{(Principal)}$$

[6] Process $\langle \psi_t \rangle$ is progressively measurable if and only if its restriction to $\Omega \times [0, t]$ is $\mathcal{F}_t \otimes \mathcal{B}([0, t])$-measurable for all $t \in [0, T]$.

$$\mathbb{E}_{\mathcal{P}}\left[e^{-r\tau}U_A(\mu_\tau) + \int_0^\tau e^{-rs}(u(\psi_s) - c)ds \,\middle|\, \mathcal{F}_0\right] \tag{Agent}$$

Constraints. Although the principal can flexibly design experiments and use transfers to incentivize the agent, she is constrained in three ways. First, the principal is subject to both limited liability and limited discretion. We record the formal assumption below.

Assumption 3. Given contract $(\langle\mu_t\rangle, \langle\psi_t\rangle, \tau)$ with respective filtered probability space $(\Omega, \mathcal{F}, \langle\mathcal{F}_t\rangle, \mathcal{P})$, transfer rule $\langle\psi_t\rangle$ satisfies the following conditions:

(i) *Limited Liability:* $0 \leq \psi_t$ almost almost everywhere for all time $t \in [0, T]$;

(ii) *Limited Discretion:* there exists an upper bound $\overline{\psi} \in \mathbb{R}_+$ such that $u(\overline{\psi}) > c + rU_A(\mu_0)$ and $\psi_t \leq \overline{\psi}$ almost everywhere for all time $t \in [0, T]$.

Intuitively, the limited liability assumption ensures that the principal cannot punish the agent by imposing fees. The limited discretion assumption inhibits how much the principal can reward the agent. These assumptions are realistic in many practical settings. For example, the NIH and NSF tend only fund principal investigators and are budget constrained in allocating grants. In terms of the model, these assumptions play a key role in ensuring the existence of a solution and satisfying a constraint qualification condition.[7]

The second key constraint is that the agent is constrained in how quickly he can generate information. To model the agent's technology, let $\chi \in (0, +\infty)$ and $H : \Delta^1 \to \mathbb{R}$ a real-valued mapping. Given history $(\mu_s, \psi_s)_{s \leq t}$, exploring rule $\langle\mu_t\rangle$ satisfies the variation constraint if and only if

$$\mathbb{E}_{\mathcal{P}}\big[H(\mu_{t'}) - H(\mu_t) \,|\, \mathcal{F}_t\big] \leq \chi \cdot (t' - t), \quad \forall t, t' \in [0, T] \text{ and } t < t' \tag{Var}$$

where $\mathcal{F}_t$ is the σ-algebra generated by (μ_s, ψ_s) for all $s \leq t$. Intuitively, $H(\cdot)$ measures how much the players' posterior belief can vary. Hence, the left-hand side of (Var) measures how much their belief changes from belief μ_t to $\mu_{t'}$. This "variation" is in part limited by capacity χ conditional on acquiring information over the time interval $[t, t']$. We impose the following additional structure on $H(\cdot)$:

Assumption 4. Flow mapping $H : \Delta^1 \to \mathbb{R}$ satisfies the following properties:

(i) *Smoothness:* twice continuously differentiable on Δ^1;

(ii) *Boundedness:* $H'(\cdot)$ is uniformly bounded on Δ^1;

(iii) *Convexity:* $H'' > 0$ on ri Δ^1, and;

[7] See Lemmas 5.2 and A.3.

Part (i) of Assumption 4 ensures that the model is tractable to analyze. The boundedness assumption excludes flow mappings which satisfy "Inada conditions" at the boundaries of the unit simplex.[8] The convexity assumption is key in that the left-hand side of (Var) provides a uniformly posterior-separable measure of information.[9]

Lastly, the third constraint describes our main strategic friction. The agent can deviate by prematurely ending the exploring process. Hence, the principal must *recommend* a stopping rule. Given history $(\mu_s, \psi_s)_{s\leq t}$, contract $(\langle\mu_t\rangle, \langle\psi_t\rangle, \tau)$ satisfies the dynamic participation constraint if and only if

$$\mathbb{E}_{\mathcal{P}}\left[e^{-r(\tau-t)}U_A(\mu_\tau) + \int_t^\tau e^{-r(s-t)}(u(\psi_s)-c)ds \,\middle|\, \mathcal{F}_t\right] \geq U_A\big(\mathbb{E}_{\mathcal{P}}[\mu_\tau \mid \mathcal{F}_t]\big),\ \forall t\in[0,T] \quad \text{(DPC)}$$

where $\tau > t$. The left-hand side of (DPC) represents the agent's continuation value at time t from following the principal's recommendation. The right-hand side represents the agent's stopping value at time t. Thus, an incentive compatible stopping rule ensures that the agent weakly prefers to continue experimenting.

Overall, the principal's full problem is to solve

$$V_{FP}(\mu_0) = \sup_{\langle\mu_t\rangle,\langle\psi_t\rangle,\tau} \mathbb{E}_{\mathcal{P}}\left[e^{-r\tau}U_P(\mu_\tau) - \int_0^\tau e^{-rs}\psi_s ds\right] \quad \text{subject to} \quad \text{(Var)}, \text{(DPC)} \qquad \text{(FP)}$$

We remark that the players have symmetric information in this strategic interaction. Experiments and its signal realizations are hard information and observable by both players. In particular, the agent cannot covertly acquire information or hide signal realizations. He also does not hold private information. Although these are realistic frictions, we argue that it is important to study a model without them. First, we can interpret our model as the symmetric information benchmark. If we considered a model with asymmetric information, we would be interested in its respective distortions (say in the players' expected utility or the optimal contract). However, to understand its distortions, we would need a well-defined benchmark. This paper provides that.

Lastly, we impose an assumption to make the analysis and proofs more tractable.

Assumption 5. We extend the flow mapping $H(\cdot)$ and the agent's indirect utility $U_A(\cdot)$ in a homogeneous of degree one manner to the positive cone $\mathbb{R}^2_+$. Abusing and reusing

[8] Examples of flow mappings which satisfy Assumption 4 include the quadratic variation $H(\mu_G,\mu_B) = 0.5\cdot(\mu_G^2+\mu_B^2)$ and the Tsallis entropy $H(\mu_G,\mu_B) = \frac{-1}{q-1}\cdot(1-\mu_G^q-\mu_B^q)$ for $q>1$. A nonexample is the Shannon entropy $H(\mu_G,\mu_B) = -\mu_G\log\mu_G - \mu_B\log\mu_B$, which has unbounded derivative near the boundaries.

[9] This way of measuring information is common in both the rational inattention and dynamic information acquisition literature. For example, see Caplin, Dean, and Leahy (2022), Zhong (2022), Hébert and Woodford (2023), and Georgiadis-Harris (2026).

notation, we let $H : \mathbb{R}^2_+ \to \mathbb{R}$ and $U_A : \mathbb{R}^2_+ \to \mathbb{R}$ denote these extensions.

For an explicit construction, if the agent's indirect utility is $\widetilde{U}_A : \Delta^1 \to \mathbb{R}$, then given $\mu = (\mu_1, \mu_2) \in \mathbb{R}^2_+$, the homogeneous degree one extension of $\widetilde{U}_A(\cdot)$ is defined by $U_A(\mu) = 0$ if $\mu_1 + \mu_2 = 0$ and $U_A(\mu) = (\mu_1 + \mu_2) \cdot \widetilde{U}_A\left(\frac{\mu_1}{\mu_1+\mu_2}, \frac{\mu_2}{\mu_1+\mu_2}\right)$ if $\mu_1 + \mu_2 > 0$. This is without loss of generality since $\widetilde{U}_A(\cdot) = U_A(\cdot)$ on Δ^1.

3 Main Result

We emphasize that the principal can flexibly design both the experiments and the flow of transfers. Despite the rich and complex contract space, our main result states that there exists a simple solution. To state our result, we first define a simple class of stochastic processes. We say $\langle \mu_t \rangle$ is a *piecewise-deterministic jump process* if and only if its paths are almost surely right-continuous and piecewise deterministic. Let $\Delta\mu_t \equiv \mu_t - \mu_{t^-}$ denote the jump in beliefs at time $t \in (0, T]$. We show that there exists an optimal contract with a simple structure.

Theorem 3.1. There exists an optimal contract for problem (FP) such that

(i) exploring rule $\langle \mu_t \rangle$ is a piecewise-deterministic jump process,

(ii) transfer rule $\langle \psi_t \rangle$ is deterministic and weakly increasing almost surely, and

(iii) stopping rule τ is such that, almost surely, for all $t \in (0, T]$, $\|\Delta\mu_t\| > 0$ implies $\tau = t$.

We defer to Section 5 for an outline of the proof.

Let us distill what Theorem 3.1 is saying. For the exploring rule, the principal acquires information via breakthroughs. If the players' belief evolves according to a piecewise-deterministic jump process, then it changes either deterministically between jumps or discontinuously at jump times. In terms of dynamic experiments, not observing a signal corresponds to a deterministic drift in beliefs. On the other hand, observing a signal means a discrete shift. Although the players could wait to observe multiple signals, after observing one breakthrough, the principal recommends that the agent stop; this is the role of the stopping rule. Lastly, for the transfer rule, the principal provides a deterministic flow of transfers and backloads it. Under this optimal contract, if the players arrive at time t, then the principal deterministically supplies a flow of transfer to the agent. The fact that the transfer rule is increasing over time means that the principal promises larger amounts of transfers towards the end.

With the interpretations at hand, we explain the economic forces. First, the exploring and stopping rules are driven by the players' risk-loving time preferences. The principal

and the agent discount time in a convex way. Hence, both love lotteries over time. The jump process captures this notion; the principal acquires information by waiting for a breakthrough, and yet, provides the agent with enough promise value to continue experimenting. At the same time, the prediction of the transfer rule is driven by the fact that the principal is risk-neutral while the agent is risk-averse over money. The principal only cares about the discounted expected transfers. However, the agent is strictly risk-averse; if principal were to use a transfer lottery, she would need to pay its certainty equivalent and risk premium. Replacing the lottery by its certainty equivalent preserves the agent's dynamic participation constraint, and the principal weakly improves. Therefore, the optimal transfer rule must be deterministic (almost surely). However, the fact that transfer rule is increasing is how the principal deals with the trade-off between risk-sharing and incentives. To see this, we state a useful corollary.

Corollary 3.2. Suppose the principal only faces an ex-ante participation constraint. Then there exists an optimal contract such that the exploring rule is a piecewise-deterministic jump process, the stopping rule is such that, almost surely, $\|\Delta\mu_t\| > 0$ implies $\tau = t$, and the transfer rule $\psi(\cdot)$ is constant almost surely.

Proof. See subsection B.1. □

If the principal only faced an ex-ante participation constraint, she would optimally give a constant flow of money. Set this as our "efficient" benchmark. When she faces a dynamic participation constraint, she cannot control the agent when to stop. Thus, a strictly increasing transfer rule means that the agent is incentivized to continue experimenting, but the trade-off is less efficient risk sharing (relative to the ex-ante participation case).

Theorem 3.1 provides a novel and simple solution within a very large contract space. We emphasize that the principal is allowed to choose history-dependent contracts and has a great degree of flexibility to acquire information. We also argue that our main result provides a key benchmark for the strategic experimentation literature. The majority of papers take the signal structure of experiments as exogenous. This paper provides a microfoundation for why experiments draw signals in a breakthrough fashion.

We remark that the prediction of the jump process is found in many dynamic information acquisition papers.[10] Although Theorem 3.1 arrives to a similar result, we argue that our setting is distinct for two reasons. First, to write the obvious, our setting involves a strategic relationship, whereas the dynamic information acquisition literature

[10]For example, see Zhong (2022), Hébert and Woodford (2023), Georgiadis-Harris (2026), and Sannikov and Zhong (2026).

only studies the decision-maker's problem. Moreover, we find it somewhat surprising that risk-loving time preferences still play an essential role in this strategic setting. Second, the principal jointly designs both the dynamic experiments and the transfer rule. This feature carries economic substance because transfers are feasible (and usually common) in such risk-sharing environments.

4 Discussion

In this section, we discuss properties of the optimal solution, and cover modifications to the original problem.

4.1 Binding Flow Constraint

Theorem 3.1 states that we can reduce our search of exploring rules to the simple class of piecewise-deterministic jump processes. However, it makes no mention of how the belief process drifts or jumps. To state our result, we introduce a key insight; instead of directly analyzing the exploring and stopping rule $(\langle \mu_t \rangle, \tau)$, we study the probability measure over the stopped belief-time pairs (μ_τ, τ).

Formally, given contract $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$, we let f denote the probability law over the belief-time stopping pairs (μ_τ, τ) (i.e., $(\mu_\tau, \tau) \sim f$). For a fixed f, let $f_\tau \equiv f_\tau(d\tau)$ denote its marginal measure over τ, $f_{\mu|\tau} \equiv f_{\mu|\tau}(d\mu \mid t)$ its conditional measure over beliefs μ given time $\tau = t$, and $F_\tau(t) = \int_{\tau \leq t} f(d\mu, d\tau)$ its cumulative function. Let $\widehat{\mu} : [0, T] \to \Delta^1$ be the continuation belief mapping; $\widehat{\mu}(t) = \mathbb{E}_f[\mu_\tau \mid \tau > t]$ if $\int_{\tau > t} f(d\mu, d\tau) > 0$ and $\widehat{\mu}(t) = 0$ if $\int_{\tau > t} f(d\mu, d\tau) = 0$.

Proposition 4.1. There exists an optimal contract $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$ such that its induced belief-time stopping law f satisfies the following: for every Borel measurable subset $B_\tau \subseteq \{t \in [0, T] : F_\tau(t) < 1\}$, we have

$$\chi \int_{\tau \in B_\tau} (1 - F_\tau(\tau)) d\tau = \int_{\tau \in B_\tau} \int \big(H(\mu) - H(\widehat{\mu}(\tau)) - \nabla H(\widehat{\mu}(\tau)) \cdot (\mu - \widehat{\mu}(\tau)) \big) f_{\mu|\tau}(d\mu \mid \tau) f_\tau(d\tau) \quad (1)$$

Proof. See subsection B.2. □

To further understand Proposition 4.1, let $B_\tau \subseteq \{t \in [0, T] : F_\tau(t) < 1\}$ and assume f_τ is absolutely continuous with respect to the Lebesgue measure on B_τ. Then for every

$t \in B_\tau$, the equality now becomes:

$$\chi = \frac{F'_\tau(t)}{1 - F_\tau(t)} \int_{\Delta^1} \big(H(\mu) - H(\widehat{\mu}(\tau)) - \nabla H(\widehat{\mu}(\tau)) \cdot (\mu - \widehat{\mu}(\tau))\big) f_{\mu|\tau}(d\mu \mid t).$$

This equality highlights that the principal faces a precision-frequency trade-off. Conditional measure $f_{\mu|\tau}$ corresponds to the "precision" because this is the jump distribution which determines the stopping utility for both players. At the same time, $\frac{F'_\tau(t)}{1-F_\tau(t)}$ corresponds to the frequency because this is the intensity rate of observing a breakthrough. If (Flow) constraint is binding, then capacity χ limits both quantities and forces the principal to trade-off between them. However, the precision-frequency trade-off can be further distorted if the agent's dynamic participation constraint is binding. Proposition 4.1 makes this intuition precise by considering measure-theoretic necessities.

4.2 Perfectly Patient Players

In the main model, we assumed both players were strictly risk-loving over time; that is, $r > 0$. As shown in the analysis, this played a key role in predicting the jump process. In this subsection, we study the problem where both players are perfectly patient. For tractability, we also assume an infinite time horizon. We explicitly record our assumptions below.

Assumption 6. Both players are perfectly patient, and the time horizon is infinite. That is, $r = 0$, $T = +\infty$, and the index set is $[0, +\infty)$.

Given prior belief $\mu_0 \in \operatorname{ri} \Delta^1$, define subset $\Pi_{\mu_0} = \{\pi \in \Delta(\Delta^1) : \mathbb{E}_\pi[\mu] = \mu_0\}$. We say $\pi \in \Pi_{\mu_0}$ is nondegenerate if $\pi \neq \delta_{\mu_0}$ where δ_{μ_0} is the Dirac measure at prior belief μ_0. Define the following static problem:

$$\begin{aligned} V_{SP}(\mu_0) = \max \quad & \mathbb{E}_\pi[U_P(\mu)] - \frac{\rho}{\chi} \cdot \mathbb{E}_\pi[H(\mu) - H(\mu_0)] \\ \text{over} \quad & \pi \in \Pi_{\mu_0} \text{ and } \rho \in [0, \overline{\psi}] \\ \text{s.t.} \quad & \mathbb{E}_\pi[U_A(\mu) - U_A(\mu_0)] + \frac{u(\rho) - c}{\chi} \cdot \mathbb{E}_\pi[H(\mu) - H(\mu_0)] \geq 0 \end{aligned} \tag{SP}$$

Given $\alpha \geq 0$ and $\pi \in \Pi_{\mu_0}$, we say stochastic process $\langle \mu_t \rangle$ is an α-dilution of π if

$$\mu_t = \mu_0 + \mathbf{1}\{N_\alpha(t) \geq 1\} \cdot (M - \mu_0), \quad \text{almost surely},$$

where M is a random variable distributed according to π and $N_\alpha(t)$ is an independent Poisson counting process with rate α. We now define an induced contract from the static

problem. Let $(\pi, \rho) \in \Pi_{\mu_0} \times [0, \overline{\psi}]$ be a feasible (but not necessarily optimal) static contract for problem (SP) such that π is nondegenerate. Define intensity rate $\alpha = \frac{\chi}{\mathbb{E}_\pi[H(\mu)] - H(\mu_0)}$. Let $M \sim \pi$ and define stopping rule $\tau \sim \text{Exp}(\alpha)$ such that it is independent of M. Let exploring rule $\langle \mu_t \rangle$ be the α-dilution of π. Choose $\rho_0 \in [0, \overline{\psi}] \setminus \{\rho\}$, and define transfer rule $\langle \psi_t \rangle$ by

$$\psi_t = \begin{cases} \rho & \text{if } t < \tau \\ \rho_0 & \text{if } t \geq \tau \end{cases}$$

We consider the natural filtration generated by $(\langle \mu_t \rangle, \langle \psi_t \rangle)$. We say contract $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$ (along with its underlying filtered probability space) is the induced contract from the static contract (π, ρ).

Taking stock, for every feasible static contract (π, ρ), we can construct a candidate induced contract $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$. With this construction at hand, we are ready to state this subsection's main result.

Proposition 4.2. Let static contract (π, ρ) solve problem (SP) such that π is nondegenerate. Let $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$ be the induced contract from (π, ρ). Under Assumption 6, contract $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$ solves problem (FP).

Proof. See subsection B.3. □

When the players are perfectly patient, Proposition 4.2 says that it suffices to study a simple static problem where we ignore the time dimension. It also illustrates that the principal does not need to consider the whole dynamic participation constraint; she only needs to ensure that the agent's ex-ante participation constraint is satisfied. Hence, a constant flow of transfers is optimal in this setting.

To illustrate the tractability and usefulness of studying the static problem (SP), we solve a toy model with explicit parameters. The purpose of this example is to delineate the effects of the agent's flow cost on the optimal solution.

Example 1. We lay out the parameters of the environment: $\chi = 0.1$, $H(\mu) = (\mu - 0.5)^2$, and $\mu_0 = 0.5$. For a fixed $c > 0$, we assume transfers are bounded above by $\overline{\psi} \equiv \overline{\psi}(c) = c + 1$; we will show that the optimal transfer rule never binds at the upper bound. For the preferences, we assume that the principal and agent share the same stopping utility. In particular, the players have conflicting preferences only in the cost of experimentation and risk-attitudes over money. Formally, $U_P(\mu) = U_A(\mu) = \max\{\mu - 0.5, 0\}$ and $u(\psi) = \sqrt{\psi}$. The agent's flow cost $c \geq 0$ is a free constant. It can be shown that the optimal

posterior belief has binary support and can be written so that $\pi_z = 0.5 \cdot \delta_{\mu_l(z)} + 0.5 \cdot \delta_{\mu_h(z)}$ where $\mu_l(z) = 0.5 - z$ and $\mu_h(z) = 0.5 + z$ for some $z \in [0, 0.5]$. For every $c \geq 0$, the optimal transfer rule and beliefs are given by:

$$(\rho, \mu_h, \mu_l)(c) = \begin{cases} (0, 1, 0), & 0 \leq c < c_0, \\ \left(\left(c - \frac{1}{5}\right)^2, 1, 0\right), & c_0 \leq c < c_1, \\ \left(\frac{c^2}{(2c+1)^2}, \frac{1}{2} + \frac{2c+1}{20c^2}, \frac{1}{2} - \frac{2c+1}{20c^2}\right), & c \geq c_1 \end{cases}$$

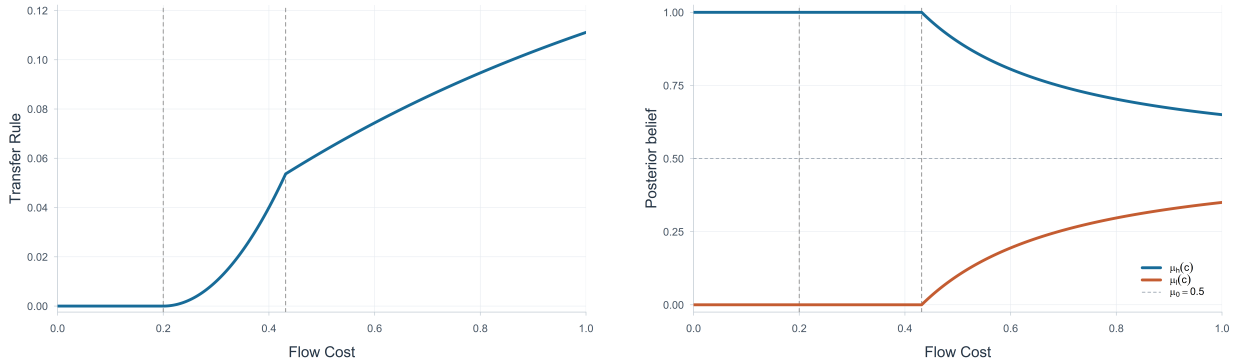
where $c_0 = 0.2$ and $c_1 = \frac{1+\sqrt{11}}{10} \approx 0.432$. Figure 1 graphically illustrates how the optimal contract changes as the agent's flow cost varies. When the agent's flow cost is sufficiently low (i.e., $0 \leq c < c_0$), the principal designs the experiment so that the agent acquires full information. Moreover, she does not use transfers to incentivize the agent. Intuitively, since the players' preferences are sufficiently aligned, the principal can obtain her first-best value (that is, the value of solving problem (SP) but ignoring the constraint). There is a middle region where the flow cost has bite (i.e., $c_0 \leq c_1 < c_1$). The principal uses the transfer rule to incentivize the agent, but she can still use the first-best experiment. Lastly, when the flow cost is sufficiently large (i.e., $c > c_1$), the principal adjusts both the experiment and transfers.

Although this is a toy model, we gain valuable insight in understanding the principal's problem. First, we have an unambiguous monotone comparative statics result on the both the optimal experiments and transfer rule. The transfer rule is weakly increasing in c while the posterior beliefs both tend to the prior μ_0. Second, notice that the transfer rule is never binding at the flow cost. In fact, for every optimal solution (π, ρ), we have that $\rho < c$. Hence, fully compensating the agent's flow cost need not be optimal. Lastly, given that the principal can fully design and commit to the exploring rule, one may conjecture that she could use transfers to obtain her first-best experiment. This example shows that this conjecture is false. While there are regions of c where she can feasibly design the fully revealing experiment, for a sufficiently large flow cost, she cannot achieve the first-best outcome.

5 Analysis

This section is dedicated to outlining the proof of theorem 3.1. Detailed arguments are found in section A.

Figure 1: Optimal transfers and posterior beliefs in Example 1.



(a) Optimal transfer rule

(b) Optimal posterior beliefs

Note: Figure (a) plots the optimal transfer rule as a function of the agent's flow cost. Figure (b) plots the corresponding optimal posterior beliefs. The upper graph in (b) corresponds to the high posterior belief $\mu_h(c) = 0.5 + \frac{2c+1}{20c^2}$ while the lower graph corresponds to the low posterior belief $\mu_l(c) = 0.5 - \frac{2c+1}{20c^2}$. The dashed vertical lines represents the thresholds $c_0 = 0.2$ and $c_1 = \frac{1+\sqrt{11}}{10} \approx 0.432$.

5.1 Relaxed Problem

Our goal in this subsection is to relax the principal's problem into more tractable choice variables and constraints. Let $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$ be a feasible contract. Instead of directly dealing with the transfer rule, we consider the indirect compensation. Formally, define stochastic process $\langle v_t \rangle$ such that it is indistinguishable from $\langle u(\psi_t) \rangle$.[11] Given Assumptions 2 and 3, $\langle v_t \rangle$ is bounded below by $u(0) = 0$ and bounded above by $\overline{v} := u(\overline{\psi})$. The principal's cost of providing v_t to the agent is $\varphi(v_t)$ where $\varphi := u^{-1}$ denotes the inverse utility. Under this transformation, the principal and agent's expected utilities are respectively

$$\mathbb{E}_{\mathcal{P}}\left[e^{-r\tau}U_P(\mu_\tau) - \int_0^\tau e^{-rs}\varphi(v_s)ds\right] \quad \text{and} \quad \mathbb{E}_{\mathcal{P}}\left[e^{-r\tau}U_A(\mu_\tau) + \int_0^\tau e^{-rs}(v_s - c)ds\right].$$

We continue to manipulate the contract. Rather than considering a general stochastic process, we focus on a deterministic one. Abusing and reusing notation, define $v : [0,T] \to [0,+\infty)$ by $v(t) = \mathbb{E}_{\mathcal{P}}[v_t \mid t < \tau]$ if $\mathcal{P}(t < \tau) > 0$ and $v(t) = 0$ if $\mathcal{P}(t < \tau) = 0$. For the exploring and stopping rules, rather than directly tackling the probability law over paths, the principal will choose the law over stopping realizations. Formally, let $f \equiv f(d\mu, d\tau) \in \Delta(\Delta^1 \times [0,T])$ denote the joint probability measure over stopping realizations.[12] An induced contract is a tuple (f, v). Whenever the context is clear, we will

[11] Given filtered probability space $(\Omega, \mathcal{F}, \langle \mathcal{F}_t \rangle, \mathcal{P})$, stochastic processes $\langle v_t \rangle$ and $\langle u(\psi_t) \rangle$ are indistinguishable if and only if the set $\{(\omega, t) \in \Omega \times [0,T] : v_t(\omega) \neq u(\psi_t(\omega))\}$ is $\mathcal{P}$-null.

[12] Given subset $Z \subseteq \mathbb{R}^n$, we let $\Delta(Z)$ denote the space of Borel probability measures on Z, with the weak topology.

also call (f, v) is a contract.

We now lay out necessary conditions for the constraints. Both players are Bayesian, and hence, must update their beliefs according to Bayes' rule from time zero. Since f must also respect this, we set a Bayes plausibility constraint; that is,

$$\int \mu f(d\mu, d\tau) = \mu_0. \tag{BP}$$

Next, we relax the variation constraint. Note that (Var) is imposed only on the exploring rule $\langle \mu_t \rangle$. That is, for all time $t, t' \in [0, T]$ with $t < t'$, constraint (Var) bounds how much $\langle \mu_t \rangle$ can "vary". Instead, we will fix only one time t and jointly consider the stopping rule τ. The thought exercise is the following: starting from the initial time, how much information can be acquired up to time $\min\{\tau, t\}$? Formally, given exploring rule $\langle \mu_t \rangle$ and stopping rule τ, we only study the stopped process $\mu_{\min\{\tau,t\}}$ for every time $t \in [0, T]$. In terms of the joint probability law f, the (relaxed) flow constraint is

$$\int_{\tau \le t} H(\mu) f(d\mu, d\tau) + H(\widetilde{\mu}_>(t)) \le H(\mu_0) + \chi \int \min\{\tau, t\} f(d\mu, d\tau), \quad \forall t \in [0, T] \tag{Flow}$$

where $\widetilde{\mu}_>(t)$ is the unweighted continuation belief at time t, excluding paths that stop at t. Formally, mapping $\widetilde{\mu}_> : [0, T] \to \mathbb{R}^2_+$ is defined by $\widetilde{\mu}_>(t) := \int_{\tau > t} \mu f(d\mu, d\tau)$ if $\int_{\tau > t} f(d\mu, d\tau) > 0$ and $\widetilde{\mu}_>(t) = 0$ if $\int_{\tau > t} f(d\mu, d\tau) = 0$. In the same way, we define $\widetilde{\mu}_\ge(t)$ as the unweighted continuation belief at time t including stopping at t. That is, mapping $\widetilde{\mu}_\ge : [0, T] \to \mathbb{R}^2_+$ is defined by $\widetilde{\mu}_\ge(t) := \int_{\tau \ge t} \mu f(d\mu, d\tau)$ if $\int_{\tau \ge t} f(d\mu, d\tau) > 0$ and $\widetilde{\mu}_\ge(t) = 0$ if $\int_{\tau \ge t} f(d\mu, d\tau) = 0$.

Lastly, we tackle (DPC). Instead of satisfying the participation constraint for every history, we aggregate the agent's continuation value over the event $\{\tau \ge t\}$ of not having stopped before t. The relaxed dynamic participation constraint is

$$\int_{\tau \ge t} \left(e^{-r(\tau - t)} U_A(\mu) + \int_t^\tau e^{-r(s-t)} (v(s) - c) ds \right) f(d\mu, d\tau) \ge U_A(\widetilde{\mu}_\ge(t)), \ \forall t \in [0, T] \tag{RDP}$$

Overall, the principal's relaxed problem is to choose induced contract (f, v) to solve

$$\begin{aligned} V_{RP}(\mu_0) = \sup \quad & \int \left(e^{-r\tau} U_P(\mu) - \int_0^\tau e^{-rs} \varphi(v(s)) ds \right) f(d\mu, d\tau) \\ \text{subject to} \quad & \text{(BP)}, \text{(Flow)}, \text{(RDP)} \end{aligned} \tag{RP}$$

The following lemma states that we indeed have a relaxed problem.

Lemma 5.1. $V_{FP}(\mu_0) \le V_{RP}(\mu_0)$ for all $\mu_0 \in \Delta^1$.

Proof. See subsection A.1. □

We relax the principal's problem one more time by using discretization. First, we discretize the participation constraint and reweight the agent's continuation and stopping values by the appropriate discount factor. Formally, let $(t_k)_{k=0}^K$ uniformly partition $[0,T]$ (where $0 = t_0 < t_1 < \cdots < t_K = T$). Given $(t_k)_{k=0}^K$, induced contract (f,v) satisfies the discretized dynamic participation constraint if and only if

$$\int_{\tau \geq t_k} \left(e^{-r\tau} U_A(\mu) + \int_{t_k}^{\tau} e^{-rs}(v(s) - c)ds \right) f(d\mu, d\tau) \geq e^{-rt_k} U_A(\widetilde{\mu}_{\geq}(t_k)) \tag{RDP$_K$}$$

for $k = 0, 1, \ldots, K-1$. Next, we relax the flow constraint by ignoring it at a relatively small neighborhood around the initial time. Formally, an induced contract (f,v) satisfies the truncated flow constraint if and only if

$$\int_{\tau \leq t} H(\mu) f(d\mu, d\tau) + H(\widetilde{\mu}_{>}(t)) \leq H(\mu_0) + \chi \int \min\{\tau, t\} f(d\mu, d\tau), \ \forall t \in [2^{-K}, T] \tag{Flow$_K$}$$

Later we will use limiting arguments as $K \to +\infty$. The idea is that in the limit, constraints (RDP$_K$) and (Flow$_K$) tend to (RDP) and (Flow), respectively. Hence, without further mention, we implicitly assume that K is large enough such that $2^{-K} < T$. The discretized relaxed problem is

$$\begin{aligned} V_{RP,K}(\mu_0) = \sup \quad & \int \left(e^{-r\tau} U_P(\mu) - \int_0^{\tau} e^{-rs} \varphi(v(s)) ds \right) f(d\mu, d\tau) \\ \text{subject to} \quad & \text{(BP)}, \text{(Flow}_K\text{)}, \text{(RDP}_K\text{)} \end{aligned} \tag{RP$_K$}$$

5.2 Strong Duality and Optimality Conditions

We study the relaxed problem with strong duality and optimality conditions. We set up notation in spirit of Luenberger (1969). Fix whole number p such that $1 < p < +\infty$. We define the sets of feasible choice variables.[13]

$$\Delta_{\mu_0} = \left\{ f \in \Delta(\Delta^1 \times [0,T]) : \int \mu f(d\mu, d\tau) = \mu_0 \right\}$$
$$\mathbb{V} = \{ v \in L^p([0,T]) : 0 \leq v(t) \leq \overline{v}, \ t \in [0,T] \text{ a.e.} \}$$

Henceforth, all induced contracts are implicitly elements of $\Delta_{\mu_0} \times \mathbb{V}$. Next, we define the constraint mappings. Recall that $\widetilde{\mu}_{>}(t)$ and $\widetilde{\mu}_{\geq}(t)$ denote the unweighted continuation

[13] Recall that $\overline{v} := \varphi(\overline{\psi})$.

beliefs at time t, excluding and including paths that stop at t, respectively. Given belief-time distribution $f \in \Delta_{\mu_0}$ and time $t \in [0,T]$, we define the flow constraint mapping $G_{Flow}(f)(t)$ by

$$G_{Flow}(f)(t) = \chi \int \min\{\tau, t\} f(d\mu, d\tau) + H(\mu_0) - \int_{\tau \leq t} H(\mu) f(d\mu, d\tau) - H(\widetilde{\mu}_{>}(t))$$

Given partition $(t_k)_{k=0}^K$ and contract $(f,v) \in \Delta_{\mu_0} \times \mathbb{V}$, we define the dynamic participation constraint mapping $G_{DP,K}(f,v) \in \mathbb{R}_+^K$ by

$$G_{DP,K}(f,v) = \left(\int_{\tau \geq t_k} \left(e^{-r\tau} U_A(\mu) + \int_{t_k}^{\tau} e^{-rs}(v(s) - c) ds \right) f(d\mu, d\tau) - e^{-rt_k} U_A(\widetilde{\mu}_{\geq}(t_k)) \right)_{k=0}^{K-1}$$

The objective function is defined by

$$\text{OBJ}(f,v) = \int \left(e^{-r\tau} U_P(\mu) - \int_0^{\tau} e^{-rs} \varphi(v(s)) ds \right) f(d\mu, d\tau)$$

Given our assumptions, our model has something to say.

Lemma 5.2. Problem (RP_K) has a solution.

Proof. See section A.2. □

We build up the tools of strong duality and optimality conditions to study relaxed problem (RP_K). Given subset $Z \subseteq \mathbb{R}^n$, let $\mathcal{B}(Z)$ denote the set of positive Borel measures over Z. Given $\lambda \in \mathcal{B}([2^{-K}, T])$, let $\Lambda(t) = \int_{s>t} d\lambda(s)$ denote its survival function. We let $\mathbb{L} = \{\lambda \in \mathcal{B}([2^{-K}, T]) : \Lambda \in L^1([2^{-K}, T])\}$ be the set of multipliers with finite measures. Let $\lambda \in \mathbb{L}$ and $\gamma \in \mathbb{R}_+^K$ be the multipliers for (Flow_K) and (RDP_K), respectively. The Lagrangian is given by

$$\mathcal{L}(f,v,\lambda,\gamma) = \text{OBJ}(f,v) + \int_{t \in [2^{-K}, T]} G_{Flow}(f)(t) d\lambda(t) + \gamma \cdot G_{DP,K}(f,v)$$

The primal and dual problems are

$$\sup_{(f,v) \in \Delta_{\mu_0} \times \mathbb{V}} \left\{ \inf_{(\lambda,\gamma) \in \mathbb{L} \times \mathbb{R}_+^K} \mathcal{L}(f,v,\lambda,\gamma) \right\} \tag{Primal_K}$$

$$\inf_{(\lambda,\gamma) \in \mathbb{L} \times \mathbb{R}_+^K} \left\{ \sup_{(f,v) \in \Delta_{\mu_0} \times \mathbb{V}} \mathcal{L}(f,v,\lambda,\gamma) \right\} \tag{Dual_K}$$

Lemma 5.3. Strong duality holds. Formally, (Primal$_K$) = (Dual$_K$) and there exist multipliers $(\lambda, \gamma) \in \mathbb{L} \times \mathbb{R}_+^K$ that solves (Dual$_K$).

Proof. See section A.3. □

With some rewriting, the Lagrangian can be given by

$$\mathcal{L}(f, v, \lambda, \gamma) = \begin{cases} \int e^{-r\tau} U_P(\mu) f(d\mu, d\tau) + \int_0^T e^{-rs}\big(\Gamma(s)v(s) - \varphi(v(s))\big)\big(1 - F(s)\big)ds \\ \qquad + \int \left(e^{-r\tau}\Gamma(\tau)U_A(\mu) - c\int_0^\tau e^{-rs}\Gamma(s)ds\right) f(d\mu, d\tau) - \sum_{k=0}^{K-1} e^{-rt_k} U_A(\widetilde{\mu}_{\geq}(t_k))\gamma_k \\ \qquad + \int_{t\in[2^{-K},T]} G_{Flow}(f)(t) d\lambda(t) \end{cases}$$

where $\Gamma(\tau) = \sum_k \gamma_k \mathbf{1}\{t_k \leq \tau\}$ denotes the cumulative multiplier and $F(s) = \int_{\tau \leq s} f(d\mu, d\tau)$ denotes the survival function of the stopping time distribution.

We are ready to derive optimality conditions. Let $\partial U_A(\cdot)$ and $\partial H(\cdot)$ denote the subdifferentials of $U_A(\cdot)$ and $H(\cdot)$, respectively.[14] Informally, the Gâteaux derivative of the Lagrangian $\mathcal{L}(f, v, \lambda, \gamma)$ with respect to mass f at point $(\mu, \tau) \in \Delta^1 \times [0, T]$ is given by

$$\ell(\mu, \tau \mid f, v, \lambda, \gamma) = \begin{cases} e^{-r\tau} U_P(\mu) + \int_0^\tau e^{-rs}\big(\Gamma(s)(v(s) - c) - \varphi(v(s))\big)ds \\ \qquad + e^{-r\tau}\Gamma(\tau)U_A(\mu) - \sum_{k=0}^{K-1} \mathbf{1}\{\tau \geq t_k\} e^{-rt_k} \cdot \big(y(t_k) \cdot \mu\big) \cdot \gamma_k \\ \qquad + \chi \int_{2^{-K}}^\tau \Lambda(t)dt - \int_{t\in[0,\tau)} \big(z(t) \cdot \mu\big) d\lambda(t) - \Lambda(\tau)H(\mu) \end{cases}$$

where $y(t_k)$ and $z(t)$ are some measurable selections from $\partial U_A(\widetilde{\mu}_{\geq}(t_k))$ and $\partial H(\widetilde{\mu}_{>}(t))$, respectively.[15] In the same manner, the Gâteaux derivative of the Lagrangian with respect to indirect compensation $v \in \mathbb{V}$ in direction $w - v$ is given by

$$\int e^{-rt}\big(1 - F(t)\big)\big(\varphi'(v(t)) - \Gamma(t)\big)(w(t) - v(t))dt.$$

Both of these derivations suggest necessary first-order conditions for an optimal contract. The next lemma makes this precise.

[14] Given convex function $f : \mathbb{R} \to \mathbb{R}$ and point $x \in \mathbb{R}$, the subdifferential of f at x, denoted by $\partial f(x)$, is the set of points $x' \in \mathbb{R}$ which satisfies the subgradient inequality: $f(y) \geq f(x) + x' \cdot (y - x)$ for all $y \in \mathbb{R}$.

[15] To formally define $\ell \equiv \ell(\mu, \tau \mid f, v, \lambda, \gamma)$, we must specify a selection from the subdifferential $\partial H(\widetilde{\mu}_{>}(t))$ for every time t. Let $\bar{t} = \sup_t \operatorname{supp} f$. If $t < \bar{t}$, choose the gradient $\nabla H(\widetilde{\mu}_{>}(t))$. If $t \geq \bar{t}$, choose a selection from $\partial H(0)$. This set is nonempty by assumption 4(i). Then, ℓ is well-defined on $\Delta^1 \times [2^{-K}, T]$.

Lemma 5.4. If multipliers (λ, γ) solve (Dual$_K$) and for all contracts (f, v) solving (RP$_K$), there exists belief multipliers $a \in \mathbb{R}^2$ and measurable selections from $\partial U_A(0)$ and $\partial H(0)$ such that the following optimality conditions hold:

(i) *First-Order Conditions.*

$$\ell(\mu, \tau \mid f, v, \lambda, \gamma) \leq a \cdot \mu, \quad \forall (\mu, \tau) \in \Delta^1 \times [0, T] \tag{FOC$_f$}$$

$$\begin{cases} \Gamma(t) \leq \varphi'(v(t)) & \text{if } v(t) = 0 \\ \Gamma(t) = \varphi'(v(t)) & \text{if } v(t) \in (0, \overline{v}) \ , \quad \text{a.e. } t \text{ with } F_\tau(t) < 1 \\ \Gamma(t) \geq \varphi'(v(t)) & \text{if } v(t) = \overline{v} \end{cases} \tag{FOC$_v$}$$

(ii) *Concave Closure Condition.*

$$\int \big(\ell(\mu, \tau \mid f, v, \lambda, \gamma) - a \cdot \mu\big) f(d\mu, d\tau) = 0 \tag{CC}$$

(iii) *Complementary Slackness Conditions.*

$$\int_{t \in [2^{-K}, T]} G_{Flow}(f)(t) d\lambda(t) = 0 \tag{CS$_{Flow}$}$$

$$\gamma \cdot G_{DP,K}(f, v) = 0 \tag{CS$_{DP}$}$$

Proof. See section A.4. □

5.3 Constructing the Optimal Solution

With strong duality and the optimality conditions, we are ready to analyze the principal's discretized relaxed problem. Our first order of business is to show the relaxed flow constraint must always be binding at an optimal solution. This is key because it predicts that the belief-time distribution must induce a piecewise-deterministic jump process.

Lemma 5.5. Let multipliers (λ, γ) solve (Dual$_K$), contract (f, v) solve (RP$_K$), and define $\overline{t} = \sup_t \text{supp } f$. Then constraint (Flow) binds for every time $t \in [2^{-K}, \overline{t}]$.

Proof. See section A.5. □

We pass on this key property to the relaxed problem (RP). To do so, we use a limiting argument to construct an optimal solution.

Lemma 5.6. There exists a contract (f, v) which solves problem (RP) and such that constraint (Flow) binds for every time $t \in [0, \overline{t}]$ (where $\overline{t} = \sup_t \text{supp } f$).

Proof. See section A.6. □

Our last step is to construct a feasible probability law over belief-time paths which induces this optimal distribution. Fix belief-time distribution $f \in \Delta_{\mu_0}$ and define $\widehat{\mu}(s) = \mathbb{E}_{f(\mu,\tau)}[\mu \,|\, \tau > s]$ whenever $\int_{\tau>s} f(d\mu, d\tau) > 0$ and $\widehat{\mu}(s) = 0$ if $\int_{\tau>s} f(d\mu, d\tau) = 0$. Define the "belief-path" mapping ϕ jumping to belief μ at time t by

$$\phi^{(\mu,t)}(s) = \begin{cases} \widehat{\mu}(s) & \text{if } s < t \\ \mu & \text{if } s \geq t \end{cases}$$

Define $\Phi : (\mu, t) \to (\phi^{(\mu,t)}, t)$ as the mapping from belief-time pair to the belief path and jump.

We are now ready to define the law over these paths. Let Ω_μ denote the set of càdlàg Δ^1-valued paths starting from μ_0. For $f \in \Delta_{\mu_0}$, define the law $\mathcal{P}$ of the corresponding belief process $\langle \mu_t^f \rangle$ and stopping rule τ^f as the pushforward of f by Φ, i.e., $\mathcal{P} = \Phi \# f$. Formally, for all Borel sets $B_\mu \subseteq \Omega_\mu$ and $B_\tau \subseteq [0, T]$:

$$\mathcal{P}(\mu_\tau^f \in B_\mu \text{ and } \tau^f \in B_\tau) = \int_{\Phi^{-1}(B_\mu \times B_\tau)} f(d\mu, d\tau)$$

Lemma 5.7. Let induced contract (f, v) solve (RP) such that (Flow) is always binding. Let $F(t) = \int_{\tau \leq t} f(d\mu, d\tau)$, and define $\psi(t) = \varphi(v(t))$ if $F(t) < 1$ and $\psi(t) = 0$ if $F(t) = 1$. Then contract $(\langle \mu_t^f \rangle, \psi, \tau^f)$ solves (FP).

Proof. See section A.7. □

6 Conclusion

In this paper, we studied a dynamic principal-agent problem where the principal must incentivize the agent to acquire information over a risky project. While the principal can flexibly design the experiments and monetary transfers, the agent lacks commitment and is constrained by how much information they can generate. Despite the rich contract space, our main result states that there exists a simple optimal contract. The principal designs the experiments by acquiring information via breakthroughs, offers deterministic and nondecreasing transfers, and recommends halting after observing a breakthrough. This prediction has broad implications in dynamic contracting and the design of information acquisition. Future work should explore extensions of this model. For instance, the

agent may hold private information about the state of the world or may covertly acquire information.

References

Arrow, Kenneth J, David Blackwell, and Meyer A Girshick (1949). "Bayes and minimax solutions of sequential decision problems." In: *Econometrica, Journal of the Econometric Society*, pp. 213–244.

Azoulay, Pierre and Danielle Li (2020). *Scientific grant funding*. Vol. 26889. National Bureau of Economic Research Cambridge, MA.

Ball, Ian (2023). "Dynamic information provision: Rewarding the past and guiding the future." In: *Econometrica* 91.4, pp. 1363–1391.

Billingsley, Patrick (2013). *Convergence of probability measures*. John Wiley & Sons.

Bryan, Kevin A and Heidi L Williams (2021). "Innovation: market failures and public policies." In: *Handbook of industrial organization*. Vol. 5. 1. Elsevier, pp. 281–388.

Bühler, Theo and Dietmar A Salamon (2018). *Functional analysis*. Vol. 191. American Mathematical Soc.

Caplin, Andrew, Mark Dean, and John Leahy (2022). "Rationally inattentive behavior: Characterizing and generalizing Shannon entropy." In: *Journal of Political Economy* 130.6, pp. 1676–1715.

Chade, Hector and Natalia Kovrijnykh (2016). "Delegated information acquisition with moral hazard." In: *Journal of Economic Theory* 162, pp. 55–92.

Che, Yeon-Koo and Konrad Mierendorff (2019). "Optimal dynamic allocation of attention." In: *American Economic Review* 109.8, pp. 2993–3029.

Georgiadis-Harris, Alkis (2026). *Preparing to act*. Tech. rep. Tech. rep., Working Paper.

Grenadier, Steven R, Andrey Malenko, and Nadya Malenko (2016). "Timing decisions in organizations: Communication and authority in a dynamic environment." In: *American Economic Review* 106.9, pp. 2552–2581.

Guo, Yingni (2016). "Dynamic delegation of experimentation." In: *American Economic Review* 106.8, pp. 1969–2008.

Halac, Marina, Navin Kartik, and Qingmin Liu (2016). "Optimal contracts for experimentation." In: *The Review of Economic Studies* 83.3, pp. 1040–1091.

Hébert, Benjamin and Michael Woodford (2023). "Rational inattention when decisions take time." In: *Journal of Economic Theory* 208, p. 105612.

Hébert, Benjamin and Weijie Zhong (2026). "Engagement maximization." In: *Theoretical Economics* 21.3, pp. 849–887.

Henry, Emeric and Marco Ottaviani (2019). "Research and the approval process: The organization of persuasion." In: *American Economic Review* 109.3, pp. 911–955.

Herresthal, Claudia (2022). "Hidden testing and selective disclosure of evidence." In: *Journal of Economic Theory* 200, p. 105402.

Iossa, Elisabetta and Francesca Stroffolini (2005). "Price cap regulation, revenue sharing and information acquisition." In: *Information Economics and Policy* 17.2, pp. 217–230.

Keller, Godfrey and Sven Rady (2010). "Strategic experimentation with Poisson bandits." In: *Theoretical Economics* 5.2, pp. 275–311.

— (2015). "Breakdowns." In: *Theoretical Economics* 10.1, pp. 175–202.

Keller, Godfrey, Sven Rady, and Martin Cripps (2005). "Strategic experimentation with exponential bandits." In: *Econometrica* 73.1, pp. 39–68.

Koh, Andrew, Sivakorn Sanguanmoo, and Weijie Zhong (2026). "Persuasion and optimal stopping." In: *arXiv preprint arXiv:2406.12278*.

Lambert, Richard A (1986). "Executive effort and selection of risky projects." In: *The RAND Journal of Economics*, pp. 77–88.

Luenberger, David G (1969). *Optimization by vector space methods*. John Wiley & Sons.

Manso, Gustavo (2011). "Motivating innovation." In: *The journal of finance* 66.5, pp. 1823–1860.

McClellan, Andrew (2022). "Experimentation and approval mechanisms." In: *Econometrica* 90.5, pp. 2215–2247.

Royden, Halsey Lawrence and Patrick Fitzpatrick (1988). *Real analysis*. Vol. 32. Macmillan New York.

Rudin, Walter (1976). *Principles of Mathematical Analysis*. 3rd. International Series in Pure and Applied Mathematics. New York: McGraw-Hill. ISBN: 007054235X.

Sannikov, Yuliy and Weijie Zhong (2026). "Exploration and Stopping." In.

Wald, A. (June 1945). "Sequential Tests of Statistical Hypotheses." In: *The Annals of Mathematical Statistics* 16.2, pp. 117–186. DOI: 10.1214/aoms/1177731118.

Yoder, Nathan (2022). "Designing incentives for heterogeneous researchers." In: *Journal of Political Economy* 130.8, pp. 2018–2054.

Zhao, Wei (2025). "Dynamic Evaluation in Learning." In: *Available at SSRN 5388111*.

Zhao, Wei et al. (2024). "Contracting over persistent information." In: *Theoretical Economics* 19.2, pp. 917–974.

Zhong, Weijie (2022). "Optimal dynamic information acquisition." In: *Econometrica* 90.4, pp. 1537–1582.

A Proofs for Theorem 3.1

A.1 Proof for Lemma 5.1

Let $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$ be an arbitrary feasible contract and $(f, \widehat{v})$ its respective induced contract. We will argue that $(f, \widehat{v})$ weakly improves the principal's expected utility (relative to $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$) and satisfies constraints (BP), (Flow), and (RDP). We start by showing all the constraints are necessary. Since $\langle \mu_t \rangle$ is a Δ^1-valued martingale and τ is bounded, by the optional sampling theorem, (BP) must hold. For the variation constraint, note that (Var) implies $H(\mu_t) - \chi \cdot t$ is a supermartingale. Since $T < +\infty$, stopping time τ is bounded. Fix time $t \in [0, T]$. By the optional sampling theorem, we get

$$\mathbb{E}_{\mathcal{P}}\left[H(\mu_{\tau \wedge t}) - H(\mu_0)\right] \leq \chi\, \mathbb{E}_{\mathcal{P}}\left[\min\{\tau, t\}\right]$$

Note that we also have

$$\mathbb{E}[\mathbf{1}\{\tau > t\} H(\mu_t)] = \mathcal{P}(\tau > t)\mathbb{E}_{\mathcal{P}}[H(\mu_t) \mid \tau > t] \geq \mathcal{P}(\tau > t) H\big(\mathbb{E}_{\mathcal{P}}[\mu_t \mid \tau > t]\big) = H(\widetilde{\mu}_{>}(t))$$

where the first inequality holds by Jensen's inequality and the last equality holds since $H(\cdot)$ is homogeneous of degree one and by definition of $\widetilde{\mu}_{>}(\cdot)$. This holds with equality if $\mathcal{P}(\tau > t) = 0$ since $H(0) = 0$. Thus,

$$\begin{aligned}
\chi \int \min\{\tau, t\} f(d\mu, d\tau) + H(\mu_0) &= \chi\, \mathbb{E}_{\mathcal{P}}\left[\min\{\tau, t\}\right] + H(\mu_0) \\
&\geq \mathbb{E}_{\mathcal{P}}\left[H(\mu_{\min\{\tau, t\}})\right] \\
&= \mathbb{E}\big[\mathbf{1}\{\tau \leq t\} H(\mu_\tau)\big] + \mathbb{E}\big[\mathbf{1}\{\tau > t\} H(\mu_t)\big] \\
&\geq \int_{\tau \leq t} H(\mu)\, f(d\mu, d\tau) + H(\widetilde{\mu}_{>}(t))
\end{aligned}$$

Therefore, f satisfies the flow constraint (Flow).

For the dynamic participation constraint, fix time $t \in [0, T]$. Let $V_t = u(\psi_t)$ and define

$$\widehat{v}(s) := \begin{cases} \mathbb{E}[V_s \mid \tau > s], & \mathcal{P}(\tau > s) > 0, \\ 0, & \mathcal{P}(\tau > s) = 0. \end{cases}$$

Since τ is a stopping time, we have $\{\tau \geq t\} \in \mathcal{F}_t$. Define the agent's continuation value at time t by

$$Y_t := e^{-r(\tau - t)} U_A(\mu_\tau) + \int_t^\tau e^{-r(s-t)} (V_s - c)\, ds,$$

and extend Y_t to be zero outside $\{\tau \geq t\}$. The full dynamic participation constraint gives $\mathbb{E}[Y_t \mid \mathcal{F}_t] \geq U_A\left(\mathbb{E}[\mu_\tau \mid \mathcal{F}_t]\right)$ on $\{\tau > t\}$. On $\{\tau = t\}$, the flow integral is zero, $\mu_\tau = \mu_t$, and

$$\mathbf{1}\{\tau = t\}\mathbb{E}[Y_t \mid \mathcal{F}_t] = \mathbf{1}\{\tau = t\}U_A(\mu_t) = \mathbf{1}\{\tau = t\}U_A\left(\mathbb{E}[\mu_\tau \mid \mathcal{F}_t]\right).$$

Thus, including paths that stop exactly at t adds an equality at the history level; it does not require a new restriction on the full contract. Multiplying by $\mathbf{1}\{\tau \geq t\}$, taking expectations, and using Jensen's inequality and homogeneity yields

$$\begin{aligned}\mathbb{E}[\mathbf{1}\{\tau \geq t\}Y_t] &\geq \mathbb{E}\left[\mathbf{1}\{\tau \geq t\}U_A\left(\mathbb{E}[\mu_\tau \mid \mathcal{F}_t]\right)\right] \\ &\geq U_A\left(\mathbb{E}\left[\mathbf{1}\{\tau \geq t\}\mathbb{E}[\mu_\tau \mid \mathcal{F}_t]\right]\right) \\ &= U_A\big(\mathbb{E}[\mathbf{1}\{\tau \geq t\}\mu_\tau]\big) = U_A(\widetilde{\mu}_{\geq}(t)).\end{aligned}$$

The Jensen step applies to the cone-valued random vector $\mathbf{1}\{\tau \geq t\}\mathbb{E}[\mu_\tau \mid \mathcal{F}_t]$, since $U_A(0) = 0$. This formulation also covers $\mathcal{P}(\tau \geq t) = 0$ without dividing by a tail probability. To express the left-hand side using $\widehat{v}(\cdot)$, Fubini gives

$$\begin{aligned}\mathbb{E}\left[\mathbf{1}\{\tau \geq t\}\int_t^\tau e^{-r(s-t)}V_s\,ds\right] &= \int_t^T e^{-r(s-t)}\mathbb{E}[\mathbf{1}\{s < \tau\}V_s]\,ds \\ &= \int_t^T e^{-r(s-t)}\mathcal{P}(\tau > s)\widehat{v}(s)\,ds \\ &= \int_{\tau \geq t}\int_t^\tau e^{-r(s-t)}\widehat{v}(s)\,ds\,df.\end{aligned}$$

For $s \geq t$, the event $\{s < \tau\}$ implies $\{\tau \geq t\}$, which explains why its indicator alone appears in the first integral. The singleton integration endpoint has zero Lebesgue duration. Substitution into the preceding inequality now gives exactly

$$\int_{\tau \geq t}\left[e^{-r(\tau-t)}U_A(\mu) + \int_t^\tau e^{-r(s-t)}(\widehat{v}(s) - c)\,ds\right]df \geq U_A(\widetilde{\mu}_{\geq}(t)).$$

Since t was arbitrary, then the induced contract $(f, \widehat{v})$ satisfies the relaxed dynamic participation constraint.

Lastly, we show that the principal weakly improves under induced contract $(f, \widehat{v})$.

$$\begin{aligned}&\mathbb{E}_{\mathcal{P}}\left[e^{-r\tau}U_P(\mu_\tau) - \int_0^\tau e^{-rs}\psi_s ds\right] \\ &= \mathbb{E}_{\mathcal{P}}\left[e^{-r\tau}U_P(\mu_\tau) - \int_0^\tau e^{-rs}\varphi(v_s) ds\right]\end{aligned}$$

$$\leq \mathbb{E}_{\mathcal{P}}\left[e^{-r\tau}U_P(\mu_\tau)\right] - \mathbb{E}_{\mathcal{P}}\left[\int_0^\tau e^{-rs}\varphi(\widehat{v}(s))ds\right]$$
$$= \int\left(e^{-r\tau}U_P(\mu) - \int_0^\tau e^{-rs}\varphi(\widehat{v}(s))ds\right)f(d\mu, d\tau)$$

where the inequality follows from Jensen's inequality and the fact that $\varphi(\cdot)$ is convex. Thus, the principal weakly improves under induced contract $(f, \widehat{v})$. This implies

$$\mathbb{E}_{\mathcal{P}}\left[e^{-r\tau}U_P(\mu_\tau) - \int_0^\tau e^{-rs}\psi_s ds\right] \leq V_{RP}(\mu_0).$$

Since $(\langle\mu_t\rangle, \langle\psi_t\rangle, \tau)$ was arbitrary, we conclude that $V_{FP}(\mu_0) \leq V_{RP}(\mu_0)$ for all $\mu_0 \in \Delta^1$.

A.2 Proof for Lemma 5.2

We prove existence using the standard Weierstrass extreme value theorem. First, we establish some notation on the constraint sets:

$$C_{Flow,K}(\mu_0) = \left\{f \in \Delta_{\mu_0} : G_{Flow}(f)(t) \geq 0, \forall t \in [2^{-K}, T]\right\}$$
$$C_{DP,K}(\mu_0) = \{(f, v) \in \Delta_{\mu_0} \times \mathbb{V} : G_{DP.K}(f, v)(t) \geq 0 \text{ for } k = 0, 1, \ldots, K-1\}$$
$$C_{RP,K}(\mu_0) = (C_{Flow,K}(\mu_0) \times \mathbb{V}) \cap C_{DP,K}(\mu_0)$$

First, $C_{RP,K}(\mu_0)$ is nonempty; no learning and zero transfers is a feasible choice (i.e., $f = \delta_{\mu_0,0}$ and $v \equiv 0$). Next, we argue that the objective function OBJ : $\Delta_{\mu_0} \times \mathbb{V} \to \mathbb{R}$ is upper semicontinuous. Rewriting it with integration by parts, we get

$$\text{OBJ}(f, v) = \int e^{-r\tau}U_P(\mu)f(d\mu, d\tau) - \int_0^T e^{-rt}\varphi(v(t))(1 - F(t))dt$$

The first additive term is upper semicontinuous because $U_P(\cdot)$ is upper semicontinuous and convergence in $d_{LP}(\cdot,\cdot)$ implies weak convergence. Next, we show that the second additive term is continuous; suppose $d_{LP}(f_n, f) \to 0$ and $\|v_n - v\|_p \to 0$. Let $\varepsilon > 0$ and choose N sufficiently large such that $\|\varphi(v_n) - \varphi(v)\|_\infty < \varepsilon \cdot (2 \cdot \overline{v} \cdot T)^{-1}$ and $\int_0^T |(1 - F_n(t)) - (1 - F(t))|dt < \varepsilon \cdot (2 \cdot \varphi(\overline{v}))^{-1}$ whenever $n \geq N$. The first inequality is justified because φ is continuous and preserves upper semicontinuity. The second inequality is justified because convergence in $d_{LP}(\cdot,\cdot)$ implies $F_n \to F$ pointwise almost everywhere and $v(\cdot) \leq \overline{v}$; hence, convergence follows by the dominated convergence theorem. Then,

for all $n \geq N$, we have

$$\begin{aligned}
&\left|\int_0^T e^{-rt}\varphi(v_n(t)) \cdot (1 - F_n(t))dt - \int_0^T e^{-rt}\varphi(v(t)) \cdot (1 - F(t))dt\right| \\
&\leq \int_0^T \left|\varphi(v_n(t)) \cdot (1 - F_n(t)) - \varphi(v(t)) \cdot (1 - F(t))\right| dt \\
&\leq \int_0^T \left|\varphi(v_n(t)) - \varphi(v(t))\right| dt + \varphi(\overline{v} + 1) \int_0^T \left|(1 - F_n(t)) - (1 - F(t))\right| dt \\
&< \int_0^T \frac{\varepsilon \cdot \overline{v}}{2 \cdot \overline{v} \cdot T} dt + \varphi(\overline{v}) \cdot \frac{\varepsilon}{2 \cdot \varphi(\overline{v})} = \varepsilon
\end{aligned}$$

We annotate the inequalities: the first is by the triangle inequality and the fact that $e^{-rt} \leq 1$ for all $t \in [0, T]$; the second is by "adding the zero term $\pm\varphi(v(t)) \cdot (1 - F_n(t))$" and another application of the triangle inequality; the third is by our choice of convergence. Thus, the second additive term is continuous, and therefore, the objective function is upper semicontinuous.

Next, we tackle the constraint set. We claim that both Δ_{μ_0} and $\mathbb{V}$ are weakly compact. For Δ_{μ_0}, since $[0, T]$ is a closed interval, then $\Delta(\Delta^1 \times [0, T])$ is tight, and hence, by Prohorov's theorem, has compact closure. For $\mathbb{V}$, note that for each $v \in \mathbb{V}$, we have

$$\|v\|_p = \left(\int_0^T |v(t)|^p dt\right)^{1/p} \leq \left(\int_0^T \overline{v}^p dt\right)^{1/p} = \overline{v} \cdot T^{1/p}$$

Thus, $\mathbb{V} \subseteq \{v \in L^p([0, T]) : \|v\|_p \leq \overline{v} \cdot T^{1/p}\}$. Since $L^p([0, T])$ is reflexive, scaling relative to $\overline{v} \cdot T^{1/p}$, we can apply the Eberlein-Šmulian theorem to claim that the closed ball $\{v \in L^p([0, T]) : \|v\|_p \leq \overline{v} \cdot T^{1/p}\}$ is weakly compact.[16] To argue that $\mathbb{V}$ is weakly compact, we only need to show that $\mathbb{V}$ is weakly closed and bounded. Evidently, it is norm-bounded by $\overline{v} \cdot T^{1/p}$. Let $(v_n) \subseteq \mathbb{V}$ be a convergent sequence with limit point $v \in L^p([0, T])$. Define $(v(t) - \overline{v})_+ = \max\{v(t) - \overline{v}, 0\}$. For each n, we have $\int_0^T (v_n(t) - \overline{v})(v(t) - \overline{v})_+ dt \leq 0$. Thus,

$$0 \geq \lim_{n \to \infty} \int_0^T (v_n(t) - \overline{v})(v(t) - \overline{v})_+ dt = \int_0^T (v(t) - \overline{v})_+^2 dt = \|(v - \overline{v})_+\|_2^2$$

Thus, $v(\cdot) \leq \overline{v}$ almost everywhere. The proof is similar to show that $v(\cdot) \geq 0$ almost everywhere. Hence, $v \in \mathbb{V}$, and therefore, $\mathbb{V}$ is weakly closed. Overall, $\mathbb{V}$ is weakly compact.

To show $C_{RP,K}(\mu_0)$ is compact, it suffices to show that both $C_{Flow,K}(\mu_0)$ and $C_{DP,K}(\mu_0)$ are closed. For $C_{DP,K}(\mu_0)$, let $d_{LP}(f_n, f) \to 0$ and $\|v_n - v\|_\infty \to 0$. We can (uniformly)

[16] For reference, see Bühler and Salamon (2018), chapter 3.4.

choose N large enough such that $G_{DP,K}(f_n, v_n) \to G_{DP,K}(f, v)$. The argument to show convergence is similar to that of the objective function. In particular, $U_A(\cdot)$ is continuous and replace $\varphi(\cdot)$ with the identity mapping. Hence, $G_{DP,K}(f_n, v_n) \geq 0$ for all $n \geq N$ implies $G_{DP,K}(f, v) \geq 0$, and therefore, $C_{DP,K}(\mu_0)$ is closed.

We directly prove that $C_{Flow,K}(\mu_0)$ is closed. Let $\{f_n\}_{n=0}^{\infty}$ be sequence such that $f_n \in C_{Flow,K}(\mu_0)$ and f_n weakly converges to f. Define subset $D_f = \{t \in \mathbb{R}_+ : f_\tau(\{t\}) > 0\}$. Note that this subset is at most countable. At every $t \in [2^{-K}, T) \setminus D_f$, the bounded integrands within $G_{Flow}(f)(t)$ are continuous (outside an f-null set). Hence, $G_{Flow}(f_n)(t) \to G_{Flow}(f)(t)$, and therefore, $G_{Flow}(f)(t) \geq 0$ for all $t \in [2^{-K}, T) \setminus D_f$. For $t \in [2^{-K}, T) \cap D_f$, choose a sequence $s_j \downarrow t$ with $s_j \in (t, T) \setminus D_f$. Right continuity of $G_{Flow}(t)$ implies $G_{Flow}(f)(t) = \lim_j G_{Flow}(f)(s_j) \geq 0$. Lastly, at the time horizon T, we have

$$G_{Flow}(f)(T) = \chi \int \tau f(d\mu, d\tau) + H(\mu_0) - \int H(\mu) f(d\mu, d\tau)$$

which is continuous in f. Conclude that $C_{Flow,K}(\mu_0)$ is closed.

A.3 Proof for Lemma 5.3

Before proving strong duality, we introduce an auxiliary problem equivalent to (RP). For the relaxed problem, we define its constraint set, value function, and the arguments of the maxima.

$$\begin{aligned}
C_{RP}(\mu_0) &= \big\{(f, v) \in \Delta_{\mu_0} \times \mathbb{V} : G_{Flow}(f)(t) \geq 0, \forall t \in [2^{-K}, T] \text{ and } G_{DP,K}(f, v) \geq 0\big\} \\
V_{RP}(\mu_0) &= \max\big\{\text{OBJ}(f, v) : (f, v) \in C_{RP}(\mu_0)\big\} \\
M_{RP}(\mu_0) &= \arg\max\big\{\text{OBJ}(f, v) : (f, v) \in C_{RP}(\mu_0)\big\}
\end{aligned}$$

We now define the auxiliary problem. Given partition $(t_k)_{k=0}^K$, belief-time distribution $f \in \Delta_{\mu_0}$ and $\eta \in L^p([0, T])$, define the auxiliary discretized dynamic participation constraint $\widetilde{G}_{DP}(f, \eta) \in \mathbb{R}^K$ by

$$\begin{aligned}
\widetilde{G}_{DP}(f, \eta) = \Bigg(\int_{\tau \geq t_k} \left(e^{-r\tau} U_A(\mu) - c \int_{t_k}^{\tau} e^{-rs} ds \right) f(d\mu, d\tau) \\
- e^{-rt_k} U_A(\widetilde{\mu}_{\geq}(t_k)) + \int_{t_k}^{T} e^{-rt} \eta(t) dt \Bigg)_{k=0}^{K-1}
\end{aligned}$$

We also denote the k-th coordinate of $\widetilde{G}_{DP}(f, \eta)$ by $\widetilde{G}_{DP}(f, \eta)(t_k)$. Given $(f, \eta) \in \Delta_{\mu_0} \times \mathbb{V}$ with respective cumulative distribution $F(t) = \int_{\tau \leq t} f(d\mu, d\tau)$, define mapping $I(f, \eta)$:

$[0,T] \to [0,+\infty]$ by

$$I(f,\eta)(t) = \begin{cases} 0 & \text{if } F(t) = 1 \text{ and } \eta(t) = 0 \\ \dfrac{\eta(t)}{1-F(t)} & \text{if } F(t) < 1 \\ +\infty & \text{if } F(t) = 1 \text{ and } \eta(t) > 0 \end{cases}$$

Define the auxiliary objective function $\widetilde{\text{OBJ}} : \Delta_{\mu_0} \times \mathbb{V} \to [-\infty, +\infty)$ by

$$\widetilde{\text{OBJ}}(f,\eta) = \int e^{-r\tau} U_P(\mu) f(d\mu, d\tau) - \int_0^T e^{-rt} \varphi\big(I(f,\eta)(t)\big)(1-F(t))dt \tag{2}$$

Given initial belief $\mu_0 \in \text{ri } \Delta^1$, define the following sets:

$$\begin{aligned} Z_{\mu_0} &:= \{(f,\eta) \in \Delta_{\mu_0} \times \mathbb{V} : 0 \le I(f,\eta)(t) \le \overline{v}, t \in [0,T] \text{ a.e.}\} \\ \widetilde{C}_{RP}(\mu_0) &= \big\{(f,\eta) \in Z_{\mu_0} : G_{Flow}(f)(t) \ge 0, \forall t \in [2^{-K}, T] \text{ and } \widetilde{G}_{DP}(f,\eta) \ge 0\big\} \\ \widetilde{V}_{RP}(\mu_0) &= \max\big\{\widetilde{\text{OBJ}}(f,\eta) : (f,\eta) \in \widetilde{C}_{RP}(\mu_0)\big\} \\ \widetilde{M}_{RP}(\mu_0) &= \arg\max\big\{\widetilde{\text{OBJ}}(f,\eta) : (f,\eta) \in \widetilde{C}_{RP}(\mu_0)\big\} \end{aligned}$$

The next lemma states that solving the auxiliary problem is necessary and sufficient for solving problem (RP).

Lemma A.1. *(Necessity)* If $(f,v) \in M_{RP}(\mu_0)$, then $(f, v\cdot(1-F)) \in \widetilde{M}_{RP}(\mu_0)$. *(Sufficiency)* If $(f,\eta) \in \widetilde{M}_{RP}(\mu_0)$, then $(f, I(f,\eta)) \in M_{RP}(\mu_0)$. Moreover, $V_{RP}(\mu_0) = \widetilde{V}_{RP}(\mu_0)$.

Proof. We start by proving that the choice variables in one problem can be tweaked to be feasible for the other problem. Let $(f,v) \in C_{RP}(\mu_0)$ with respective cumulative function $F(t) = \int_{\tau \le t} f(d\mu, d\tau)$. We argue that $(f, v\cdot(1-F))$ is feasible for the auxiliary problem; that is, $(f, v\cdot(1-F)) \in \widetilde{C}_{RP}(\mu_0)$. Clearly, $f \in \Delta_{\mu_0}$ and $G_{Flow}(f)(\cdot) \ge 0$. If $F(t) = 1$ for some t, then $v(t)\cdot(1-F(t)) = 0$. Hence,

$$I(f, v\cdot(1-F))(t) = \begin{cases} 0 & \text{if } F(t) = 1 \\ v(t) & \text{if } F(t) < 1 \end{cases}$$

Hence, $0 \le I(f, v\cdot(1-F)) \le \overline{v}$. With a substitution, $G_{DP}(f,v) \ge 0$ implies $\widetilde{G}_{DP}(f, v\cdot(1-F)) \ge 0$. Hence, we can claim that $(f, v\cdot(1-F)) \in \widetilde{C}_{RP}(\mu_0)$. Starting the other direction, let $(f,\eta) \in \widetilde{C}_{RP}(\mu_0)$. We show that $(f, I(f,\eta)) \in C_{RP}(\mu_0)$. Obviously, $f \in \Delta_{\mu_0}$ and $G_{Flow}(f)(\cdot) \ge 0$. Evidently, $I(f,\eta) \in \mathbb{V}$. Lastly, using substitution, $\widetilde{G}_{DP}(f,\eta) \ge 0$ implies $G_{DP}(f, I(f,\eta)) \ge 0$. Hence, $(f, I(f,\eta)) \in C_{RP}(\mu_0)$.

Necessity. Let $(f, v) \in M_{RP}(\mu_0)$. We argue that $(f, v \cdot (1-F))$ attains $\widetilde{V}_{RP}(\mu_0)$. Given an arbitrary $(f, \eta) \in \widetilde{C}_{RP}(\mu_0)$, we get

$$\widetilde{\text{OBJ}}(f, \eta) = \text{OBJ}(f, I(f, \eta)) \leq \text{OBJ}(f, v) = \widetilde{\text{OBJ}}(f, v \cdot (1-F))$$

where the inequality holds since $(f, I(f, \eta)) \in C_{RP}(\mu_0)$ and $(f, v) \in M_{RP}(\mu_0)$. Hence, $(f, v \cdot (1-F))$ attains $\widetilde{V}_{RP}(\mu_0)$.

Sufficiency. Let $(f, \eta) \in \widetilde{M}_{RP}(\mu_0)$. Given an arbitrary $(f, v) \in C_{RP}(\mu_0)$, we get

$$\text{OBJ}(f, v) = \widetilde{\text{OBJ}}(f, v \cdot (1-F)) \leq \widetilde{\text{OBJ}}(f, \eta) = \text{OBJ}(f, I(f, \eta))$$

where, again, we emphasize the inequality holds since $(f, v \cdot (1-F)) \in \widetilde{C}_{RP}(\mu_0)$ and $(f, \eta) \in \widetilde{M}_{RP}(\mu_0)$. Thus, $(f, I(f, \eta))$ attains $V_{RP}(\mu_0)$. The necessity and sufficiency proofs imply $V_{RP}(\mu_0) = \widetilde{V}_{RP}(\mu_0)$. □

Now to prove strong duality, our goal is to apply theorem 1 from Luenberger (1969), chapter 8.6, page 224. To carry this out, we show that our setup satisfies the assumptions of the theorem. Define mapping $\widetilde{G}(f, v)$ from $[2^{-K}, T] \times \{(t_k)_{k=0}^{K-1}\}$ into $\mathbb{R}^{K+1}$ by

$$\widetilde{G}(f, \eta)(t) = \big(G_{Flow}(f)(t), \widetilde{G}_{DP}(f, \eta)\big)$$

The auxiliary Lagrangian is

$$\widetilde{\mathcal{L}}_K(f, \eta, \lambda, \gamma) = \begin{cases} \int e^{-r\tau} U_P(\mu) f(d\mu, d\tau) + \int_0^T e^{-rs}\big(\Gamma(s)\eta(s) - \varphi(I(f,\eta)(s)) \cdot (1 - F(s))\big) ds \\ + \int \left(e^{-r\tau}\Gamma(\tau) U_A(\mu) - c \int_0^\tau e^{-rs}\Gamma(s) ds \right) f(d\mu, d\tau) - \sum_{k=0}^{K-1} e^{-rt_k} U_A(\widetilde{\mu}_{\geq}(t_k))\gamma_k \\ + \int_{t \in [2^{-K}, T]} G_{Flow}(f)(t) d\lambda(t) \end{cases}$$

Let $UC([2^{-K}, T])$ be the set of uniformly continuous mappings on $[2^{-K}, T]$. Define subset $Z_{\mu_0}^C \subseteq \Delta_{\mu_0} \times \mathbb{V}$ by:

$$Z_{\mu_0}^C := \{(f, \eta) \in Z_{\mu_0} : f(\Delta^1 \times \{t\}) = 0 \text{ for every } t \in [0, T]\}.$$

Lastly, when referencing the positive cone of $UC([2^{-K}, T]) \times \mathbb{R}_+^K$, we use the norm $\|\cdot\|_\infty + \|\cdot\|_E$ where $\|\cdot\|_\infty$ denotes the essential supremum and $\|\cdot\|_E$ the Euclidean norm.

We verify that the objective function $\widetilde{\text{OBJ}}(f, \eta)$ is continuous and concave on $\Delta_{\mu_0} \times \mathbb{V}$. As per the arguments in section A.2, the objective function is continuous. Next we show

concavity. The first additive term is concave since f enters linearly. For the second term, we argue that the composite mapping $t \mapsto \varphi(I(f,\eta)(t))(1-F(t))$ is convex. Let $f, f' \in \Delta_{\mu_0}$ with respective cumulative functions $F(t) = \int_{\tau \le t} f(d\mu, d\tau)$ and $F'(t) = \int_{\tau \le t} f'(d\mu, d\tau)$, respectively. Fix $\eta, \eta' \in \mathbb{V}$. Given $\alpha \in [0,1]$, define the convex combinations $f^\alpha = \alpha f + (1-\alpha) f'$, $\eta^\alpha = \alpha\eta + (1-\alpha)\eta'$, and $F^\alpha = \alpha F + (1-\alpha)F'$. For $t \in [0,T]$ such that $F(t) < 1$ and $F'(t) < 1$, we have

$$
\begin{aligned}
&\varphi\big(I(f^\alpha, \eta^\alpha)(t)\big) \cdot (1 - F^\alpha(t)) \\
&= \varphi\left(\frac{\alpha F(t)}{1 - F^\alpha(t)} \cdot \frac{\eta(t)}{F(t)} + \frac{(1-\alpha)F'(t)}{1 - F^\alpha(t)} \cdot \frac{\eta'(t)}{F'(t)}\right) \cdot (1 - F^\alpha(t)) \\
&\le \left(\frac{\alpha F(t)}{1 - F^\alpha(t)} \cdot \varphi\left(\frac{\eta(t)}{F(t)}\right) + \frac{(1-\alpha)F'(t)}{1 - F^\alpha(t)} \cdot \varphi\left(\frac{\eta'(t)}{F'(t)}\right)\right) \cdot (1 - F^\alpha(t)) \\
&= \alpha \cdot \Big(\varphi\big(I(f,\eta)(t)\big) \cdot (1 - F(t))\Big) + (1-\alpha) \cdot \Big(\varphi\big(I(f',\eta')(t)\big) \cdot (1 - F'(t))\Big)
\end{aligned}
$$

where the inequality holds since φ is convex. If $t \in [0,T]$ such that $1 - F'(t) = \eta'(t) = 0$ and $1 > F(t)$, then

$$
\begin{aligned}
&\varphi\big(I(f^\alpha, \eta^\alpha)(t)\big) \cdot (1 - F^\alpha(t)) \\
&= \varphi\left(\frac{\alpha \cdot \eta(t)}{\alpha \cdot (1 - F(t))}\right) \cdot \alpha \cdot (1 - F(t)) \\
&= \alpha \cdot \Big(\varphi\big(I(f,\eta)(t)\big) \cdot (1 - F(t))\Big) + (1-\alpha) \cdot 0 \\
&= \alpha \cdot \Big(\varphi\big(I(f,\eta)(t)\big) \cdot (1 - F(t))\Big) + (1-\alpha) \cdot \Big(\varphi\big(I(f',\eta')(t)\big) \cdot (1 - F'(t))\Big)
\end{aligned}
$$

For the last case, if $t \in [0,T]$ such that $\eta(t) > 0$ and $F(t) = 1$, then $\varphi(I(f,\eta)(t)) \cdot (1 - F(t)) = +\infty$. Overall, the composite mapping $t \mapsto \varphi(I(f,\eta)(t))(1 - F(t))$ is convex, and therefore the second additive term is concave.

Next, we tackle the constraint set. By lemma A.2, operator $\widetilde{G}$ is concave. Moreover, lemma A.3 says there exists $(f,\eta) \in Z^C_{\mu_0}$ such that $\widetilde{G}(f,\eta)$ is in the interior of the positive cone of $UC([2^{-K}, T]) \times \mathbb{R}^K_+$. Note that $G_{Flow}(f) \in UC([2^{-K}, T])$ implies that the appropriate dual space is $\mathcal{B}([2^{-K}, T]) \times \mathbb{R}^K_+$. Thus, by theorem 1 of Luenberger (1969), chapter 8.6, page 224, we have

$$
\sup_{(f,\eta) \in Z^C_{\mu_0}, \widetilde{G}(f,\eta) \ge 0} \widetilde{\mathrm{OBJ}}(f,\eta) = \min_{(\lambda,\gamma) \in \mathcal{B} \times \mathbb{R}^K_+} \sup_{(f,\eta) \in Z^C_{\mu_0}} \widetilde{\mathcal{L}}(f,\eta,\lambda,\gamma)
$$

where the minimum is achieved by $(\lambda, \gamma) \in \mathcal{B}([2^{-K}, T]) \times \mathbb{R}^K$.

Then, for all $(f,\eta) \in Z_{\mu_0}$, $\lambda \in \mathcal{B}([2^{-K},T])$, $\gamma \in \mathbb{R}^K$, and $\varepsilon > 0$, Lemma A.4 implies there exists $\widehat{f} \in \Delta_{\mu_0}$ such that $G_{Flow}(\widehat{f}) \in UC([2^{-K},T])$ and $\widetilde{\mathcal{L}}_K(\widehat{f},\widehat{\eta},\lambda,\gamma) \geq \widetilde{\mathcal{L}}_K(f,\eta,\lambda,\gamma) - \varepsilon$. Therefore, we obtain

$$\begin{aligned}
\sup_{(f,\eta)\in Z_{\mu_0}^C} \inf_{(\lambda,\gamma)\in\mathcal{B}([2^{-K},T])\times\mathbb{R}^K} \widetilde{\mathcal{L}}_K(f,\eta,\lambda,\gamma) &\leq \sup_{(f,\eta)\in Z_{\mu_0}} \inf_{(\lambda,\gamma)\in\mathcal{B}([2^{-K},T])\times\mathbb{R}^K} \widetilde{\mathcal{L}}_K(f,\eta,\lambda,\gamma) \\
&\leq \min_{(\lambda,\gamma)\in\mathcal{B}([2^{-K},T])\times\mathbb{R}^K} \sup_{(f,\eta)\in Z_{\mu_0}} \widetilde{\mathcal{L}}_K(f,\eta,\lambda,\gamma) \\
&= \min_{(\lambda,\gamma)\in\mathcal{B}([2^{-K},T])\times\mathbb{R}^K} \sup_{(f,\eta)\in Z_{\mu_0}^C} \widetilde{\mathcal{L}}_K(f,\eta,\lambda,\gamma)
\end{aligned}$$

By lemma A.1 and lemma 5.2, there exists $(f,v) \in M_{RP}(\mu_0)$ such that $(f, v\cdot(1-F))$ also solves (Primal$_K$). Hence, we have strong duality.

The rest of the section is dedicated to proving our lemmas. We fix some notations. Let $E_K := C([2^{-K},T]) \times \mathbb{R}^K$ and $E_{K,+} := C_+([2^{-K},T]) \times \mathbb{R}^K_+$ where $C_+([2^{-K},T])$ consists of the nonnegative continuous functions.

Lemma A.2. $Z_{\mu_0}^C$ is convex, $\widetilde{G}$ is well-defined as a mapping from $Z_{\mu_0}^C$ into E_K, and $\widetilde{G}$ is concave with respect to the order induced by $E_{K,+}$.

Proof. The upper bound on transfers within Z_{μ_0} is equivalent to its definition using the mapping $I(f,\eta)$. Where $F_f(t) < 1$, dividing the cap inequality by $1 - F_f(t)$ gives $0 \leq I(f,\eta)(t) \leq \overline{v}$. Where $F_f(t) = 1$, the cap forces $\eta(t) = 0$, and the definition of I assigns zero. Conversely, finiteness and the bounds on I, together with $\eta \geq 0$, imply the cap inequality. Also, the cap implies $0 \leq \eta \leq \overline{v}$ almost everywhere, so $\eta \in \mathbb{V}$.

Fix $(f_1,\eta_1),(f_2,\eta_2) \in Z_{\mu_0}^C$ and $\alpha \in [0,1]$, and put

$$f_\alpha := \alpha f_1 + (1-\alpha) f_2, \qquad \eta_\alpha := \alpha\eta_1 + (1-\alpha)\eta_2.$$

The measure f_α is a probability measure and satisfies Bayes plausibility, since

$$\int \mu \, df_\alpha = \alpha \int \mu \, df_1 + (1-\alpha) \int \mu \, df_2 = \mu_0.$$

Moreover, $\eta_\alpha \in L^p([0,T])$, and linearity of integration yields

$$F_{f_\alpha}(t) = \alpha F_{f_1}(t) + (1-\alpha) F_{f_2}(t).$$

Therefore, almost everywhere,

$$0 \leq \eta_\alpha(t) \leq \overline{v}\,[\alpha(1-F_{f_1}(t)) + (1-\alpha)(1-F_{f_2}(t))] = \overline{v}(1-F_{f_\alpha}(t)).$$

Thus, $(f_\alpha, \eta_\alpha) \in Z_{\mu_0}$. For every $t \in [0,T]$, atomlessness is also preserved:

$$f_\alpha(\Delta^1 \times \{t\}) = \alpha f_1(\Delta^1 \times \{t\}) + (1-\alpha) f_2(\Delta^1 \times \{t\}) = 0.$$

Consequently, $(f_\alpha, \eta_\alpha) \in Z^C_{\mu_0}$, proving convexity of the domain.

We now verify that the constraint mapping takes values in E_K. Fix $(f,\eta) \in Z^C_{\mu_0}$, and let $t_n \to t$ in $[0,T]$. Outside the f-null set $\Delta^1 \times \{t\}$, the indicators of $\{\tau \le t_n\}$ and $\{\tau > t_n\}$ converge to those of $\{\tau \le t\}$ and $\{\tau > t\}$, respectively. Boundedness of H and μ on the simplex and dominated convergence give

$$\int_{\tau \le t_n} H(\mu)\, df \longrightarrow \int_{\tau \le t} H(\mu)\, df, \qquad \widetilde{\mu}^f_>(t_n) \longrightarrow \widetilde{\mu}^f_>(t).$$

The "accumulating time" term is continuous because

$$\left| \int \min\{\tau, t_n\}\, df - \int \min\{\tau, t\}\, df \right| \le |t_n - t|.$$

Continuity of the homogeneous extension of $H(\cdot)$ consequently implies that $G_{Flow}(f) \in C([0,T])$, including at T and at any time when the continuation mass vanishes. In particular, its restriction belongs to $C([2^{-K}, T])$. Each auxiliary participation coordinate is finite because U_A is bounded on the simplex, the horizon is finite, and η is bounded by $\overline{v}$. Hence, $\widetilde{G}(f,\eta) \in E_K$.

It remains to prove concavity. The continuation-belief vectors are linear in the underlying law:

$$\begin{aligned} \widetilde{\mu}^{f_\alpha}_>(t) &= \alpha \widetilde{\mu}^{f_1}_>(t) + (1-\alpha)\widetilde{\mu}^{f_2}_>(t), \\ \widetilde{\mu}^{f_\alpha}_\ge(t_k) &= \alpha \widetilde{\mu}^{f_1}_\ge(t_k) + (1-\alpha)\widetilde{\mu}^{f_2}_\ge(t_k). \end{aligned}$$

For every $t \in [2^{-K}, T]$, the other terms in the flow residual are affine in f. Thus,

$$\begin{aligned} &G_{Flow}(f_\alpha)(t) - \alpha G_{Flow}(f_1)(t) - (1-\alpha) G_{Flow}(f_2)(t) \\ &\quad = \alpha H(\widetilde{\mu}^{f_1}_>(t)) + (1-\alpha) H(\widetilde{\mu}^{f_2}_>(t)) - H\left(\alpha \widetilde{\mu}^{f_1}_>(t) + (1-\alpha)\widetilde{\mu}^{f_2}_>(t)\right) \ge 0, \end{aligned}$$

where the inequality follows from convexity of H on $\mathbb{R}^2_+$. Since all three flow residuals are continuous, their displayed difference belongs to $C_+([2^{-K}, T])$.

For each $k \in \{0, \dots, K-1\}$, all terms in $\widetilde{G}_{DP}$ other than the negative outside-option

term are linear in (f, η). Therefore,

$$\begin{aligned}\widetilde{G}_{DP}(f_\alpha, \eta_\alpha)(t_k) - \alpha\widetilde{G}_{DP}(f_1, \eta_1)(t_k) - (1-\alpha)\widetilde{G}_{DP}(f_2, \eta_2)(t_k) \\ = e^{-rt_k}\Big[\alpha U_A(\widetilde{\mu}_{\geq}^{f_1}(t_k)) + (1-\alpha)U_A(\widetilde{\mu}_{\geq}^{f_2}(t_k)) \\ - U_A\left(\alpha\widetilde{\mu}_{\geq}^{f_1}(t_k) + (1-\alpha)\widetilde{\mu}_{\geq}^{f_2}(t_k)\right)\Big] \geq 0,\end{aligned}$$

by convexity of U_A on $\mathbb{R}_+^2$. Every coordinate of the participation difference is therefore nonnegative. Combining this with the flow inequality shows that the full constraint difference belongs to $E_{K,+}$, proving the asserted concavity. □

Lemma A.3. There exists $(\widehat{f}, \widehat{\eta}) \in Z_{\mu_0}^C$ such that $\widetilde{G}(\widehat{f}, \widehat{\eta}) \in \mathrm{int}_{E_K}(E_{K,+})$.

Proof. The upper bound on transfers gives $\overline{v} = u(\overline{\psi}) > c + rU_A(\mu_0)$. Choose

$$\varepsilon_* := \frac{\overline{v} - c - rU_A(\mu_0)}{2} > 0, \qquad v_* := c + rU_A(\mu_0) + \varepsilon_* < \overline{v}.$$

Define the probability law and auxiliary transfer by

$$\widehat{f}(d\mu, d\tau) := \delta_{\mu_0}(d\mu)\,\frac{d\tau}{T}, \qquad \tau \in [0,T], \qquad \widehat{\eta}(s) := v_*\left(1 - \frac{s}{T}\right), \quad s \in [0,T].$$

The belief marginal is concentrated at μ_0, so $\int \mu\, \widehat{f}(d\mu, d\tau) = \mu_0$. Its stopping-time marginal is uniform on $[0,T]$ and therefore has no atoms, including at 0 and T. Its cumulative distribution is

$$\widehat{F}(s) = \int_{\tau \leq s} \widehat{f}(d\mu, d\tau) = \frac{s}{T}, \qquad s \in [0,T].$$

The function $\widehat{\eta}$ is bounded and continuous, hence belongs to $L^p([0,T])$, and satisfies

$$0 \leq \widehat{\eta}(s) = v_*(1 - \widehat{F}(s)) \leq \overline{v}(1 - \widehat{F}(s)) \leq \overline{v}.$$

In terms of the mapping I defining Z_{μ_0},

$$I(\widehat{f}, \widehat{\eta})(s) = \begin{cases} v_*, & 0 \leq s < T, \\ 0, & s = T, \end{cases}$$

where the value at T follows from $\widehat{F}(T) = 1$ and $\widehat{\eta}(T) = 0$. Thus, $(\widehat{f}, \widehat{\eta}) \in Z_{\mu_0}^C$. For this law, the two unweighted continuation beliefs coincide:

$$\widetilde{\mu}_{>}(t) = \widetilde{\mu}_{\geq}(t) = \left(1 - \frac{t}{T}\right)\mu_0, \qquad t \in [0,T].$$

By degree-one homogeneity of $H(\cdot)$,

$$\int_{\tau\le t} H(\mu)\, d\widehat{f} = \frac{t}{T} H(\mu_0), \qquad H(\widetilde{\mu}_{>}(t)) = \left(1 - \frac{t}{T}\right) H(\mu_0).$$

Consequently, the information terms in the flow residual cancel, and

$$\begin{aligned} G_{Flow}(\widehat{f})(t) &= \frac{\chi}{T} \int_0^T \min\{\tau, t\}\, d\tau \\ &= \frac{\chi}{T} \left(\int_0^t \tau\, d\tau + \int_t^T t\, d\tau \right) = \chi \left(t - \frac{t^2}{2T} \right). \end{aligned}$$

This function is continuous on $[0, T]$ and increasing there, since its derivative is $\chi(1 - t/T) \ge 0$. Hence,

$$\inf_{t\in[2^{-K},T]} G_{Flow}(\widehat{f})(t) = \chi \left(2^{-K} - \frac{2^{-2K}}{2T} \right) > 0.$$

Next fix $k \in \{0, \dots, K-1\}$. Fubini's theorem gives

$$\int_{t_k}^T e^{-rs} \widehat{\eta}(s)\, ds = \frac{1}{T} \int_{t_k}^T \int_{t_k}^{\tau} e^{-rs} v_*\, ds\, d\tau.$$

Substituting this identity into the auxiliary participation residual and using homogeneity of $U_A(\cdot)$, we obtain

$$\begin{aligned} \widetilde{G}_{DP}(\widehat{f}, \widehat{\eta})(t_k) = \frac{1}{T} \int_{t_k}^T \left(e^{-r\tau} U_A(\mu_0) + \int_{t_k}^{\tau} e^{-rs} (v_* - c)\, ds \right) d\tau \\ - e^{-rt_k} \left(1 - \frac{t_k}{T} \right) U_A(\mu_0). \end{aligned}$$

Since $v_* - c = rU_A(\mu_0) + \varepsilon_*$, for every $\tau \ge t_k$,

$$e^{-r\tau} U_A(\mu_0) + \int_{t_k}^{\tau} e^{-rs} (v_* - c)\, ds = e^{-rt_k} U_A(\mu_0) + \varepsilon_* \int_{t_k}^{\tau} e^{-rs}\, ds.$$

The constant term cancels the outside option. Applying Fubini once more therefore yields

$$\begin{aligned} \widetilde{G}_{DP}(\widehat{f}, \widehat{\eta})(t_k) &= \frac{\varepsilon_*}{T} \int_{t_k}^T \int_{t_k}^{\tau} e^{-rs}\, ds\, d\tau \\ &= \varepsilon_* \int_{t_k}^T e^{-rs} \left(1 - \frac{s}{T} \right) ds > 0. \end{aligned}$$

Strict positivity holds because $t_k < T$ and $\varepsilon_* > 0$. There are finitely many participation

coordinates, so

$$\rho_K := \min\left\{\chi\left(2^{-K} - \frac{2^{-2K}}{2T}\right), \min_{0\le k<K} \widetilde{G}_{DP}(\widehat{f},\widehat{\eta})(t_k)\right\} > 0.$$

Let $(h,b) \in E_K$ satisfy $\|(h,b) - \widetilde{G}(\widehat{f},\widehat{\eta})\|_{E_K} = \sup_{t\in[2^{-K},T]} |h(t)| + \|b\|_2 < \rho_K$. For every $t \in [2^{-K},T]$,

$$h(t) \ge G_{Flow}(\widehat{f})(t) - \|h - G_{Flow}(\widehat{f})\|_\infty > 0.$$

Likewise, for every participation coordinate,

$$b_k \ge \widetilde{G}_{DP}(\widehat{f},\widehat{\eta})(t_k) - \|b - \widetilde{G}_{DP}(\widehat{f},\widehat{\eta})\|_2 > 0.$$

Thus, the open ball of radius ρ_K about $\widetilde{G}(\widehat{f},\widehat{\eta})$ is contained in $E_{K,+}$. The constraint image is therefore an interior point of the positive cone, as required. □

Lemma A.4. For every $(f,\eta) \in Z_{\mu_0}$, every finite positive Borel measure λ on $[2^{-K},T]$, every $\gamma \in \mathbb{R}^K_+$, and every $\varepsilon > 0$, there exists $(\widehat{f},\widehat{\eta}) \in Z^C_{\mu_0}$ such that $G_{Flow}(\widehat{f})$ is uniformly continuous on $[2^{-K},T]$ and

$$\widetilde{\mathcal{L}}_K(\widehat{f},\widehat{\eta},\lambda,\gamma) \ge \widetilde{\mathcal{L}}_K(f,\eta,\lambda,\gamma) - \varepsilon.$$

Proof. Fix $(f,\eta) \in Z_{\mu_0}$. Throughout the proof, we write $F_g(s) = \int_{\tau\le s} g(d\mu,d\tau)$ for the stopping-time cumulative distribution of a law g. The upper bound on transfers implies

$$0 \le \eta(s) \le \overline{v}(1 - F_f(s)) \quad \text{for almost every } s \in [0,T].$$

Choose a nonnegative representative of η. We will construct $(f_\delta,\eta_\delta) \in Z^C_{\mu_0}$ for

$$0 < \delta < \frac{T - t_{K-1}}{3}$$

and show that its Lagrangian has the required lower approximation as $\delta \downarrow 0$. On the product probability space with law $f(d\mu,d\tau)\,du$, where $u \in [0,1]$, let U be the uniform coordinate and define

$$\tau_\delta := \begin{cases} \tau + \dfrac{\delta}{2}(1+U), & \tau \le T - 2\delta, \\ T - \delta + \dfrac{\delta}{2}U, & \tau > T - 2\delta. \end{cases}$$

Let f_δ be the law of (μ, τ_δ). Equivalently, for every bounded Borel function g,

$$\int g(\mu, s)\, f_\delta(d\mu, ds) = \int \int_0^1 g(\mu, \tau_\delta(\tau, u))\, du\, f(d\mu, d\tau).$$

The map defining this pushforward is Borel measurable. For every fixed original (μ, τ), the new stopping time is uniform on an interval of length $\delta/2$. Hence, its conditional probability of equaling any fixed time is zero, and integration over f shows that f_δ has an atomless stopping-time marginal. Also,

$$\frac{\delta}{2} \leq \tau_\delta \leq T - \frac{\delta}{2}, \qquad |\tau_\delta - \tau| \leq 2\delta.$$

The belief coordinate is unchanged. Consequently, f_δ has the same belief marginal as f, satisfies (BP), and converges weakly to f as $\delta \downarrow 0$. The last assertion follows by applying uniform continuity to any continuous test function on the compact belief-time space. Set

$$\eta_\delta(s) := \min\{\eta(s), \overline{v}(1 - F_{f_\delta}(s))\}.$$

This is measurable, belongs to $L^p([0,T])$, and satisfies the cap associated with f_δ. Hence, $(f_\delta, \eta_\delta) \in Z^C_{\mu_0}$. By the preceding concave-mapping lemma, $G_{Flow}(f_\delta)$ is continuous, and therefore uniformly continuous, on $[2^{-K}, T]$.

We first establish convergence of the auxiliary transfers. For $s \leq T - 2\delta$, the construction implies that $\tau > s$ entails $\tau_\delta > s$: stopping times outside the terminal band are delayed, and every new stopping time from the terminal band exceeds $T - 2\delta$. Thus,

$$1 - F_{f_\delta}(s) \geq 1 - F_f(s), \qquad s \leq T - 2\delta.$$

It follows that $\eta_\delta = \eta$ almost everywhere on $[0, T - 2\delta]$. Since both functions lie between zero and $\overline{v}$ almost everywhere,

$$\|\eta_\delta - \eta\|_{L^p([0,T])} \leq \overline{v}(2\delta)^{1/p} \longrightarrow 0.$$

They also converge in L^1, and $\eta_\delta(s) \to \eta(s)$ for almost every s. Moreover, $|\tau_\delta - \tau| \to 0$ implies

$$F_{f_\delta}(s) \longrightarrow F_f(s)$$

at every time that is not an atom of the original stopping-time marginal. That marginal has at most countably many atoms, so this convergence holds for Lebesgue-almost every s.

Next fix a participation time t_k, with $k < K$. The terminal-band draws remain above t_{K-1}, since $T - \delta > t_{K-1}$. If the original stopping time equals t_k, then, for all sufficiently small δ, it lies in the delay branch and is moved strictly to the right. Together with $\tau_\delta \to \tau$, this proves

$$\mathbf{1}\{\tau_\delta \geq t_k\} \longrightarrow \mathbf{1}\{\tau \geq t_k\} \quad \text{almost surely.}$$

In particular, original atoms at participation times are retained in the limiting weak tail. Dominated convergence gives

$$\widetilde{\mu}_{\geq}^{f_\delta}(t_k) \longrightarrow \widetilde{\mu}_{\geq}^{f}(t_k)$$

and convergence of both the discounted stopping-utility integral and the duration-cost integral in $\widetilde{G}_{DP}$. These integrands are bounded, and their time-dependent terms converge under the coupling. Continuity of the homogeneous extension of $U_A(\cdot)$ handles the outside-option term. Finally,

$$\left| \int_{t_k}^{T} e^{-rs} (\eta_\delta(s) - \eta(s))\, ds \right| \leq \|\eta_\delta - \eta\|_{L^1([0,T])} \longrightarrow 0.$$

Consequently,

$$\widetilde{G}_{DP}(f_\delta, \eta_\delta) \longrightarrow \widetilde{G}_{DP}(f, \eta) \quad \text{in } \mathbb{R}^K.$$

We now prove convergence of the auxiliary objective. The principal's stopping payoff $U_P(\cdot)$ is bounded: it is either zero or one of the bounded affine project payoffs on the simplex. Since the belief coordinate is unchanged,

$$\left| \int e^{-rs} U_P(\mu)\, f_\delta(d\mu, ds) - \int e^{-r\tau} U_P(\mu)\, f(d\mu, d\tau) \right| \leq 2r\delta\, \|U_P\|_\infty \longrightarrow 0.$$

This estimate does not require continuity of U_P in the belief. For the transfer cost, define the capped perspective

$$\Pi(x, y) := \begin{cases} x\varphi(y/x), & x > 0, \\ 0, & x = y = 0, \end{cases} \qquad 0 \leq x \leq 1, \quad 0 \leq y \leq \overline{v}x.$$

This function is continuous where $x > 0$. It is also continuous at $(0, 0)$, because

$$0 \leq \Pi(x, y) \leq x\varphi(\overline{v}).$$

The established almost-everywhere convergence of the cumulative distributions and trans-

fers therefore yields

$$\Pi(1 - F_{f_\delta}(s), \eta_\delta(s)) \longrightarrow \Pi(1 - F_f(s), \eta(s)) \quad \text{for almost every } s.$$

The same bound supplies the integrable constant $\varphi(\overline{v})$. Dominated convergence gives convergence of the discounted transfer-cost integral. Together with the stopping-payoff estimate, this proves

$$\widetilde{\mathrm{OBJ}}(f_\delta, \eta_\delta) \longrightarrow \widetilde{\mathrm{OBJ}}(f, \eta).$$

It remains to control the flow term against λ. Fix $t \in [2^{-K}, T)$. For all sufficiently small δ, an original atom at t lies in the delay branch and moves strictly past t. Hence

$$\mathbf{1}\{\tau_\delta \le t\} \longrightarrow \mathbf{1}\{\tau < t\}, \qquad \mathbf{1}\{\tau_\delta > t\} \longrightarrow \mathbf{1}\{\tau \ge t\} \quad \text{almost surely.}$$

The occupation term converges because $|\min\{\tau_\delta, t\} - \min\{\tau, t\}| \le 2\delta$. Dominated convergence for the remaining integrals and continuity of H thus imply

$$\begin{aligned} G_{Flow}(f_\delta)(t) \longrightarrow \chi \int \min\{\tau, t\}\, f(d\mu, d\tau) + H(\mu_0) \\ - \int_{\tau<t} H(\mu)\, f(d\mu, d\tau) - H(\widetilde{\mu}^f_{\ge}(t)). \end{aligned}$$

This limit is at least $G_{Flow}(f)(t)$. To see this, use

$$\widetilde{\mu}^f_{\ge}(t) = \widetilde{\mu}^f_{>}(t) + \int_{\tau=t} \mu\, f(d\mu, d\tau)$$

and the inequalities

$$H(\widetilde{\mu}^f_{\ge}(t)) \le H(\widetilde{\mu}^f_{>}(t)) + H\left(\int_{\tau=t} \mu\, f(d\mu, d\tau)\right) \le H(\widetilde{\mu}^f_{>}(t)) + \int_{\tau=t} H(\mu)\, f(d\mu, d\tau).$$

The first inequality is subadditivity, which follows from convexity and degree-one homogeneity; the second is Jensen's inequality with homogeneity. Subtracting the original flow residual from the displayed limit gives precisely

$$\int_{\tau=t} H(\mu)\, f(d\mu, d\tau) + H(\widetilde{\mu}^f_{>}(t)) - H(\widetilde{\mu}^f_{\ge}(t)) \ge 0.$$

At $t = T$, all mass is counted as stopped under both laws, and the belief marginal is

identical. Therefore,

$$G_{Flow}(f_\delta)(T) - G_{Flow}(f)(T) = \chi\,\mathbb{E}[\tau_\delta - \tau], \qquad |G_{Flow}(f_\delta)(T) - G_{Flow}(f)(T)| \le 2\chi\delta \longrightarrow 0.$$

We have proved a pointwise lower-limit inequality at every time in $[2^{-K}, T]$, including its endpoints. To pass to the multiplier integral, let $M_H := \sup_{\mu\in\Delta^1} H(\mu) < \infty$. For any probability law g, homogeneity gives

$$\int_{\tau\le t} H(\mu)\,dg + H(\widetilde{\mu}^g_>(t)) \le M_H.$$

Indeed, the first term is bounded by $M_H F_g(t)$, and the second by $M_H(1 - F_g(t))$. Since the occupation term and $H(\mu_0)$ are nonnegative, $G_{Flow}(g)(t) \ge -M_H$. Fatou's lemma applied to $G_{Flow}(f_\delta) + M_H$ against the finite positive measure λ consequently yields

$$\liminf_{\delta\downarrow 0} \int_{[2^{-K},T]} G_{Flow}(f_\delta)(t)\,d\lambda(t) \ge \int_{[2^{-K},T]} G_{Flow}(f)(t)\,d\lambda(t).$$

This argument permits λ to have atoms, including at participation times and at T. Combining the objective convergence, participation-vector convergence, and flow-integral inequality gives

$$\liminf_{\delta\downarrow 0} \widetilde{\mathcal{L}}_K(f_\delta, \eta_\delta, \lambda, \gamma) \ge \widetilde{\mathcal{L}}_K(f, \eta, \lambda, \gamma).$$

For the prescribed $\varepsilon > 0$, choose δ sufficiently small and set $(\widehat{f}, \widehat{\eta}) = (f_\delta, \eta_\delta)$. This pair has all the required properties and satisfies the asserted inequality. □

A.4 Proof for Lemma 5.4

We show the optimality conditions are necessary. Let multipliers (λ, γ) solve (Dual$_K$) and contract (f, v) solve problem (RP). Without loss of generality, assume $\bar{t}_f \ge \sup_t \operatorname{supp}\,\lambda$.[17] We show directly that constraint (FOC$_v$) is necessary. Let $w \in \mathbb{V}$ be an indirect compensation rule. The Lagrangian functional $\mathcal{L}(f, \cdot, \lambda, \gamma)$ is concave. We evaluate the Gâteaux derivative of $\mathcal{L}$ at contract (f, v) in direction (f, w); denote this by $\partial\mathcal{L}(f, v, \lambda, \gamma; f, w)$. Since

[17] We elaborate on this. Let multipliers (λ, γ) solve (Dual$_K$), mechanism (f, v) solve (Primal$_K$), and suppose $\bar{t} < \sup_t \operatorname{supp}\,\lambda$. Define an alternative multiplier $\widetilde{\lambda}(t) = \lambda(t)\mathbf{1}\{t \le \bar{t}\}$. Then $\int_{t\in[2^{-K},T]} G_{Flow}(f)(t)d\lambda(t) = \int_{t\in[2^{-K},T]} G_{Flow}(f)(t)d\widetilde{\lambda}(t)$. This implies $\mathcal{L}(f, v, \lambda, v) = \mathcal{L}(f, v, \widetilde{\lambda}, v)$. Therefore, $(\widetilde{\lambda}, \gamma)$ also solves (Dual$_K$).

$\mathcal{L}$ is concave and (f, v) solves the primal problem, we obtain

$$\partial\mathcal{L}(f, v, \lambda, \gamma; f, w) = \int_0^T e^{-rt}(1 - F(t))(\Gamma(t) - \varphi'(v(t)))(w(t) - v(t))dt \leq 0$$

We consider the case when $v(t) = 1$. Let $\varepsilon > 0$. For any measurable subset $E \subseteq \{t \in [0, T] : v(t) = \overline{v}\}$, define $w(t) = v(t) - \varepsilon \cdot \mathbf{1}\{t \in E\}$. The above inequality implies

$$-\varepsilon \int_{t\in E} e^{-rt}(1 - F(t))(\Gamma(t) - \varphi'(v(t)))dt \leq 0$$

Since $\varepsilon > 0$ was arbitrary, we get $\Gamma(t) \geq \varphi'(v(t))$ almost everywhere on $\{t \in [0, T] : v(t) = \overline{v}\}$. To show $\Gamma(t) \leq \varphi'(v(t))$ almost everywhere on $\{t \in [0, T] : v(t) = 0\}$, the argument is similar; choose $w(t) = v(t) + \varepsilon \cdot \mathbf{1}\{t \in E\}$ and again use the above inequality. Lastly, we show $\Gamma(t) = \varphi'(v(t))$ whenever $v(t) \in (0, \overline{v})$. For $n \geq 2$, define

$$D_n = \left\{t \in [0, T] : \frac{1}{n} < v(t) < \overline{v} - \frac{1}{n}\right\}$$

Hence, $\{t \in [0, T] : v(t) \in (0, \overline{v})\} = \bigcup_{n\geq 2} D_n$. For a measurable subset $E \subseteq \bigcup_{n\geq 2} D_n$, define $v^*(t) = v(t) + \varepsilon \cdot \mathbf{1}\{t \in E\}$ and $v_*(t) = v(t) - \varepsilon \cdot \mathbf{1}\{t \in E\}$. We get at the same time,

$$\varepsilon \int_{t\in E} e^{-rt}(1 - F(t))(\Gamma(t) - \varphi'(v(t)))dt \leq 0$$
$$\text{and } -\varepsilon \int_{t\in E} e^{-rt}(1 - F(t))(\Gamma(t) - \varphi'(v(t)))dt \leq 0$$

This implies $\Gamma(t) = \varphi'(v(t))$ almost everywhere on $\{t \in [0, T] : v(t) \in (0, \overline{v})\}$. Thus, we get our desired optimal first-order condition.

Next, we show that there exists a belief multiplier $a \in \mathbb{R}^2$ such that (FOC$_f$) holds. Given $g \in \Delta_{\mu_0}$, define $\widetilde{\mu}^g_>(t) := \int_{\tau>t} \mu g(d\mu, d\tau)$ and $\widetilde{\mu}^g_\geq(t) := \int_{\tau\geq t} \mu g(d\mu, d\tau)$. Each vector is zero when its corresponding tail has zero mass. Define $\widehat{\ell} : \Delta^1 \to \overline{\mathbb{R}}$ as follows:[18]

$$\widehat{\ell}(\nu) = \sup\left\{\mathbb{E}_{g(\mu',\tau')}[\ell(\mu', \tau' \mid f, \lambda, \gamma)] : g \in \Delta_\nu\right\} \tag{3}$$

where we also choose bounded and measurable selections $y(\cdot) \in \partial U_A(\widetilde{\mu}^g_\geq(\cdot))$ and $z(\cdot) \in \partial H(\widetilde{\mu}^g_>(\cdot))$; such selections exist by assumption 4 (i). Since $y(\cdot), z(\cdot)$ are bounded and $\lambda \in \mathbb{L}$, we can say $\ell(\cdot \mid f, \lambda, \gamma)$ is bounded and upper semicontinuous. This implies problem (3) has a solution, and therefore, $\widehat{\ell}$ is well-defined and finite-valued. We argue

[18] We let $\overline{\mathbb{R}}$ denote the extended real line.

that $\widehat{\ell}$ is concave via a concavification argument. Define auxiliary mapping $J : \Delta^1 \to \mathbb{R}$ by $J(\mu) = \sup_{\tau \in [0,T]} \ell(\mu, \tau \mid f, \lambda, \gamma)$, and define mapping $\widetilde{\ell} : \Delta^1 \to \mathbb{R}$ by

$$\widetilde{\ell}(\nu) = \sup \{z \in \mathbb{R} : (\mu, z) \in \text{conv}(J)\}$$

where $\text{conv}(J)$ denotes the convex hull of the graph of J. We show that $\widehat{\ell}(\nu) = \widetilde{\ell}(\nu)$ for all $\nu \in \Delta^1$. Starting with proving $\widehat{\ell}(\nu) \leq \widetilde{\ell}(\nu)$, let $f \in \Delta_\mu$ be feasible and f_μ its marginal distribution over Δ^1. We have,

$$\mathbb{E}_{f(\mu',\tau')}[\ell(\mu', \tau' \mid f, \lambda, \gamma)] \leq \mathbb{E}_{f_\mu(\mu')}[J(\mu')] \leq \mathbb{E}_{f_\mu(\mu')}[\widetilde{\ell}(\mu')] \leq \widetilde{\ell}\big(\mathbb{E}_{f_\mu(\mu')}[\mu']\big) = \widetilde{\ell}(\mu)$$

where the third inequality follows from Jensen's inequality. Since f was arbitrary, we get our desired inequality. To show $\widehat{\ell}(\nu) \geq \widetilde{\ell}(\nu)$, let $\varepsilon > 0$ be your favorite smallest number. Using Carathéory's theorem, there exists finitely many points $\mu_1, \ldots, \mu_n \in \Delta^1$ and weights $\alpha_1, \ldots, \alpha_n \geq 0$ with $\sum_{i=1}^n \alpha_i = 1$ such that

$$\sum_{i=1}^n \alpha_i \mu_i = \nu \qquad \text{and} \qquad \sum_{i=1}^n \alpha_i J(\mu_i) \geq \widetilde{\ell}(\nu) - \varepsilon$$

Since ℓ is bounded and upper semincontinuous, for every μ_i, let τ_i be the time that attains $J(\mu) = \ell(\mu_i, \tau_i \mid f, \lambda, \gamma)$. Hence, the distribution $f \equiv \sum_i \alpha_i \delta_{(\mu_i, \tau_i)}$ is feasible, and its payoff is

$$\mathbb{E}_f[\ell(\mu, \tau \mid f, \lambda, \gamma)] = \sum_i \alpha_i \ell(\mu_i, \tau_i \mid f, \lambda, \gamma) = \sum_i \alpha_i J(\mu_i) \geq \widetilde{\ell}(\nu) - \varepsilon$$

Since $\varepsilon > 0$ was arbitrary, we get $\widehat{\ell}(\nu) \geq \widetilde{\ell}(\nu)$, and therefore, we have equality. Evidently, $\widetilde{\ell}$ is concave, and therefore, $\widehat{\ell}$ is concave. Whence, we can choose belief multipliers $a \in \mathbb{R}^2$ such that $a \cdot \mu$ is the supporting hyperplane of $\widehat{\ell}$ at prior belief μ_0. By the definition of supremum, $\ell(\mu, \tau \mid f, \lambda, \gamma) \leq a \cdot \mu$ for all $(\mu, \tau) \in \Delta^1 \times [0, T]$.

Next, we argue that $\int \ell(\mu, \tau \mid f, \lambda, \gamma) f(d\mu, d\tau) = a \cdot \mu_0$. Suppose for the sake of contradiction that $\int \ell(\mu, \tau \mid f, \lambda, \gamma) f(d\mu, d\tau) < a \cdot \mu_0$. This implies that there exists $\varepsilon > 0$ such that $\int \ell(\mu, \tau \mid f, \lambda, \gamma) f(d\mu, d\tau) \leq a \cdot \mu_0 - \varepsilon$. Since problem (3) can attain a solution, let $g \in \Delta_{\mu_0}$ be one. By construction, we have $\mathbb{E}_{g(\mu',\tau')}[\ell(\mu', \tau' \mid f, \lambda, \gamma)] = a \cdot \mu_0$. Let $y(\cdot) \in \partial U_A(\widetilde{\mu}_{\geq}(\cdot))$ and $z(\cdot) \in \partial H(\widetilde{\mu}_{>}(\cdot))$ be measurable selections. Then we have

$$\liminf_{\alpha \downarrow 0} \frac{\mathcal{L}(\alpha g + (1-\alpha) f, v, \lambda, \gamma) - \mathcal{L}(f, v, \lambda, \gamma)}{\alpha}$$

$$
\begin{aligned}
&= \left\{ \begin{aligned}
&\liminf_{\alpha \downarrow 0} \frac{1}{\alpha} \Bigg\{ \int \left(e^{-r\tau} \big(U_P(\mu) + \Gamma(\tau) U_A(\mu) \big) + \int_0^\tau e^{-rs} \big(\Gamma(s)(v(s) - c) - \varphi(v(s)) \big) ds \right) (\alpha(g - f))(d\mu, d\tau) \\
&\qquad - \sum_{k=0}^{K-1} e^{-rt_k} \left(U_A \left(\int_{\tau > t_k} \mu(\alpha g + (1-\alpha) f)(d\mu, d\tau) \right) - U_A \left(\int_{\tau > t_k} \mu f(d\mu, d\tau) \right) \right) \gamma_k \\
&\qquad + \int_{t \in [2^{-K}, T]} \left(\chi \int \min\{\tau, t\} (\alpha(g - f))(d\mu, d\tau) \int_{\tau \le t} H(\mu)(\alpha(g - f))(d\mu, d\tau) \right) d\lambda(t) \\
&\qquad - \int_{t \in [2^{-K}, T]} \left(H \left(\int_{\tau > t} \mu(\alpha g + (1-\alpha) f)(d\mu, d\tau) \right) - H \left(\int_{\tau > t} \mu f(d\mu, d\tau) \right) \right) d\lambda(t) \Bigg\}
\end{aligned} \right. \\
&= \left\{ \begin{aligned}
&\int \left(e^{-r\tau} \big(U_P(\mu) + \Gamma(\tau) U_A(\mu) \big) + \int_0^\tau e^{-rs} \big(\Gamma(s)(v(s) - c) - \varphi(v(s)) \big) ds \right) (g - f)(d\mu, d\tau) \\
&\qquad - \limsup_{\alpha \downarrow 0} \sum_{k=0}^{K-1} e^{-rt_k} \frac{U_A(\int_{\tau > t_k} \mu(\alpha g + (1-\alpha) f)(d\mu, d\tau)) - U_A(\int_{\tau > t_k} \mu f(d\mu, d\tau))}{\alpha} \gamma_k \\
&\qquad + \int_{t \in [2^{-K}, T]} \left(\chi \int \min\{\tau, t\} (g - f)(d\mu, d\tau) \int_{\tau \le t} H(\mu)(g - f)(d\mu, d\tau) \right) d\lambda(t) \\
&\qquad - \limsup_{\alpha \downarrow 0} \int_{t \in [2^{-K}, T]} \frac{H(\int_{\tau > t} \mu(\alpha g + (1-\alpha) f)(d\mu, d\tau)) - H(\int_{\tau > t} \mu f(d\mu, d\tau))}{\alpha} d\lambda(t)
\end{aligned} \right. \\
&\geq \left\{ \begin{aligned}
&\int \left(e^{-r\tau} \big(U_P(\mu) + \Gamma(\tau) U_A(\mu) \big) + \int_0^\tau e^{-rs} \big(\Gamma(s)(v(s) - c) - \varphi(v(s)) \big) ds \right) (g - f)(d\mu, d\tau) \\
&\qquad - \sum_{k=0}^{K-1} e^{-rt_k} \left(y(t_k) \int_{\tau > t_k} \mu(g - f)(d\mu, d\tau) \right) \gamma_k \\
&\qquad + \int_{t \in [2^{-K}, T]} \left(\chi \int \min\{\tau, t\} (g - f)(d\mu, d\tau) \int_{\tau \le t} H(\mu)(g - f)(d\mu, d\tau) \right) d\lambda(t) \\
&\qquad - \int_{t \in [2^{-K}, T]} \left(z(t) \int_{\tau > t} \mu(g - f)(d\mu, d\tau) \right) d\lambda(t)
\end{aligned} \right. \\
&= \left\{ \begin{aligned}
&\int \left(e^{-r\tau} \big(U_P(\mu) + \Gamma(\tau) U_A(\mu) \big) + \int_0^\tau e^{-rs} \big(\Gamma(s)(v(s) - c) - \varphi(v(s)) \big) ds \right) (g - f)(d\mu, d\tau) \\
&\qquad - \int \left(\sum_{k=0}^{K-1} \mathbf{1}\{\tau > t_k\} e^{-rt_k} y(t_k) \gamma_k \right) (g - f)(d\mu, d\tau) \\
&\qquad + \int_{t \in [2^{-K}, T]} \left(\chi \int \min\{\tau, t\} (g - f)(d\mu, d\tau) \int_{\tau \le t} H(\mu)(g - f)(d\mu, d\tau) \right) d\lambda(t) \\
&\qquad - \int \left(\int_{\tau \in (0, \tau)} (z(t) \cdot \mu) d\lambda(t) \right) (g - f)(d\mu, d\tau)
\end{aligned} \right. \\
&= \ell(\mu, \tau \mid f, v, \lambda, \gamma)(g - f)(d\mu, d\tau) \geq \varepsilon > \frac{\varepsilon}{2} > 0
\end{aligned}
$$

We annotate this series: the first equality is by the definition of the Lagrangian functional; the second holds by distributing the limit inferior operation; the inequality holds because $U_A(\cdot)$ and $H(\cdot)$ are convex; the third holds by an application of Fubini's theorem, and the last series hold by the definition of ℓ and the construction of g. This contradicts that f was optimal for (RP_K). Hence, we may conclude that $\int \ell(\mu, \tau \mid f, \lambda, \gamma) f(d\mu, d\tau) = a \cdot \mu_0$. This completes the proof that the optimality conditions are necessary.

A.5 Proof for Lemma 5.5

Let multipliers (λ, γ) solve (Dual$_K$) and mechanism (f, v) solve (RP$_K$). To prove that (Flow) is always binding on $[2^{-K}, \bar{t}]$, we argue that $\Lambda(\cdot)$ is strictly decreasing on $[2^{-K}, \bar{t}]$. Let $t, t' \in [0, \bar{t}]$ with $t < t'$. Since $\Lambda(\cdot)$ is a survival function, by construction, it is weakly decreasing on $[2^{-K}, T]$. Lemma A.5 states that $\Lambda(\cdot)$ is strictly decreasing on $(2^{-K}, \bar{t}) \setminus (\{t_k\}_{k=0}^K \cap (0, \bar{t}))$. Hence, there are points s, s' such that $t < s < s' < t'$ and $s, s' \notin \{t_k\}_{k=0}^K \cup \{2^{-K}, \bar{t}\}$. Therefore, we get $\Lambda(t) \geq \Lambda(s) > \Lambda(s') \geq \Lambda(t')$. Thus, $\Lambda(\cdot)$ is strictly decreasing on $[2^{-K}, \bar{t}]$.

Lemma A.5. $\Lambda(\cdot)$ is strictly decreasing on $(2^{-K}, \bar{t}) \setminus (\{t_k\}_{k=0}^K \cap (0, \bar{t}))$.

Proof. We use contradiction. Assume that $\Lambda(\cdot)$ is not strictly decreasing. This means that there exists $s_1, s_2 \in (0, \bar{t}) \setminus (\{t_k\}_{k=0}^K \cap (0, \bar{t}))$ with $s_1 < s_2$ such that $\Lambda(t) = C_0 \in \mathbb{R}_+$ for every $t \in [s_1, s_2]$. Without loss of generality, we further assume that $[s_1, s_2]$ is a maximal interval such that this is true; that is, $C_0 \geq \Lambda(t)$ for all $t \geq s_2$ and strictly so whenever $t > s_2$.

The key object of this proof is the minimal gap in equation (FOC$_f$). Formally, it is the mapping $\xi : [0, \bar{t}] \to \mathbb{R}$ defined by $\xi(t) = \max_{\mu \in \Delta^1} \{\ell(\mu, t) - a \cdot \mu\}$ where $\ell(\mu, t) \equiv \ell(\mu, t \mid f, v, \lambda, \gamma)$ and some measurable selections $y(\cdot) \in \partial U_A(\cdot)$ and $z(\cdot) \in \partial H(\cdot)$. This is a key object because it will help us get our contradiction. In step 1, we look at the behavior of ξ at exactly time s_2. We will argue that $G_{Flow}(f)(s_2) = 0$ and $\xi(s_2) = 0$. Next, in step 2, we look at the behavior of ξ at the left-hand limit s_2^-. We will derive that $\xi'(s_2^-) > 0$. Equation (FOC$_f$) implies $\xi(t) \leq 0$ for all $t \in [0, \bar{t}]$. And yet, in step 3, we will show at the right-hand limit that $\xi(s_2^+) > 0$.

Step 1. $G_{Flow}(f)(s_2) = 0$ and $\xi(s_2) = 0$.

Since $[s_1, s_2]$ is the maximal interval where Λ is constant, $G_{Flow}(f)(s_2) = 0$. Otherwise, it would violate the complementary slackness condition. Thus, constraint (Flow) binds at time s_2. We now argue that $\xi(s_2) = 0$. For the sake of contradiction, suppose $\xi(s_2) < 0$. This implies that there would exist an open neighborhood where f places no mass. Formally, there exists a large enough $N \geq 0$ such that $\int_{\Delta^1 \times (s_2 - 2^{-N}, s_2 + 2^{-N})} f(d\mu, d\tau) = 0$. For all $n \geq N$, we have

$$G_{Flow}(s_2 - 2^{-n}) - G_{Flow}(s_2) = -2^{-n} \chi \int_{\tau > s_2} \tau f(d\mu, d\tau)$$

Given that $s_2 \in (0, \bar{t})$ implies $\int_{\tau > s_2} \tau f(d\mu, d\tau) > 0$ and $G_{Flow}(f)(s_2) = 0$, we have $G_{Flow}(s_2 - 2^{-n}) < G_{Flow}(s_2) = 0$ for all $n \geq N$. This violates (Flow), and therefore, we may claim that $\xi(s_2) = 0$.

Step 2. $\xi'(s_2^-) > 0$.

Since $(t_k)_{k=0}^K$ partitions $[0,T]$, let $[t_k, t_{k+1})$ be the interval that contains time s_2. The cumulative multiplier Γ is piecewise constant on partition $(t_k)_{k=0}^K$. First-order condition (FOC$_v$) says $\Gamma(t) = \varphi'(v(t))$ for almost every t with $F(t) < 1$. Since the agent is strictly risk-averse over transfers, φ' is strictly increasing, and therefore, implies that $v \in \mathbb{V}$ is piecewise constant on partition $(t_k)_{k=0}^K$. To emphasize this property, we abuse notation and let $v_k \equiv v(\tau)$ when $\tau \in [t_k, t_{k+1})$. In the same way, let $\Gamma_k \equiv \Gamma(\tau)$ when $\tau \in [t_k, t_{k+1})$. Let $S_k = \sum_{j=0}^k (\Gamma_j(v_j - c) - \varphi(v_j))$. Define mapping $B_k : \Delta^1 \to \mathbb{R}$ by

$$B_k(\mu) = \begin{cases} \dfrac{e^{-rt_k} S_k}{r} + \displaystyle\int_0^{t_k} e^{-rs}\big(\Gamma(s)(v(s) - c) - \varphi(v(s))\big) ds \\ \qquad + \displaystyle\sum_{k=0}^{K-1} \mathbf{1}\{\tau > t_k\} e^{-rt_k} \cdot \big(y(t_k) \cdot \mu\big) \cdot \gamma_k \end{cases}$$

for $\mu \in \Delta^1$. Fix $\tau \in (t_k \vee s_1, s_2)$,[19] and let $\theta = e^{-r\tau}$. We define an auxiliary mapping $\widetilde{\xi}(\theta)$ by

$$\begin{aligned} \widetilde{\xi}(\theta) &= \max_{\mu \in \Delta^1} \{\ell(\mu, \tau) - a \cdot \mu\} \\ &= \begin{cases} \max\limits_{\mu \in \Delta^1} \Big\{ e^{-r\tau} U_P(\mu) + e^{-r\tau} \Gamma_k \, U_A(\mu) - \dfrac{e^{-rt_k}}{r} \big(\Gamma_k(v_k - c) - \varphi(v_k)\big) \\ \qquad + B_k(\mu) + \chi \displaystyle\int_0^\tau \Lambda(s) ds - \int_{s \in [0,\tau)} \big(z(s) \cdot \mu\big) d\lambda(s) - \Lambda(\tau) H(\mu) - a \cdot \mu \Big\} \end{cases} \\ &= \begin{cases} \max\limits_{\mu \in \Delta^1} \Big\{ \theta U_P(\mu) + \theta \Gamma_k \, U_A(\mu) + \chi \, C_0 \left(-\dfrac{\ln \theta}{r} - s_1 \right) - \dfrac{e^{-rt_k}}{r} \big(\Gamma_k(v_k - c) - \varphi(v_k)\big) \\ \qquad + B_k(\mu) + \chi \displaystyle\int_0^{s_1} \Lambda(s) ds - \int_{s \in [0,\tau)} \big(z(s) \cdot \mu\big) d\lambda(s) - \Lambda(\tau) \, C_0 - a \cdot \mu \Big\} \end{cases} \\ &= \begin{cases} -\dfrac{\chi \, C_0}{r} \ln \theta + \max\limits_{\mu \in \Delta^1} \Big\{ \theta U_P(\mu) + \theta \Gamma_k \, U_A(\mu) - \dfrac{e^{-rt_k}}{r} \big(\Gamma_k(v_k - c) - \varphi(v_k)\big) \\ \quad - \chi \, C_0 \, s_1 + B_k(\mu) + \chi \displaystyle\int_0^{s_1} \Lambda(s) ds - \int_{s \in [0,\tau)} \big(z(s) \cdot \mu\big) d\lambda(s) - \Lambda(\tau) C_0 - a \cdot \mu \Big\} \end{cases} \end{aligned}$$

By construction, we have $\widetilde{\xi}(\theta) = \xi(\tau)$. Moreover, $\widetilde{\xi}(\cdot)$ is strictly convex on $(0, \infty)$; in the last equality, the first additive is strictly convex, while the second additive term is convex as it is a maximum of affine mappings. Let $\theta_1 = e^{-r(t_k \vee s_1)}$ and $\theta_2 = e^{-rs_2}$. From step 1, we can say that $\widetilde{\xi}(\theta_2) = 0$. By (FOC$_f$), we have $\widetilde{\xi}'(\theta) \leq 0$ for all $\theta \in (\theta_2, \theta_1)$. The strict convexity of $\widetilde{\xi}(\cdot)$ implies the quotient $\frac{\widetilde{\xi}'(\theta) - \widetilde{\xi}'(\theta_2)}{\theta - \theta_2}$ is strictly increasing in θ' and with limit

[19] We let $x \vee y = \max\{x, y\}$ denote the join of real numbers x and y.

$\widetilde{\xi}'(\theta_2^+)$ as $\theta' \downarrow \theta_2$. Hence,

$$0 \geq \widetilde{\xi}'(\theta) = \widetilde{\xi}'(\theta_2) + \frac{\widetilde{\xi}'(\theta) - \widetilde{\xi}'(\theta_2)}{\theta - \theta_2}(\theta - \theta_2) > \widetilde{\xi}'(\theta_2^+)(\theta - \theta_2)$$

Given $\theta > \theta_2$, this implies $\widetilde{\xi}'(\theta_2^+) < 0$. In terms of the gap equation, evaluating the derivative of $\xi(\cdot)$ at $\tau = s_2^-$ gives

$$\xi'(s_2^-) = -r\theta_2\widetilde{\xi}'(\theta_2^+) = -re^{-rs_2}\widetilde{\xi}'(\theta_2^+) > 0. \tag{4}$$

Step 3. $\xi(s_2^+) > 0$.

With $\xi(s_2) = 0$ and condition (CC), there exists a belief $\mu \in \Delta^1$ such that (μ, s_2) is an element of $\operatorname{supp} f$. Let $\mu^*_{s_2} \in \Delta^1$ denote a belief maximizer that attains $\xi(s_2)$. Define mapping $\beta_k : \Delta^1 \to \mathbb{R}$ by

$$\beta_k(\mu) = U_P(\mu) + \Gamma_k U_A(\mu) - \frac{1}{r}\big(\Gamma_k\,(v_k - c) - \varphi(v_k)\big), \quad \text{for } \mu \in \Delta^1.$$

At last, we evaluate $\xi(s_2^+)$. For a large enough n, we have

$$\begin{aligned}
&\xi(s_2 + 2^{-n}) \\
&= \ell(\mu^*_{s_2}, s_2 + 2^{-n}) - a \cdot \mu^*_{s_2} - \xi(s_2) && (5)\\
&\geq \ell(\mu^*_{s_2}, s_2 + 2^{-n}) - \ell(\mu^*_{s_2}, s_2) && (6)\\
&= \begin{cases} \big(e^{-r(s_2+2^{-n})} - e^{-rs_2}\big)\beta_k(\mu^*_{s_2}) \\ \qquad + \chi \displaystyle\int_{s_2}^{s_2+2^{-n}} \Lambda(s)ds + \int_{s\in(s_2,s_2+2^{-n}]} \big(H(\mu^*_{s_2}) - z(s)\cdot\mu^*_{s_2}\big)\lambda(s)ds \end{cases} && (7)\\
&\geq \big(e^{-r(s_2+2^{-n})} - e^{-rs_2}\big)\beta_k(\mu^*_{s_2}) + \chi \int_{s_2}^{s_2+2^{-n}} \Lambda(s)ds && (8)\\
&= -re^{-rs_2}\beta_k(\mu^*_{s_2}) + \chi\, C_0\, 2^{-n} + o(2^{-n}) && (9)\\
&= \xi'(s_2^-)\, 2^{-n} + o(2^{-n}) && (10)
\end{aligned}$$

We annotate the above: equality (5) holds by the definition of ξ and that $\xi(s_2) = 0$; inequality (6) is true because $\mu^*_{s_2}$ is a maximizer at time s_2; equality (7) is simply the definition of $\ell(\mu, \tau)$ and canceling terms up to time s_2. Inequality (8) holds since H is both convex and homogeneous of degree one. We get equality (9) by first-order Taylor approximation.[20] Lastly, applying equation (4) yields equality (10).

[20] We say $f(x) = o(g(x))$ if for all $\epsilon > 0$ there exists $x_0 > 0$ such that $|f(x)| \leq \varepsilon\,|g(x)|$ whenever $x > x_0$.

Putting everything together, we have argued that $\xi'(s_2^-) > 0$. Our series of inequalities imply that $\xi(s_2^+) > 0$. However, this contradicts (FOC$_f$). Thus, Λ must be strictly decreasing on $(0, \bar{t}) \setminus (\{t_k\}_{k=0}^K \cap (0, \bar{t}))$. □

A.6 Proof for Lemma 5.6

For each $K \geq 0$, let mechanism (f_K, v_K) solve (Primal$_K$). Consider the sets $\{f_K\}_{K=1}^\infty$ and $\{v_K\}_{K=1}^\infty$. Since each f_K has compact support, the set $\{f_K\}_{K=1}^\infty$ is tight. Therefore, by Prokhorov's theorem, it is also relatively compact. Hence, there is a subsequence $\{f_{K_J}\}_{J=1}^\infty$ and limit point $f \in \Delta(\Delta^1 \times [0,T])$ such that f_{K_J} converges weakly to f. In addition, since each v_K is bounded and increasing, by Helly's selection theorem, there exists a subsequence $\{v_{K_J}\}_{J=1}^\infty$ and limit point $v \in \mathbb{V}$ such that v_{K_J} converges pointwise to v.[21] Note that $\{v_{K_J}\}_{J=1}^\infty$ is bounded in $L^p([0,T])$; so, it also weakly converges to v.[22] Each (f_{K_J}, v_{K_J}) provides an upper bound for the principal's value. Since her objective function is continuous in (f, v), by the Portmanteau theorem[23], (f, v) is also an upper bound.

We check that (f, v) is feasible for (RP). It satisfies (BP) by weak convergence. We next pass the truncated flow constraints and their binding property to the limit. Let

$$D := \{t \in [0,T] : f(\Delta^1 \times \{t\}) > 0\}, \qquad \bar{t} := \sup \operatorname{supp} f_\tau.$$

The set D is at most countable. For every $t \notin D$, weak convergence gives

$$G_{Flow}(f_{K_j})(t) \longrightarrow G_{Flow}(f)(t),$$

because the bounded integrands containing stopping indicators are continuous outside the f-null set $\Delta^1 \times \{t\}$, and the homogeneous extension of $H(\cdot)$ is continuous. At T, convergence holds regardless of a terminal atom, since

$$G_{Flow}(f_{K_j})(T) = \chi \int \tau f_{K_j}(d\mu, d\tau) + H(\mu_0) - \int H(\mu) f_{K_j}(d\mu, d\tau),$$

whose integrands are continuous and bounded. For any probability law g, the function $G_{Flow}(g)$ is right-continuous and has left limits. This follows from the corresponding properties of its cumulative integrals, continuity of H, and continuity of the occupation

[21] For reference, see exercise 13 of Rudin (1976), page 167.
[22] For reference, see theorem 12 of Royden and Fitzpatrick (1988), page 168.
[23] For reference, see theorem 2.1 of Billingsley (2013), page 16.

term. For $t > 0$, its jump satisfies

$$G_{Flow}(g)(t) - G_{Flow}(g)(t^-) = -\Big[\int_{\tau=t} H(\mu)g(d\mu, d\tau) + H(\widetilde{\mu}^g_>(t)) - H(\widetilde{\mu}^g_\geq(t))\Big] \leq 0.$$

The inequality follows from subadditivity and homogeneous Jensen's inequality for $H(\cdot)$. Now fix $t \in (0,T) \setminus D$. Since $2^{-K_j} \to 0$, eventually $2^{-K_j} < t$. The constraint (Flow_{K_j}) therefore gives $G_{Flow}(f_{K_j})(t) \geq 0$, so $G_{Flow}(f)(t) \geq 0$. Approximate each remaining $t \in [0,T)$ by nonatom times decreasing to t; right-continuity extends this inequality to those times, including 0. Direct convergence at T gives the inequality there as well. Thus f satisfies the full flow constraint on $[0,T]$. At zero, Bayes plausibility and sublinearity also imply

$$H(\mu_0) \leq \int_{\tau=0} H(\mu)f(d\mu, d\tau) + H(\widetilde{\mu}^f_>(0)).$$

Hence $G_{Flow}(f)(0) \leq 0$, and feasibility yields $G_{Flow}(f)(0) = 0$. To pass binding to the limit, fix $t < \bar{t}$. The relatively open set $\Delta^1 \times (t,T]$ has positive f-mass. Portmanteau gives

$$\liminf_{j\to\infty} f_{K_j}(\Delta^1 \times (t,T]) \geq f(\Delta^1 \times (t,T]) > 0.$$

Consequently, $\bar{t}_{K_j} > t$ eventually. If also $t \in (0,\bar{t}) \setminus D$, then eventually $2^{-K_j} < t < \bar{t}_{K_j}$. The truncated binding-flow property gives $G_{Flow}(f_{K_j})(t) = 0$, and passing to the limit gives $G_{Flow}(f)(t) = 0$. Right-continuity extends this equality to every $t \in [0,\bar{t})$. If $\bar{t} > 0$, its left limit at $\bar{t}$ is zero; the absence of upward jumps and feasibility imply

$$0 \leq G_{Flow}(f)(\bar{t}) \leq G_{Flow}(f)(\bar{t}^-) = 0.$$

This covers a stopping atom at $\bar{t}$, including $\bar{t} = T$. If $\bar{t} = 0$, binding was already established at zero. Thus, the full flow constraint binds throughout $[0,\bar{t}]$.

Let $G_{DP}(f,v)(t)$ denote the agent's continuation value from time t. Formally,

$$\widehat{G}(\mu,\tau) = e^{-r\tau}U_A(\mu) + \int_t^\tau e^{-rs}\big(v(s) - c\big)ds$$
$$G_{DP}(f,v)(t) = \int_{\tau>t} \widehat{G}(\mu,\tau)f(d\mu,d\tau) - e^{-rt}U_A\left(\int_{\tau>t} \mu f(d\mu,d\tau)\right)$$

We argue that (f,v) satisfies constraint (RDP). Let $\varepsilon > 0$ be your favorite smallest number.

Given $t \in [0,T)$ and $0 < \varepsilon \leq T - t$, define auxiliary mapping $Q_{t,\varepsilon} : [0,T] \to [0,1]$ by

$$Q_{t,\varepsilon}(\tau) = \min\left\{T, \frac{(\tau - t)_+}{\varepsilon}\right\}$$

Notice that the integrand $\widehat{G}(\mu,\tau)$ is continuous and bounded[24]; hence, the product $\tau \mapsto Q_{t,\varepsilon}(\tau) \cdot \widehat{G}(\mu,\tau)$ is also continuous and bounded in τ. As $\varepsilon \downarrow 0$, the mapping $Q_{t,\varepsilon}(\cdot)$ converges pointwise to the indicator function $\mathbf{1}\{\tau > t\}$. Finally,

$$\begin{aligned}
&G_{DP}(f,v)(t) \\
&= \lim_{\varepsilon \downarrow 0}\left\{\int Q_{t,\varepsilon}(\tau)\widehat{G}(\mu,\tau)f(d\mu,d\tau) - e^{-rt}U_A\left(\int Q_{t,\varepsilon}(\tau)\mu f(d\mu,d\tau)\right)\right\} \\
&\geq \lim_{\varepsilon \downarrow 0}\left\{\limsup_{J\to\infty}\left\{\int Q_{t,\varepsilon}(\tau)\widehat{G}(\mu,\tau)f_{K_J}(d\mu,d\tau) - e^{-rt}U_A\left(\int Q_{t,\varepsilon}(\tau)\mu f_{K_J}(d\mu,d\tau)\right)\right\}\right\} \\
&\geq 0
\end{aligned}$$

The first equality follows from the dominated convergence theorem. The first inequality follows from the Portmanteau theorem. The second inequality follows from the fact that each (f_{K_J}, v_{K_J}) satisfies constraint (RDP). This shows that (f,v) is feasible for (RP).

A.7 Proof for Lemma 5.7

To prove the lemma, we verify that $\langle \mu_t^f \rangle$ is a martingale and that $(\langle \mu_t^f \rangle, \tau^f, v)$ satisfies (DPC) and (Var). We start by showing that $\langle \mu_t^f \rangle$ is a martingale. Equip the construction with the filtration $\mathcal{F}_t := \sigma(\mathbf{1}\{\tau^f \leq s\}, \mu_{\tau^f}^f \mathbf{1}\{\tau^f \leq s\} : 0 \leq s \leq t)$. On $\{\tau^f \leq t\}$, this filtration reveals the terminal belief $\mu_{\tau^f}^f$. On $\{\tau^f > t\}$, all its generators are zero, so it reveals only survival beyond t. Since the joint law of $(\mu_{\tau^f}^f, \tau^f)$ is f, it follows that

$$\mathbb{E}_{\mathcal{P}}[\mu_{\tau^f}^f \mid \mathcal{F}_t] = \mu_{\tau^f}^f \mathbf{1}\{\tau^f \leq t\} + \widehat{\mu}(t)\mathbf{1}\{\tau^f > t\} = \mu_t^f \qquad \mathcal{P}\text{-almost surely.}$$

If $\mathcal{P}(\tau^f > t) = 0$, the continuation event is null and its assigned belief is immaterial. The process is adapted and bounded. For $0 \leq t \leq t' \leq T$, the tower property therefore gives

$$\mathbb{E}_{\mathcal{P}}[\mu_{t'}^f \mid \mathcal{F}_t] = \mathbb{E}_{\mathcal{P}}\Big[\mathbb{E}_{\mathcal{P}}[\mu_{\tau^f}^f \mid \mathcal{F}_{t'}] \,\Big|\, \mathcal{F}_t\Big] = \mathbb{E}_{\mathcal{P}}[\mu_{\tau^f}^f \mid \mathcal{F}_t] = \mu_t^f.$$

[24]This is justified since $U_A(\cdot)$ is continuous and bounded, and $v \in \mathbb{V}$.

Hence, $\langle \mu_t^f \rangle$ is a martingale in $\langle \mathcal{F}_t \rangle$, and also in its natural filtration, which is contained in $\langle \mathcal{F}_t \rangle$. Finally, the process starts at μ_0 almost surely. Binding flow at zero gives

$$\int_{\tau=0} H(\mu) f(d\mu, d\tau) + H(\widetilde{\mu}_>(0)) = H(\mu_0).$$

Together with Bayes plausibility, this is equality in Jensen's inequality for the beliefs stopping at zero and, when it has positive probability, the surviving conditional mean. Strict convexity of H on Δ^1 therefore forces $\mu = \mu_0$ almost surely on $\{\tau = 0\}$ and $\widehat{\mu}(0) = \mu_0$ whenever $\mathcal{P}(\tau^f > 0) > 0$. If all mass stops at zero, the same Jensen equality forces every stopping belief to equal μ_0. In either case, $\mu_0^f = \mu_0$ almost surely.

Next we show our candidate satisfies (DPC). We have for every time $t \in [0, T]$,

$$\begin{aligned}
&\mathcal{P}(\tau^f > t) \cdot \mathbb{E}_{\mathcal{P}}\left[e^{-r(\tau^f - t)} U_A(\mu^f_{\tau^f}) + \int_t^{\tau^f} e^{-r(s-t)} \big(v(s) - c\big) ds \,\middle|\, \mathcal{F}_t \right] \\
&= \mathcal{P}(\tau^f > t) \cdot \mathbb{E}_{\mathcal{P}}\left[e^{-r(\tau^f - t)} U_A(\mu^f_{\tau^f}) + \int_t^{\tau^f} e^{-r(s'-t)} \big(v(s') - c\big) ds' \,\middle|\, \tau^f > t \right] \\
&= \int_{\tau > t} \left(e^{-r(\tau - t)} U_A(\mu) + \int_t^{\tau} e^{-r(s-t)} \big(v(s) - c\big) ds \right) f(d\mu, d\tau) \geq U_A(\widetilde{\mu}_>(t))
\end{aligned}$$

Lastly, we verify that $\langle \mu_t^f \rangle$ satisfies (Var). For all $t, t' \in [0, T]$ with $t < t'$, we have

$$\begin{aligned}
&\mathcal{P}(\tau^f > t) \cdot \mathbb{E}_{\mathcal{P}}[H(\mu_{t'}) - H(\mu_t) \mid \mathcal{F}_t] \\
&= \left(\int_{\tau > t} f(d\mu, d\tau) \right) \cdot \mathbb{E}_f\big[H(\mu_{t'}) - H(\mu_t) \mid \mathcal{F}_t\big] \\
&= \left(\int_{\tau > t} f(d\mu, d\tau) \right) \cdot \mathbb{E}_f\big[H(\mu_{t'}) - H(\mu_t) \mid \tau > t\big] \\
&= \int_{t < \tau \leq t'} H(\mu) f(d\mu, d\tau) + H\left(\int_{\tau > t'} \mu f(d\mu, d\tau) \right) - H\left(\int_{\tau > t} \mu f(d\mu, d\tau) \right) \\
&= \chi \int_{\tau > t} \min\{\tau - t, t' - t\} f(d\mu, d\tau) \\
&\leq \chi \int_{\tau > t} (t' - t) f(d\mu, d\tau) \\
&= \mathcal{P}(\tau^f > t) \cdot \chi \cdot (t' - t)
\end{aligned}$$

Thus, $\langle \mu_t^f \rangle$ satisfies (Var). This completes the proof.

B Proofs

B.1 Proof for Corollary 3.2

Suppose the principal only faces an ex-ante participation constraint. With this, we only need to consider the partition $\{0, T\}$. Let $(f, v) \in \Delta_{\mu_0} \times \mathbb{V}$ solve problem (RP_K) and (λ, γ_0) solve (Dual_K) where $\gamma_0 \in \mathbb{R}_+$; γ_0 represents the multiplier for the agent's ex-ante participation constraint. By lemma 5.4, the indirect compensation rule $v(\cdot)$ must satisfy

$$\begin{cases} \gamma_0 \leq \varphi'(v(t)) & \text{if } v(t) = 0 \\ \gamma_0 = \varphi'(v(t)) & \text{if } v(t) \in (0, \overline{v}) \ , \quad \text{a.e. } t \text{ with } F_\tau(t) < 1 \\ \gamma_0 \geq \varphi'(v(t)) & \text{if } v(t) = \overline{v} \end{cases}$$

Evidently, the indirect compensation rule must be constant almost surely for all $t \in [0, T]$ with $F_\tau(t) < 1$. Therefore, the transfer rule $\psi(\cdot)$ is constant almost surely.

B.2 Proof for Proposition 4.1

Let $B_\tau \subseteq \{t \in [0, T] : F_\tau(t) < 1\}$ be a Borel measurable subset and $z(\cdot) \in \partial H(\widehat{\mu}(\cdot))$ a measurable selection. Recall mapping $G_{Flow}(f)(t)$ defined in Subsection 5.2. By lemma 5.5, $G_{Flow}(f)(t)$ is always binding on $[0, T]$.

First we consider the easier case when f_τ has no atoms on B_τ. For all $t \in B_\tau$,

$$\begin{aligned} &\frac{d}{ds}\Big[G_{Flow}(f)(s)\Big]_{s=t} = 0 \\ &\implies \chi \cdot (1 - F_\tau(t))d\tau = \bigg(\int H(\mu) f_{\mu|\tau}(d\mu \mid t) - z(t) \cdot \int \mu f_{\mu|\tau}(d\mu \mid t) \bigg) f_\tau(d\tau) \\ &\implies \chi \cdot (1 - F_\tau(t))d\tau = \int \Big(H(\mu) - H(\widehat{\mu}(t)) - z(t) \cdot (\mu - \widehat{\mu}(t)) \Big) f_{\mu|\tau}(d\mu \mid t) f_\tau(d\tau) \end{aligned}$$

This is our desired result. Now suppose at time t, there is an atom of mass $m > 0$. This implies

$$m \cdot \int H(\mu) f_{\mu|\tau}(d\mu \mid t) + H\bigg(\int_{\tau > t} \mu f(d\mu, d\tau) \bigg) = H\bigg(m \cdot \int \mu f_{\mu|\tau}(d\mu \mid t) + \int_{\tau > t} \mu f(d\mu, d\tau) \bigg)$$

Convexity and homogeneity of degree one yields

$$H\bigg(m \cdot \int \mu f_{\mu|\tau}(d\mu \mid t) + \int_{\tau > t} \mu f(d\mu, d\tau) \bigg)$$

$$\leq H\bigg(\int_{\tau>t} \mu f(d\mu, d\tau) + m \cdot H\bigg(\int \mu f_{\mu|\tau}(d\mu \mid t)\bigg)\bigg)$$
$$\leq H\bigg(\int_{\tau>t} \mu f(d\mu, d\tau)\bigg) + m \cdot \int H(\mu) f_{\mu|\tau}(d\mu \mid t)$$

Since H is strictly convex, this implies $\int \mu f_{\mu|\tau}(d\mu \mid t) = \widehat{\mu}(t)$, and therefore, $\mu = \widehat{\mu}(t)$ $f_{\mu|\tau}(\cdot \mid t)$-a.s. The right-hand side of (1) equals to zero, and therefore, we still get our desired equality.

B.3 Proof for Proposition 4.2

We prove Proposition 4.2 with two lemmas: lemma B.1 provides an upper bound for problem (FP) and lemma B.2 gives an attainment argument. Throughout this section, it is understood that both lemmas impose Assumption 6.

Lemma B.1. $V_{FP}(\mu_0) \leq V_{SP}(\mu_0)$.

Proof. Fix a feasible full contract $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$. Since U_P is bounded on Δ^1, a contract for which $\mathbb{E}_{\mathcal{P}}\left[\int_0^\tau \psi_t \, dt\right] = +\infty$ gives the principal payoff $-\infty$ and satisfies the desired upper bound trivially. Hence, suppose $\mathbb{E}_{\mathcal{P}}\left[\int_0^\tau \psi_t \, dt\right] < +\infty$. Concavity of u and $u(0) = 0$ imply $u(x) \leq u'(0) \cdot x$. Therefore, $\mathbb{E}_{\mathcal{P}}\left[\int_0^\tau u(\psi_t) \, dt\right] \leq u'(0) \cdot \mathbb{E}_{\mathcal{P}}\left[\int_0^\tau \psi_t \, dt\right] < +\infty$. Since $\{\tau = 0\} \in \mathcal{F}_0$, $\mu_\tau = \mu_0$ and all time integrals vanish on this event, we may extend the time-zero dynamic participation constraint there by equality. Taking expectations gives

$$\mathbb{E}_{\mathcal{P}}[U_A(\mu_\tau)] + \mathbb{E}_{\mathcal{P}}\left[\int_0^\tau u(\psi_t) \, dt\right] - c \cdot \mathbb{E}_{\mathcal{P}}[\tau] \geq \mathbb{E}_{\mathcal{P}}\left[U_A\left(\mathbb{E}_{\mathcal{P}}[\mu_\tau \mid \mathcal{F}_0]\right)\right].$$

Since U_A is bounded, this inequality cannot hold if $\mathbb{E}_{\mathcal{P}}[\tau] = +\infty$. Thus, $\mathbb{E}_{\mathcal{P}}[\tau] < +\infty$, and in particular $\tau < +\infty$ almost surely.

Let π be the law of μ_τ. Apply optional sampling first at $\tau \wedge n$. Since $\langle \mu_t \rangle$ is a bounded martingale, dominated convergence gives

$$\mathbb{E}_\pi[\mu] = \mathbb{E}_{\mathcal{P}}[\mu_\tau] = \mu_0. \tag{11}$$

Hence, $\pi \in \Pi_{\mu_0}$. Moreover, the variation constraint makes $H(\mu_t) - \chi t$ a supermartingale. Optional sampling at $\tau \wedge n$, followed by bounded convergence for $H(\mu_{\tau \wedge n})$ and monotone convergence for $\tau \wedge n$, yields

$$\mathbb{E}_\pi[H(\mu) - H(\mu_0)] \leq \chi \cdot \mathbb{E}_{\mathcal{P}}[\tau]. \tag{12}$$

Since U_A and H are convex, equation (11) and Jensen's inequality imply

$$\mathbb{E}_\pi[U_A(\mu) - U_A(\mu_0)] \geq 0 \quad \text{and} \quad \mathbb{E}_\pi[H(\mu) - H(\mu_0)] \geq 0. \tag{13}$$

Optional sampling also gives $\mathbb{E}_\mathcal{P}[\mu_\tau \mid \mathcal{F}_0] = \mu_0$. Taking expectations in the time-zero dynamic participation constraint gives

$$\mathbb{E}_\pi[U_A(\mu) - U_A(\mu_0)] + \mathbb{E}_\mathcal{P}\left[\int_0^\tau u(\psi_t)\,dt\right] - c \cdot \mathbb{E}_\mathcal{P}[\tau] \geq 0. \tag{14}$$

If $\mathbb{E}_\mathcal{P}[\tau] = 0$, then $\tau = 0$ almost surely and equations (12)–(13) imply $\mathbb{E}_\pi[H(\mu) - H(\mu_0)] = 0$. In fact, $\mu_\tau = \mu_0$ almost surely, so $\pi = \delta_{\mu_0}$. The static pair $(\pi, 0)$ is feasible and has the same payoff as the full contract, which proves the desired comparison in this case. Suppose henceforth that $\mathbb{E}_\mathcal{P}[\tau] > 0$.

Define the average transfer utility by

$$q \equiv \frac{\mathbb{E}_\mathcal{P}\left[\int_0^\tau u(\psi_t)\,dt\right]}{\mathbb{E}_\mathcal{P}[\tau]}.$$

Since $0 \leq \psi_t \leq \overline{\psi}$, we have $q \in [0, \overline{v}]$, and hence $\rho \equiv \varphi(q) \in [0, \overline{\psi}]$. Convexity of φ and Jensen's inequality under the normalized measure $\mathbf{1}\{t < \tau(\omega)\}d\mathcal{P}(\omega)dt/\mathbb{E}_\mathcal{P}[\tau]$ give

$$\begin{aligned}\frac{\mathbb{E}_\mathcal{P}\left[\int_0^\tau \psi_t\,dt\right]}{\mathbb{E}_\mathcal{P}[\tau]} &= \frac{\mathbb{E}_\mathcal{P}\left[\int_0^\tau \varphi(u(\psi_t))\,dt\right]}{\mathbb{E}_\mathcal{P}[\tau]} \\ &\geq \varphi\left(\frac{\mathbb{E}_\mathcal{P}\left[\int_0^\tau u(\psi_t)\,dt\right]}{\mathbb{E}_\mathcal{P}[\tau]}\right) = \varphi(q) = \rho.\end{aligned} \tag{15}$$

We claim that (π, ρ) is feasible for static problem (SP). Its participation expression is

$$\mathbb{E}_\pi[U_A(\mu) - U_A(\mu_0)] + \frac{q-c}{\chi}\mathbb{E}_\pi[H(\mu) - H(\mu_0)].$$

If $q \geq c$, this expression is nonnegative by equation (13). If $q < c$, equations (12) and (14) imply

$$\begin{aligned}&\mathbb{E}_\pi[U_A(\mu) - U_A(\mu_0)] + \frac{q-c}{\chi}\mathbb{E}_\pi[H(\mu) - H(\mu_0)] \\ &\quad \geq \mathbb{E}_\pi[U_A(\mu) - U_A(\mu_0)] + (q-c)\mathbb{E}_\mathcal{P}[\tau] \\ &\quad = \mathbb{E}_\pi[U_A(\mu) - U_A(\mu_0)] + \mathbb{E}_\mathcal{P}\left[\int_0^\tau u(\psi_t)\,dt\right] - c\mathbb{E}_\mathcal{P}[\tau]\end{aligned}$$

$$\geq 0.$$

Thus, the static participation constraint holds in either case.

Finally, equations (12) and (15), together with $\rho \geq 0$, imply

$$\begin{aligned}
\mathbb{E}_{\mathcal{P}}\left[U_P(\mu_\tau) - \int_0^\tau \psi_t\, dt\right] &= \mathbb{E}_\pi[U_P(\mu)] - \mathbb{E}_{\mathcal{P}}\left[\int_0^\tau \psi_t\, dt\right] \\
&\leq \mathbb{E}_\pi[U_P(\mu)] - \rho \cdot \mathbb{E}_{\mathcal{P}}[\tau] \\
&\leq \mathbb{E}_\pi[U_P(\mu)] - \frac{\rho}{\chi}\mathbb{E}_\pi[H(\mu) - H(\mu_0)] \\
&\leq V_{SP}(\mu_0).
\end{aligned}$$

Taking the supremum over all feasible full contracts proves the result. □

Lemma B.2. Let $(\pi, \rho) \in \Pi_{\mu_0} \times [0, \overline{\psi}]$ solve static problem (SP) where π is nondegenerate. Define $\alpha = \frac{\chi}{\mathbb{E}_\pi[H(\mu)] - H(\mu_0)}$. Let $M \sim \pi$ and define stopping rule $\tau \sim \text{Exp}(\alpha)$ to be independent of M. Define exploring rule $\langle \mu_t \rangle$ by

$$\mu_t = \mu_0 + \mathbf{1}\{\tau \leq t\} \cdot (M - \mu_0)$$

Choose $\rho_0 \in [0, \overline{\psi}] \setminus \{\rho\}$. Define transfer rule $\langle \psi_t \rangle$ by

$$\psi_t = \begin{cases} \rho & \text{if } t < \tau \\ \rho_0 & \text{if } t \geq \tau \end{cases}$$

We consider the natural filtration generated by $\langle \mu_t \rangle, \langle \psi_t \rangle$. Then $(\langle \mu_t \rangle, \langle \psi_t \rangle, \tau)$ is feasible and optimal for problem (FP).

Proof. Since π is nondegenerate and Bayes plausible, strict convexity of H and strict Jensen's inequality imply $\mathbb{E}_\pi[H(\mu) - H(\mu_0)] > 0$. Thus, α is positive and finite. Let $\mathcal{F}_t = \sigma(\mu_s, \psi_s : 0 \leq s \leq t)$ denote the natural filtration, or its usual augmentation. Because $\rho_0 \neq \rho$, $\{\tau \leq t\} = \{\psi_t = \rho_0\} \in \mathcal{F}_t$. Hence, τ is a stopping time, including when π has an atom at μ_0. Both processes are bounded, adapted, and càdlàg. In particular, $\langle \psi_t \rangle$ is progressive and satisfies limited liability and the upper transfer bound.

Fix $0 \leq s < t$. On $\{\tau \leq s\}$, both μ_s and μ_t equal M. On $\{\tau > s\}$, the observed history reveals no information about M. Independence, exponential memorylessness, and $\mathbb{E}_{\mathcal{P}}[M] = \mu_0$ therefore give

$$\mathbb{E}_{\mathcal{P}}[\mu_t \mid \mathcal{F}_s] = \mathbf{1}\{\tau \leq s\} M$$

$$+ \mathbf{1}\{\tau > s\} \left(e^{-\alpha(t-s)}\mu_0 + (1 - e^{-\alpha(t-s)})\mathbb{E}_{\mathcal{P}}[M]\right) = \mu_s.$$

Thus, $\langle \mu_t \rangle$ is a martingale. The same argument yields

$$\begin{aligned}\mathbb{E}_{\mathcal{P}}[H(\mu_t) - H(\mu_s) \mid \mathcal{F}_s] &= \mathbf{1}\{\tau > s\}(1 - e^{-\alpha(t-s)})\mathbb{E}_{\pi}[H(\mu) - H(\mu_0)] \\ &\leq \alpha(t-s)\mathbb{E}_{\pi}[H(\mu) - H(\mu_0)] \\ &= \chi(t-s),\end{aligned}$$

where the inequality uses $1 - e^{-x} \leq x$. Hence, the variation constraint holds.

It remains to verify dynamic participation. At every active history $\{\tau > t\}$, independence and memorylessness imply

$$\mathbb{E}_{\mathcal{P}}[M \mid \mathcal{F}_t] = \mu_0 \quad \text{and} \quad \mathbb{E}_{\mathcal{P}}[\tau - t \mid \mathcal{F}_t] = \frac{1}{\alpha}.$$

Consequently, the left-hand side of constraint (DPC) minus its right-hand side equals

$$\begin{aligned}&\mathbb{E}_{\pi}[U_A(\mu) - U_A(\mu_0)] + \frac{u(\rho) - c}{\alpha} \\ &\quad = \mathbb{E}_{\pi}[U_A(\mu) - U_A(\mu_0)] + \frac{u(\rho) - c}{\chi}\mathbb{E}_{\pi}[H(\mu) - H(\mu_0)] \geq 0,\end{aligned}$$

where the inequality is precisely the static participation constraint. Thus, the induced contract is feasible.

Finally, $\mathbb{E}[\tau] = \frac{1}{\alpha} = \frac{\mathbb{E}_{\pi}[H(\mu) - H(\mu_0)]}{\chi}$, and the value of ρ_0 at and after τ does not affect the Lebesgue integral. Therefore, the principal's payoff is

$$\begin{aligned}\mathbb{E}_{\mathcal{P}}\left[U_P(\mu_\tau) - \int_0^\tau \psi_t \, dt\right] &= \mathbb{E}_{\pi}[U_P(\mu)] - \rho \cdot \mathbb{E}[\tau] \\ &= \mathbb{E}_{\pi}[U_P(\mu)] - \frac{\rho}{\chi} \cdot \mathbb{E}_{\pi}[H(\mu) - H(\mu_0)] = V_{SP}(\mu_0).\end{aligned}$$

Lemma B.1 shows that no feasible full contract can give the principal more than $V_{SP}(\mu_0)$. Hence, the induced contract is optimal for problem (FP). □