Optimal Discovery and Influence Through Selective Sampling*

Arjada Bardhi†

April 15, 2018

Abstract

Most decisions– from a job seeker appraising a job offer to a policymaker assessing a novel social program– involve the consideration of numerous attributes of an object of interest. This paper studies the optimal evaluation of a complex project of uncertain quality by sampling a limited number of its attributes. The project is described by a unit mass of correlated attributes, of which only one is observed initially. Optimal sampling and adoption is characterized for both single-agent and principal-agent evaluation. In the former, sampling is guided by the initial attribute but it is unaffected by its realization. Sequential and simultaneous sampling are equivalent. The optimal sample balances variability of sampled attributes with the importance of neighboring unsampled ones. Under principal-agent evaluation, the realization of the initial attribute informs sampling so as to better influence adoption. Sampling hinges on (i) its informativeness for the principal, and (ii) the variation of the agent's posterior belief explained by the principal's posterior belief. Optimal sampling is not necessarily a compromise between the players' ideal samples. I identify conditions under which mild disagreement leads to excessively risky or conservative sampling. Yet, drastic disagreement always induces compromise.

JEL: *D83, D81, D72, D04*

Keywords: *Multi-attribute object, evaluation, optimal sampling, adoption decision, Brownian motion, sampling capacity, influence.*

*I am deeply grateful to Bruno Strulovici, Asher Wolinsky, Georgy Egorov, and Yingni Guo for their insightful guidance and support. I have benefited from comments and suggestions by Nemanja Antic, Aislinn Bohren, Gent Carrabregu, Gabriel Carroll, Eddie Dekel, Jeff Ely, Peter Klibanoff, Meg Meyer, Natalya Naumenko, Wojciech Olszewski, Alessandro Pavan, Debraj Ray, Eran Shmaya, Marciano Sinischalci, Bela Szabadi, Quitze Valenzuela-Stookey, and Rajkamal Vasu, as well as from various seminar audiences at Northwestern and Oxford. Any remaining errors are mine.

†Department of Economics, Northwestern University, abardhi@u.northwestern.edu.

Contents

1 Introduction

Whether one evaluates a job candidate or a job offer, contemplates a purchase, or appraises the benefits and shortcomings of a public policy –each of these decision situations involves many considerations. Understanding how one chooses to explore some considerations rather than others prior to making the decision is of fundamental importance in economics. An employer evaluating a prospective employee cares about her proficiency across a wide range of skills – from soft skills such as self-confidence, job ethic, and effective communication to hard skills such as computer programming, proficiency in a foreign language, and data processing. Given time constraints, however, the employer verifies only a few skills before the hiring decision. A policymaker considering the large-scale implementation of a social program weighs the potential impact of the policy on all communities. However she learns about the effect on some particular communities through pilot projects or focus groups. An investigative journalist analyzing a long piece of written evidence is certainly interested in the relevance and authenticity of the entire document before including it in a published report. Yet she has to decide which aspects or passages of the document to focus on.[1] An instructor contemplating whether to adopt a new textbook selectively skims through its treatment of some of the major topics of the subject before deciding.

All of these examples share some key features. A decisionmaker considers the adoption of a multi-attribute object (an employee, a policy, an evidence source, or a textbook) with correlated attributes. Due to her inability to verify all relevant attributes, she perfectly samples a subset of these attributes before making an adoption decision. This is *the optimal evaluation problem* subject to a capacity constraint on attribute discovery. It consists of (i) optimal sampling of attributes, and (ii) optimal adoption based on the sampled evidence.

This paper studies the evaluation of a single multi-attribute project in two natural scenarios: when evaluation is performed by a single agent, and when it is shared among a principal and an agent who disagree on the relative importance of the attributes. In the single-agent problem, we seek to understand what criteria guide optimal sampling and the subsequent adoption decision. How does the capacity constraint, prior knowledge about some attributes, and the relative importance of attributes affect optimal sampling? What shapes the agent's preference for depth versus breadth of sampling?

The second part of the paper studies the nature of the distortions that arise in principal-agent evaluation. How does prior disagreement on the expected value of the project guide such sampling? Is the sampled evidence always a compromise, in the sense that its informativeness for one player cannot be improved without hurting its informativeness for the other? Is the sampling size ever purposely restricted?

Examples of shared evaluation between players with different attribute preferences are abundant. Organizations routinely rely on evaluation units to produce evidence about day-to-day decisions. Similarly, U.S. Congress relies on federal agencies to produce evidence on

[1] A similar observation can be made regarding an investigative journalist's choice of databases. The journalist cares about the potential evidence in all databases, but she has to decide which databases to focus on.

the effectiveness of potential federal social programs.[2] Such programs are expected to have heterogeneous impact on different communities. The differences in how U.S. Congress and the federal agency weigh the affected communities influences which communities are explored prior to the large-scale implementation of the program. Similarly, in the textbook choice example, the review of the textbook might be conducted by a teaching assistant, while the instructor decides whether to adopt it based on the sampled material. The teaching assistant might want to learn about textbook's treatment of topics that are more distant from her research area (which require greater effort by her), whereas the instructor might be interested on topics that will be revisited in students' future coursework.[3]

The first contribution of the paper consists in identifying a novel and tractable way of modeling correlated attributes in the evaluation problem. We employ methods similar to those used by **?** and **?** in models of search and strategic experimentation: the realizations of a unit mass of attributes are assumed to follow an unknown Brownian path. The players know perfectly the realization of a single attribute, and have limited sampling opportunities. That is, the policymaker and the evaluation unit witness the effect of the policy under evaluation in one community, and can afford to expand it in a limited number of other communities in the form of small-scale pilot studies. The evidence generated by the pilot studies informs the policymaker's decision of whether to launch the program at large scale. The chosen correlation structure captures some natural features. Communities are ranked along a policy-relevant dimension, for instance the median household income. Two communities of comparable income are expected to witness similar outcomes from the policy. Inferences for an untreated community are drawn from treated communities that are most comparable to it. The more comparable an untreated community is to a treated one, the less uncertain is its outcome. The magnitude of this uncertainty is independent of observed outcomes. Hence in a nutshell, inference is local and based on linear interpolation for untreated communities.

This framework, and this is in turn the second contribution of the paper, enables a complete characterization of optimal single-agent evaluation for both simultaneous and sequential attribute sampling (section 3). We make precise the intuition that optimal sampling balances the variability of sampled attributes with how informative these attributes are about unsampled neighboring attributes. A policymaker who knows the outcome of the policy on a community of income x_0 faces a sampling tradeoff. On the one hand, she prefers to sample (or treat) communities with income x_1 very different from x_0, as the outcome in such communities is highly uncertain. On the other hand, we show that the greater the difference

[2] Starting with Weiss (1973), a growing literature on evaluation studies explores the political motivations behind the evaluation of social and educational programs. For instance, Lipsky et al. (2007) study the interaction between sponsors of ADR (Alternative Dispute Resolution) workplace programs and academic evaluators. Chelimsky (1995) aptly pointed out that "a theory of evaluation has to be as much a theory of political interaction as a theory of how to determine facts or how knowledge is constructed."

[3] In a similar vein, Amazon offers the *Look Inside the Book* program, which allows publishers to make a few pages of the book vieweable by prospective readers. That is, the publisher samples some pages, perfectly reveals them to the reader, and the reader decides whether to buy the book. In contrast to examples modeled in this paper, the publisher is arguably (i) primarily interested in selling the book, rather than discovering its quality, and (ii) privately informed about the quality of the book.

$|x_1 - x_0|$, the less informative the treated community is about other untreated communities. This tradeoff induces the policymaker to never treat the richest and the poorest community. Moreover, a larger sampling capacity leads to more dispersed exploration: if the policymaker can sample more communities, the poorest sampled community is poorer and the richest sampled one is richer than with fewer sampled communities (proposition 3.4). We also provide a number of other monotone comparative statics results with respect to the sampling capacity, the initially known attribute, and the weight function over attributes (subsection 3.3).

The single-agent analysis further establishes that optimal sampling of attributes is informed by the initially known attribute, but not by its particular observed realization. This realization informs only the adoption decision; the informativeness of a sample for the agent's decision does not depend on the first realization. That is, whether the program was successful in the initial community is immaterial to which communities receive the pilot project: there is an objective way of distributing them across communities despite the initial success of the program (proposition 3.2). Moreover, this insight is useful for establishing the equivalence of optimal simultaneous sampling and optimal sequential sampling in single-agent evaluation. The optimal samples coincide under these two modes of sampling (proposition 3.5). The policymaker does not benefit from the flexibility of rolling out pilot projects sequentially. This equivalence is due to the irrelevance of the expected value of the project to the optimal sample, regardless of the number of attributes already sampled. That is, having sampled k attributes, the current expected value for the project does not inform the choice of the remaining $n-k$ attributes.

The third contribution of the paper lies in the analysis of shared evaluation between a principal (endowed with adoption authority) and an agent (endowed with sampling authority) in section 4. We rule out the presence of any informational asymmetries by assuming that at the start of the game both players observe only the initial attribute and its realization. Because the players differ in the importance they attach to different attributes, their initial expected values for the project might differ.[4] Ex-ante disagreement takes two forms: it is *mild* if the players agree on whether the project is initially promising, and it is *drastic* otherwise. The agent uses this disagreement to decide on the optimal sample: accordingly, the initial realization, which is irrelevant for sampling in the single-player problems, gains persuasive value. Because the sample choice now depends on the expected values held by the two players, the equivalence between simultaneous and sequential sampling breaks down.

A key observation regarding the agent's problem, established in proposition 4.1, is that each feasible sample is described by two sufficient statistics: the informativeness of the sample for the principal, and the variation of the agent's posterior belief explained by the principal's posterior belief. The latter captures how well the principal's decision reflects the agent's interests given the acquired sample. This explained variation consists of the informativeness of the sample for the agent, weighted by the correlation that the sample induces on the two

[4]In the case of a driftless Brownian process, the two expected values are the same.

players' posterior expected values. *Ceteris paribus*, the agent prefers higher explained variation, because that makes him more optimistic after conditioning on the principal's adoption. We show that central to the agent's problem is the comparison of (i) the principal's expected share of adopted projects with (ii) the share of projects that the agent would adopt if he observed the principal's posterior expected value but not particular realizations of attributes within the sample. That is, if the agent had access only to the object based on which principal decides on adoption –namely, the principal's posterior expected value– what's the expected share of projects he would adopt? When the principal is ex-ante indifferent between adoption and rejection, the agent samples the attributes that maximize the explained variation (propositions 4.1 and 4.7).

With a minimal capacity of one additional sampled attribute, the sufficient statistics simplify to the informativeness of the sampled attribute for each player, because the posterior expected values held by the two players are perfectly correlated. Hence, the agency conflict takes the form of whether the principal adopts more or less often than what the agent would after observing the realization of the sampled attribute. We introduce a concept of *local compromise*: a sampled attribute is a compromise if it is between the single-player local optima.[5] We show that under drastic disagreement, the optimal attribute is always a compromise. Due to the difference in the players' interpretation of the initial evidence, the agent prefers more informative attributes for both himself and the principal. Revisiting the federal program example, if Congress initially deems the program to be unpromising and the federal agency deems it promising, the agency chooses to treat a community of income between what the two players would prefer so as to overturn Congress's initial bias against the program.

Moreover, we show that for any two given weight functions, there is a region of initial mild disagreement that induces the agent to either overshoot or undershoot both single-player optima (proposition 4.4). That is, the optimal attribute is either of excessively high variance or of excessively low variance. For instance, if the federal agency in isolation from Congress samples richer communities, it might be willing to treat a community even richer when facing Congress. The outcome of the treated community is less informative than the agency's (or Congress's) preferred community for *both* players.

Such sampling arises whenever the principal reacts less strongly than the agent for any sampled attribute, that is, if she adopts a smaller (resp., larger) share of projects than the agent when both players deem the project promising (resp., unpromising). To align the players' expected share of adopted projects, the agent is willing to reduce the informativeness of the sampled attribute for both players. For instance, this is always the case when the principal is relatively more uncertain about the project than the agent. A principal who is ex-ante indifferent between adoption or rejection (hence, perfectly uncertain about the project) is always offered the most preferred attribute by the agent. If the principal is slightly more enthusiastic about the project, while the agent is moderately more enthusiastic about it, the agent prefers to move away from his most preferred attribute by reducing the informativeness of

[5]Local optima are guaranteed to be unique under mild conditions on the weight functions of the two players.

the sample for both himself and the principal, so as to induce a higher adoption rate by the more skeptical principal. The unavoidable loss in the sample's informativeness for the agent is a second-order effect. We further illustrate that such overshooting/undershooting arises even under ex-ante agreement when the weight functions are linear (example 2). Moreover, the phenomenon extends to capacities greater than one as well: the agent might settle for an optimal sample that can be locally modified to have both a higher informativeness for the principal and a higher explained variation for the agent (proposition 4.8).

Section 2 introduces the model and discusses key assumptions. The main analysis is presented in sections 3 and 4, which analyze single-agent evaluation and principal-agent evaluation respectively. We extend the analysis to two other configuration of sampling and adoption authorities (i.e., collective adoption and preemptive sampling) in section 5. Section 6 discusses an alternative model with finitely many correlated attributes, where many of the main results of our main model arise naturally. A comprehensive discussion of the related literature is presented in section 7. Results in the main text are proved in appendices A and B. Appendix C collects proofs and calculations related to the presented examples.

2 Model

2.1 Setup

An agent $\mathscr{A}$ (he) and a principal $\mathscr{P}$ (she) jointly consider the adoption of a multi-attribute project of unknown quality. The principal decides whether to adopt ($d = 1$) or reject ($d = 0$) the project. The agent decides which attributes are to be sampled prior to the adoption decision.

Project attributes. The project is characterized by a unit mass of attributes indexed by $a \in [0,1]$. The value of each attribute, $B(a)$, is determined by an unknown mapping $B : [0,1] \to \mathbb{R}$. This mapping is drawn according to the canonical probability measure from the set of one-dimensional Brownian paths with drift μ, variance σ^2, and constrained to go through $(a_0, B(a_0))$ for some $a_0 \in [0,1]$ (**?**).[6] The mapping B is drawn prior to the game between the principal and the agent. The players perfectly know (μ, σ). Without loss, we set $\sigma = 1$.[7]

The principal and the agent know $(a_0, B(a_0))$ but do not know the realizations of other attributes.[8] Let $B_0 := B(a_0)$ denote the realization for this known attribute. All other attribute realizations are distributed normally:

$$B(a) \sim \mathcal{N}\left(B_0 + \mu(a - a_0), \quad |a - a_0|\right) \tag{1}$$

[6] Note that unlike in the standard Brownian motion, $B(0)$ need not be equal to zero.

[7] The variance σ scales the variance of the posterior expected value pertaining to each sample of attributes, but does not affect optimal sampling.

[8] Fixing the realization of a_0, attribute realizations in $[0, a_0]$ and $[a_0, 1]$ follow two independent Brownian motions, each of which starts at $(a_0, B(a_0))$.

Note that B_0 determines the expected realization for all other attributes, but not their variances. Figure 1 illustrates this.

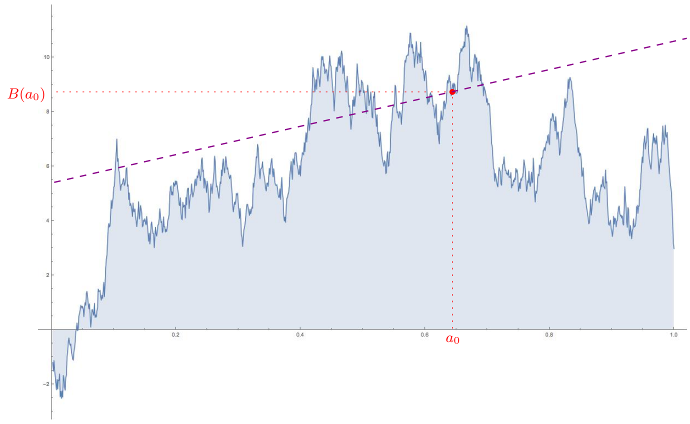

Figure 1: Initial evidence (a_0, B_0) with drift $\mu > 0$.

Timing and actions. Time is discrete and consists of two periods. At $t = 1$, the agent decides on a finite set of attributes that is to be sampled publicly. The agent perfectly commits to not sample further after the realizations of this sample are observed by both players. At $t = 2$, the principal makes an approval decision $d \in \{0, 1\}$. All actions are observable by both players.

Payoffs and conflict of interest. If the project is rejected $(d = 0)$, each player obtains a sure payoff of zero. If the project is adopted $(d = 1)$, the payoff of player $i \in \{\mathscr{A}, \mathscr{P}\}$ is given by: [9]

$$v_i(B) := \int_0^1 \omega_i(a) B(a) da. \tag{2}$$

Player i's payoff from the project is additive on the weighted realizations of all attributes. The weight $\omega_i(a)$ captures the importance of attribute a to player i. The weight functions satisfy the following assumption.

Assumption 1. *For $i \in \{\mathscr{P}, \mathscr{A}\}$, the weight function $\omega_i(\cdot) : [0, 1] \to [0, \infty)$ is continuous, weakly positive, and there exists $a \in [0, 1]$ such that $\omega_i(a) > 0$.*

The continuity of ω coupled with a compact attribute space guarantees that the weight function is bounded. Nearby attributes are attached similar weights. Both players prefer positive realizations for all attributes.

The ex-ante expected value of the project to player i is

$$v_0^i := \mathbb{E}[v_i(B) \mid (a_0, B_0)] = \int_0^1 (B_0 + \mu(a - a_0))\omega_i(a) da \tag{3}$$

$$= B_0 \int_0^1 \omega_i(a) da + \mu \int_0^1 \omega_i(a)(a - a_0) da \tag{4}$$

[9]Due to the almost sure continuity and boundedness of a realized Brownian path on $[0, 1]$, the payoff function is a well-defined Riemann integral pathwise.

We refer to v_0^i as the initial promise of the project. The project is deemed as *initially promising* for player i if $v_i^0 \geq 0$ and as initially unpromising otherwise. The sensitivity of v_0^i with respect to B_0 is captured by the sum of the weights of all attributes. Without loss, the analysis normalizes these sums to unity.[10] That is,

$$\int_0^1 \omega_{\mathscr{A}}(a)da = \int_0^1 \omega_{\mathscr{P}}(a)da = 1.$$

This normalization implies that when $\mu = 0$, the two players fully agree about the ex-ante value of the project, equal to the initial observation B_0. The two ex-ante values are generically different for $\mu \neq 0$.[11]

Sampling of attributes by the agent is costly. We assume that the agent is bound by an exogenous finite capacity $q \in \mathbb{N}$. That is, the cost function of sampling any n attributes is:

$$c(k) = \begin{cases} 0 & \text{for } n \leq q \\ +\infty & \text{for } n > q. \end{cases} \tag{5}$$

Information and sampling strategies. The players are symmetrically informed at any point of the game. All actions are publicly observable. Define the set of samples of size n as

$$\mathscr{S}_n := \left\{(a_1, \ldots, a_n) \in [0,1]^n : a_1 < a_2 < \ldots < a_n\right\}.$$

Let $\mathbf{s}$ denote an element in $\mathscr{S}_n$, and let $\mathscr{S}_n(a_0) \subset \mathscr{S}_n$ denote the subset of samples that include a_0. A sampling strategy of the agent $s_{\mathscr{A}}$ maps initial evidence and capacity (a_0, B_0, q) into the space of samples of size less than or equal to $q+1$:

$$s_{\mathscr{A}}: \quad [0,1] \times \mathbb{R} \times \mathbb{N} \to \bigcup_{n \leq q+1} \mathscr{S}_n(a_0).$$

Let $\{s_{\mathscr{P}}^n\}_{n \in \mathbb{N}}$, where $s_{\mathscr{P}}^n : \mathscr{S}_n \times \mathbb{R}^n \to \{0,1\}$, denote the adoption strategy for the principal. Strategy $s_{\mathscr{P}}^n$ maps a sample $\mathbf{s}$ of n distinct attributes and their respective realizations $\mathbf{B}(\mathbf{s})$ into an adoption decision.

2.2 Basic observations on inference

This subsection establishes some useful facts regarding the posterior belief about the value of the project after sampling takes place. Let $\mathbf{s} = (a_1, \ldots, a_n) \in \mathscr{S}_n$ and respective realizations

[10] Formally, starting from two arbitrary weight functions $(\omega_{\mathscr{A}}, \omega_{\mathscr{P}})$ with respective sum of attributes $(\Omega_{\mathscr{A}}, \Omega_{\mathscr{P}})$, define $\tilde{\omega}_i(a) := \omega_i(a)/\Omega_i$ for each player. This modification scales v_0^i and the standard deviation of the posterior expected value induced by a sample by $1/\Omega_i$ for each player. It does not affect the correlation among the players' posterior expected values. It scales each player's payoff by $1/\Omega_i$, but does not affect their decisions.

[11] The normalization does not guarantee that the adjustment term $\int_0^1 \omega_i(a)(a - a_0)da$ is the same for the two players. E.g., if $\mu > 0$ a player who values attributes close to one will have a higher ex-ante expected value than one who values attributes close to zero.

$\mathbf{B}(\mathbf{s}) = (B(a_1), \ldots, B(a_n))$. The realizations of other attributes are distributed normally:

$$B(a) \sim \begin{cases} \mathcal{N}\left(B(a_1) - \mu(a_1 - a), a_1 - a\right) & \text{for } 0 < a < a_1 \\ \mathcal{N}\left(B(a_i)\frac{a_{i+1}-a}{a_{i+1}-a_i} + B(a_{i+1})\frac{a-a_i}{a_{i+1}-a_i}, \frac{(a-a_i)(a_{i+1}-a)}{a_{i+1}-a_i}\right) & \text{for } a_i < a < a_{i+1}, \forall i \in [1, n-1], \\ \mathcal{N}\left(B(a_n) + \mu(a - a_n), a - a_n\right) & \text{for } a_n < a < 1. \end{cases}$$

The known realizations $\mathbf{B}(\mathbf{s})$ and μ appear in the mean of the distributions, but not in their variances. Realizations of attributes between any two sampled attributes follow a Brownian bridge. A sampled attribute $(a_i, B(a_i))$ enters only the distribution of attributes in $[a_{i-1}, a_i] \cup [a_i, a_{i+1}]$: hence, inference to unsampled attributes is local. Attributes beyond the farmost sampled attributes a_1 and a_n follow a Brownian motion with drift μ. The distribution of attributes in $[0, a_1]$ and $[a_n, 1]$ is informed only by a_1 and a_n respectively.

Let $v(\mathbf{s}, \mathbf{B}(\mathbf{s}))$ denote the realized posterior expected value of the project given the realizations of sample $\mathbf{s}$. This posterior expected value is a linear combination of all observed realizations.

Lemma 2.1. *The expected value of the project given sample* $\mathbf{s}$ *and realizations* $\mathbf{B}(\mathbf{s})$ *is:*

$$v(\mathbf{s}, \mathbf{B}(\mathbf{s})) = \sum_{i=1}^{n} B(a_i)\tau(a_i; \mathbf{s}) + \mu\left(\int_{a_n}^{1} (a - a_n)\omega(a)da - \int_{0}^{a_1} (a_1 - a)\omega(a)da\right), \tag{6}$$

where

$$\tau(a_i; \mathbf{s}) = \begin{cases} \int_0^{a_1} \omega(a)da + \int_{a_1}^{a_2} \frac{a_2-a}{a_2-a_1}\omega(a)da & \textit{for} \quad i = 1, \\ \int_{a_{i-1}}^{a_i} \frac{a-a_{i-1}}{a_i-a_{i-1}}\omega(a)da + \int_{a_i}^{a_{i+1}} \frac{a_{i+1}-s}{a_{i+1}-a_i}\omega(a)da & \textit{for} \quad i = 2, \ldots, n-1, \\ \int_{a_{n-1}}^{a_n} \frac{a-a_{n-1}}{a_n-a_{n-1}}\omega(a)da + \int_{a_n}^{1} \omega(a)da & \textit{for} \quad i = n. \end{cases}$$

Each realization $B(a_i)$ is scaled by a sample-dependent coefficient $\tau(a_i; \mathbf{s})$. This coefficient reflects the interval of attributes in the vicinity of a_i the distribution of which is affected by the realized $B(a_i)$. For instance, consider a two-attribute sample (a_1, a_2). The realization of a_1 informs the agent's belief about attributes smaller than a_1 and about those between a_1 and a_2. The first term of $\tau(a_1; (a_1, a_2))$ captures the fact that $B(a_1)$ is the agent's only source of information for extrapolation to attibutes in $[0, a_1]$; hence the weights of these attributes are fully accounted in $\tau(a_1; (a_1, a_2))$. On the other hand, the mean of attributes between a_1 and a_2 is a weighted sum of both $B(a_1)$ and $B(a_2)$. Hence, the second term of $\tau(a_1; (a_1, a_2))$ reflects the importance of $B(a_1)$ for forming beliefs about attributes in $[a_1, a_2]$: the weights of these attributes are "prorated" by the share of $B(a_1)$ in the inferred means of these attributes. $B(a_1)$ is irrelevant for inference to attributes in $[a_2, 1]$, hence the weights of these attributes do not appear in $\tau(a_1; (a_1, a_2))$. Similarly, $\tau(a_2; (a_1, a_2))$ reflects the importance of realization $B(a_2)$ for inference to attributes in $[a_1, a_2]$ and $[a_2, 1]$. In particular, notice that the sensitivity of the expected value to each observed realization, captured by τ, is unaffected by the drift.

The expected value $v(\mathbf{s}, \mathbf{B}(s))$ also features an additional term that reflects the importance of the drift for extrapolation to peripheral attributes in $[0, a_1)$ and $(a_n, 1]$. For any other attribute strictly between two sampled attributes, inference is based solely on the realizations of these sampled attributes (but not on the drift) due to the created Brownian bridge. To understand better this term, consider $\mu > 0$. Realizations of attributes in $[a_n, 1]$ are expected to be strictly greater than $B(a_n)$, while those of attributes in $[0, a_1]$ are expected to be strictly lower than $B(a_1)$. Therefore, the last term takes the difference between the sums of weights in these two intervals of attributes.

Given a sequence of samples of varying sizes $(\mathbf{s}_t)_{t=0}^{\infty}$, define the random process

$$v_t := \mathbb{E}\left[v(B) \mid \mathbf{s_0}, \dots, \mathbf{s_t}, \mathbf{B}(\mathbf{s_0}), \dots, \mathbf{B}(\mathbf{s_t})\right],$$

which describes the path followed by the posterior expected value as samples are taken in order. Note that $\mathbf{s}_0 = a_0$. The following lemma establishes three useful properties of this process.

Lemma 2.2 (Evolution of the posterior expected value)**.** *Fix a sampling sequence* $(\mathbf{s}_t)_{t=1}^{\infty}$. *The stochastic process* $(v_t)_{t\in\mathbb{N}}$ *is such that:*

(i) v_t *is normally distributed for any* t*, given* $(\mathbf{s}_0, \dots, \mathbf{s}_k)$ *and* $(\mathbf{B}(\mathbf{s}_0), \dots, \mathbf{B}(\mathbf{s}_k))$ *for any* $k < t$*;*

(ii) $(v_t)_{t\in\mathbb{N}}$ *is a martingale:* $\mathbb{E}(v_{t+1} \mid v_0, v_1, \dots, v_t) = v_t$ *and* $\mathbb{E}(|v_t|) < \infty$ *for any* t*;*

(iii) $(v_t)_{t\in\mathbb{N}}$ *is Markov:* $\Pr(v_t = \bar{v} \mid v_0, v_1, \dots, v_{t-1}) = \Pr(v_t = \bar{v} \mid v_{t-1})$.

2.3 Discussion of model features

Correlation structure across attributes. The model assumes the correlation structure among attributes to be described by a Brownian process with constant drift and variance.[12] This assumption captures four main features: (i) attributes close together are not expected to have dramatically different realizations, (ii) a conjecture about an unsampled attribute is formed based only on the realizations of the closest sampled attribute(s), (iii) the closer an unsampled attribute is to a sampled one, the less uncertain is its realization, and (iv) the magnitude of this uncertainty is independent of observed realizations. While modeling based on a particular stochastic process is unavoidably restrictive, this process is tractable and captures features that are natural in the motivating examples.

Common prior knowledge about a single attribute. The prior knowledge about (a_0, B_0) provides a context for further evaluation of the project by $\mathscr{P}$ and $\mathscr{A}$. Not only does such knowledge determine the players' prior belief about the quality of the project, but it also pins down the variability of all other attributes. By limiting this baseline knowledge to a single attribute, we capture real-world examples in which individuals evaluate projects which they know little

[12] One might wonder whether our analysis can be extended to a Brownian motion with variable (e.g., attribute-dependent) drift, e.g., $\mu(a) = a$. The main barrier is that, when pinned down at two sampled attributes $(a_1, B(a_1))$ and $(a_2, B(a_2))$, such processes might no longer result in Brownian bridges. Hence, our methods and results do not immediately extend to such processes.

about, for instance due to limited exposure or experience with similar projects in the past. Also, restricting attention to a single attribute permits a clearer analysis of how the centrality of a_0 in $[0,1]$ (e.g., distance from the median attribute $1/2$) affects optimal sampling.

Yet, this assumption precludes two other interesting scenarios: (i) no attribute realizations are known, but $\mathscr{P}$ and $\mathscr{A}$ share a common prior on $B(a)$ for any $a \in [0,1]$, and (ii) two or more realizations are initially known. The first scenario requires and depends upon the specification of a sensible prior over Brownian paths. Appendix D.1 offers some preliminary remarks in this direction. Regarding the second scenario, the key insights presented in our analysis extend to prior knowledge of multiple attributes.

Moreover, we assume that (a_0, B_0) is observable to both players. This informational symmetry focuses the analysis on the distortions in the production of new information rather than in the truthful reporting of existing information. Nonetheless, an implication of our analysis is that if $\mathscr{P}$ is privately informed about some realizations and $\mathscr{A}$ knows which attributes $\mathscr{P}$ is informed about, $\mathscr{A}$ is unable to elicit truthfully this information from $\mathscr{P}$ by conditioning the sampled attributes on $\mathscr{P}$'s report of the realizations.

Capacity-constrained sampling. The assumption of costless sampling up to an exogeneous capacity q is a key feature of the model. The agent $\mathscr{A}$ is bound by limited time, material and human resources, and/or cognitive capabilities to perform more than a given amount of sampling. For instance, when asked to evaluate a piece of legislation, a legislative committee has a limited number of subcommittees and teams to allocate to the investigation of particular aspects of the legislation. The committee staff is typically fixed, so evaluation cannot involve more issues than what can be investigated by the fixed staff. $\mathscr{A}$ may also be constrained by the information that $\mathscr{P}$ can process within the decision timeframe. For example, if Senate members can only deliberate on a limited number of aspects due to floor debate limits, the informing committee has to consider this limit.[13]

Tangentially to the analysis of capacity-based sampling, appendix D.3 offers some results on the single-player problem with a fixed cost per sampled attribute. It fully characterizes optimal simultaneous sampling, and offers some insights into optimal sequential sampling. The analysis reveals the added difficulties that this cost specification presents.

3 Single-agent evaluation

This section analyzes optimal sampling and adoption when both decisions are made by a single agent. The ex-ante expected value of the project is:

$$v_0 = B_0 + \mu \int_0^1 (a - a_0)\omega(a)\, da.$$

Subsection 3.1 characterizes optimal simultaneous sampling and it analyzes the effect of a larger capacity on optimal sampling. Subsection 3.2 establishes an equivalence result be-

[13]For instance, per U.S. Senate rules, a so-called reconciliation bill is limited to 20 hours of floor debate.

tween sequential and simultaneous sampling. Subsection 3.3 analyzes how the optimal sample changes with shifts in the weight function ω and in the initial attribute a_0. The final subsection 3.4 entertains a one-directional interpretation of the attribute domain. It draws implications for the special case in which attributes are reintepreted as inspection times of a project of uncertain performance over time.

3.1 Optimal simultaneous sampling

Given the initial evidence (a_0, B_0) and capacity q, the agent first decides which additional attributes $\mathbf{s} \in \mathscr{S}_{q+1}(a_0)$ to sample. Upon observing their realizations, she adopts the project only if the posterior expected value of the project is above zero. Lemma 2.1 derived the contribution of each attribute realization in the posterior expected value. Naturally, the closer a sampled attribute is to other attributes in the sample and the smaller are the weights of unsampled attributes in its vicinity, the smaller is its contribution to the posterior expected value. Moreover, the drift serves only in the extrapolation to unsampled attributes beyond the outermost attributes in the sample.

3.1.1 Features of the attribute-specific adoption threshold

Suppose the agent already knows the realizations of attributes $a_0, a_1, \ldots, a_{n-1}$ –let them be $B(a_0)$, $B(a_1)$, ..., $B(a_{n-1})$,– and is deciding how to judge the realization of a newly sampled attribute a_n. We call *adoption threshold* of a_n the critical realization of a_n, denoted by $\bar{B}(a_n)$, that leaves the agent indifferent between adoption and rejection. Naturally, this adoption threshold depends on the other attributes $a_0, \ldots, a_n$ and their realizations, which are surpressed in the notation of the adoption threshold.

Proposition 3.1, presented below, shows that even if all observed realizations are negative (resp., positive)–hence, all the collected evidence is in itself discouraging (resp., encouraging), – the project is still adopted (resp., rejected) if the drift is sufficiently steep and at least one of the endpoint attributes ($a = 0$ or $a = 1$) remains unsampled. The second part of proposition 3.1 points to the possibility of radically different adoption standards even for attributes very close to each other. For instance, suppose that $q = 1$ and consider the *adoption thresholds* of $a_0 - \varepsilon$ and $a_0 + \varepsilon$, i.e. the respective realizations that leave the agent indifferent between adopting and rejecting. The two adoption thresholds are generically different, due to the radically different implications that the agent draws from sampling these attributes. A realization from $a_0+\varepsilon$ (resp., $a_0-\varepsilon$) sets the expectation for all $a \in (a_0, 1]$ (resp., $[0, a_0)$), while B_0 sets the expectation for all $[0, a_0)$ (resp., $(a_0, 1]$). Hence, the sampling of two attributes that are arbitrarily close to each other and to a_0, but on different sides of a_0, generates very different knowledge about the project. Figure 2 illustrates this discontinuity for $\mu = 0$ and $B_0 > 0$.

Proposition 3.1 (Optimal adoption)**.** *Consider sample* $\mathbf{s} \in \mathscr{S}_{q+1}(a_0)$ *with realizations* $\mathbf{B}(\mathbf{s}) \in \mathbb{R}^{q+1}$.

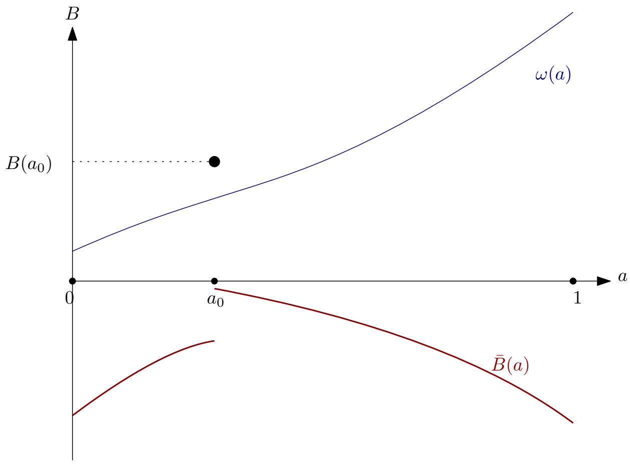

Figure 2: Discontinuity of the adoption threshold around a_0 for $B_0 > 0$ and $\mathbf{s} = a_0$.

(i) Suppose that all realizations are positive (resp., negative), i.e. $\mathbf{B}(\mathbf{s}) < \mathbf{0}$ *(resp.,* $\mathbf{B}(\mathbf{s}) > \mathbf{0}$*). If* $\mathbf{s}$ *does not include both* $a = 0$ *and* $a = 1$*, there exists* $\mu \neq 0$ *with* $|\mu|$ *sufficiently large that leads to adoption (resp., rejection) of the project.*

(ii) Fix any $a \notin \mathbf{s}$*. The adoption threshold* $\bar{B}(a)$ *is discontinuous at* $a = a'$ *for any* $a' \in \mathbf{s}$*.*

Proposition 3.1 highlights two features of optimal sampling that, in the context of the applications mentioned in the introduction, might be naively mistaken as signs of lack of due diligence in the evaluation process. First, adoption (rejection) in the face of discouraging (encouraging) news might be justified if inference to unsampled attributes is sufficiently positive (negative). Second, different adoption thresholds for arbitrarily close attributes might be optimal when the sampling of the two attributes leads to very different inferences for unsampled attributes.

3.1.2 Optimal sampling

Let us remind the reader that $v(\mathbf{s}, \mathbf{B}(\mathbf{s}))$ denotes the posterior expected value calculated from the observed realizations $\mathbf{B}(\mathbf{s})$ from sample $\mathbf{s}$. For any discovered sample $\mathbf{s} \in \mathscr{S}_{q+1}(a_0)$, the project is adopted if and only if the posterior expected value $v(\mathbf{s}, \mathbf{B}(\mathbf{s})) \geq 0$. Therefore, given capacity q, the problem of the agent consists in simultaneously choosing a sample of attributes that maximizes the expected value of the project conditional on adoption scaled by the probability of adoption:

$$\max_{\mathbf{s} \in S_{q+1}(a_0)} \quad \Pr\Big(v(\mathbf{s}, \mathbf{B}(\mathbf{s})) \geq 0\Big)\mathbb{E}\Big[v(\mathbf{s}, \mathbf{B}(\mathbf{s})) \mid v(\mathbf{s}, \mathbf{B}(\mathbf{s})) \geq 0\Big] \tag{7}$$

From lemma 2.2, we know that for any sample $\mathbf{s} \in S_{q+1}(a_0)$ the posterior expected value is normally distributed with mean v_0. Let $\sigma^2(\mathbf{s})$ denote the variance of $v(\mathbf{s}, \mathbf{B}(\mathbf{s}))$. The adoption probability is given by $\Phi\left(\frac{v_0}{\sigma(\mathbf{s})}\right)$, where Φ denotes the cdf of the standard normal distribution. Intuitively, an initially promising project is more likely to be adopted than not after sampling, i.e. it has an adoption probability greater than 1/2. Moreover, when conditioning on future adoption, the expected value of the project is strictly higher than v_0:

$$\mathbb{E}\left[v(\mathbf{s},\mathbf{B}(\mathbf{s})) \mid v(\mathbf{s},\mathbf{B}(\mathbf{s})) \geq 0\right] = v_0 + \sigma(\mathbf{s}) \frac{\phi\left(\frac{v_0}{\sigma(\mathbf{s})}\right)}{\Pr(v(\mathbf{s},\mathbf{B}(\mathbf{s})) \geq 0)},$$

where ϕ is the pdf of the standard normal distribution. Hence, the objective of the agent can be rewritten as:[14]

$$\max_{\mathbf{s}\in S_{q+1}(a_0)} \quad v_0 \Phi\left(\frac{v_0}{\sigma(\mathbf{s})}\right) + \sigma(\mathbf{s})\phi\left(\frac{v_0}{\sigma(\mathbf{s})}\right). \qquad \text{(Sampling problem)}$$

This objective is increasing in $\sigma(\mathbf{s})$ despite the ex-ante expected value v_0. Therefore, the agent seeks to discover the sample with the highest induced variance on the posterior expected value. This variance $\sigma^2(\mathbf{s})$ depends on a_0 but not on B_0: hence the ranking of any two samples according to how informative they are for the agent's problem does not depend on the initial promise of the project. Therefore, the agent would be able to identify the optimal sample even if B_0 were unobservable to him (but not the optimal adoption decision).

Put differently, the adoption decision that the agent takes is a random variable before sample realizations are observed. The agent seeks to discover samples that maximize the variance of the adoption decision, which is just

$$\Phi\left(\frac{v_0}{\sigma(\mathbf{s})}\right)\left(1-\Phi\left(\frac{v_0}{\sigma(\mathbf{s})}\right)\right).$$

The adoption decision is most uncertain whenever $\Phi(v_0/\sigma(\mathbf{s}))$ is close to 1/2, which corresponds to samples with the highest possible $\sigma(\mathbf{s})$.

Proposition 3.2 (Optimal sampling under capacity q)**.**

(i) Given capacity q, an optimal sample solves

$$\mathbf{s}^* \in \underset{\mathbf{s}\in\mathscr{S}_{q+1}(a_0)}{\arg\max} \quad \sigma^2(\mathbf{s}).$$

(ii) For any optimal sample $\mathbf{s}^* = (a_1^*, a_2^*, \ldots, a_{q+1}^*)$, where $a_n^* < a_{n+1}^*$ for any $n \in [1, q+1]$ and $a_k^* = a_0$ for some $k \in [1, q+1]$, the outermost sampled attributes are interior, i.e. if $k < q+1$, then $a_{q+1}^* < 1$, and if $k > 1$, then $a_1^* > 0$.

The variance of $v(\mathbf{s},\mathbf{B}(\mathbf{s}))$ is pinned down by the attributes in the sample and the weight function. Given $\mathbf{s}$ let us relabel the attributes in increasing distance from a_0 as $a_1^\ell > \ldots > a_{k-1}^\ell$ in $[0, a_0]$ and $a_1^r < \ldots < a_{q+1-k}^r$ in $[a_0, 1]$. Lemma A.1 in appendix A provides an expression for the variance of the posterior:

$$\sigma^2(\mathbf{s}) := \sum_{i=1}^{k-1} \tau^2\left(a_i^\ell; (a_i^\ell, a_{i-1}^\ell)\right)(a_{i-1}^\ell - a_i^\ell) + \sum_{i=1}^{q+1-k} \tau^2\left(a_i^r; (a_{i-1}^r, a_i^r)\right)(a_i^r - a_{i-1}^r), \tag{8}$$

[14] This expression is derived in the proof of proposition 3.2.

where for any $m \in \{\ell, r\}$, $\tau(a_i^m; (a_i^m, a_{i-1}^m))$ denotes the coefficient corresponding to $B(a_i^m)$ in the posterior expected value *if only two attributes were known*, a_i^m and a_{i-1}^m. The variance $\sigma^2(\mathbf{s})$ is pairwise separable on any two adjacent attributes in the sample. Realizations of sampled attributes on the same side of a_0 are correlated; attributes closer to a_0 are more informative as their realizations set the expectation for attributes further away from a_0.

In identifying an optimal sample, the agent faces a tradeoff between the variability of a sampled attribute–determined by the distance of this attribute from other known attributes–and how much this attribute is informative about other sorrounding attributes, captured by τ. To see this, suppose that $q = 1$. An attribute further away from a_0 is more uncertain, and hence, in the absence of any correlation among attributes, more appealing to be discovered. When factoring in the correlation structure, we observe that the further an attribute a is from a_0, the smaller is its $\tau(a; (a_0, a))$, hence its realization is weighted by less in forming the posterior expected value. A more distant attribute from a_0 (i) is less informative about attributes between a_0 and itself, and (ii) serves as the only basis of inference for fewer attributes that are further away from a_0. The agent resolves this tradeoff by sampling interior attributes only, as proposition 3.2 *(ii)* establishes.

The following examples illustrates optimal simultaneous sampling if the agent weighs all attributes equally.

Example 1 (Quasi-representative sampling with uniform weights)**.** *Let* $\omega(a) = 1$ *for all* $a \in [0,1]$*. Fixing the number of attributes sampled in* $[0, a_0]$ *to* $k \le q$*, the optimal sample consists of:*

$$a_i^\ell = a_0 - \frac{2i}{2k+1} a_0 \qquad \text{for } 1 \le i \le k$$
$$a_i^r = a_0 + \frac{2i}{2(q-k)+1}(1-a_0) \qquad \text{for } 1 \le i \le q-k.$$

Two remarks are in order. First, attributes are sampled equidistantly between $a_0/(2k+1)$ *and* a_0*, and between* a_0 *and* $(2(q-k)+a_0)/(2(q-k)+1)$*. This is akin to representative sampling of the attribute continuum. The peripheral regions of attributes close to* $a = 0$ *and* $a = 1$ *remain unsampled. Figure 3 illustrates optimal sampling for* $a_0 = 1/3$ *as* q *varies from 1 to 5 attributes. Note that the smallest (largest) sampled attribute is smaller (larger) under a greater capacity. Moreover, as* q *increases, exploration becomes denser.*

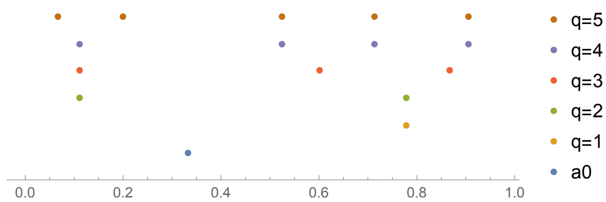

Figure 3: Quasi-representative optimal sampling in example 1 for $a_0 = 1/3$ and $q \le 5$.

Second, small changes in a_0 *might bring about very different optimal samples. Take, for* *instance,* $a_0 = 1/2 - \varepsilon$ *and* $a_0' = 1/2 + \varepsilon$ *with* $q = 1$*, where* ε *is arbitrarily small. As* ϵ *shrinks to zero, the optimal sample under* a_0 *tends to* $5/6$*, while that under* a_0' *tends to* $1/6$*. Two ex-ante arbitrarily similar projects are evaluated very differently.*[15]

Naturally, the agent prefers a greater capacity, because it allows her to gain a more thorough understanding of the project by discovering more attributes. What is less immediate is whether the added benefit from a marginally greater capacity (i.e., willingness to pay for one additional attribute) strictly decreases as the capacity gets larger. The following proposition confirms this. Let $V(q)$ denote the added value from sampling under capacity q, i.e.

$$V(q) := \max_{\mathbf{s}\in\mathscr{S}_{q+1}(a_0)} \nu_0 \Phi\left(\frac{\nu_0}{\sigma(\mathbf{s})}\right) + \sigma(\mathbf{s})\phi\left(\frac{\nu_0}{\sigma(\mathbf{s})}\right)$$

Naturally, this added value is strictly increasing and strictly convex in the maximum attained variance of the posterior expected value under capacity q. The following proposition compares V across capacities. The nontrivial part of the proof of proposition 3.3 establishes that that increase in the maximum attained variance as capacity increases for $q-1$ to q is greater than the increase in the maximum attained variance as capacity goes from q to $q+1$. An allowance of one additional attribute is less valuable (in terms of the gain in the variance of the posterior) the more thorough exploration is under the current capacity.[16]

Proposition 3.3 (Diminishing returns from exploration)**.** *Optimal sampling exhibits decreasing marginal returns from a marginally greater capacity, i.e. for any* $q \in \mathbb{N}$*,*

$$V(q+1) - V(q) < V(q) - V(q-1).$$

Figure 3 suggests that a larger sample is, roughly speaking, more spread out, denser, and with a weakly larger number of attributes in each of the initial intervals $[0, a_0]$ and $[a_0, 1]$. These features are not specific to the case of equal weights; they generalize to any weight function satisfying assumption 1. As capacity increases by exactly one attribute, the agent allocates the additional attributes to either $[0, a_0]$ or $[a_0, 1]$, without lowering the number of attributes taken in any of these two intervals. In the interval that is allocated the additional attribute, the smallest (resp., largest) sampled attribute becomes strictly smaller (larger). The span of exploration strictly expands in this region. Moreover, any sampled attribute in the larger sample falls between two sampled attributes in the smaller sample. Exploration becomes strictly denser. Proposition 3.4 formalizes these observations.

[15] Such discontinuities arise for $q > 1$ as well. Let $q = 2$. For $a_0 = 0.3$, the optimal sample is $(a_0, 0.58, 0.86)$, but for $a_0 = 0.31$, it is $(0.10\bar{3}, a_0, 0.77)$.

[16] Although motivated by a different problem, **?** uses similar arguments to study optimal interval division for capacity-constrained problems. The approach differs technically from **?** in three respects: (i) **?** assumes a value function that is pairwise additive; here such a value function (i.e. variance of the posterior expected value which is pairwise additive in adjacent attributes) arises endogeneously from the Brownian process, (ii) the variance is strictly supermodular, so methods from **?** and **?** can be invoked directly, and (iii) sampling attributes $a = 0$ and $a = 1$ in the sampling problem here is less trivial to deal with than partitioning the state space at $a = 0$ and $a = 1$ in the optimal division problem (where the partition simply creates degenerate cells).

Proposition 3.4 (Gradual and expansive exploration)**.** *Given optimal samples* $\mathbf{s} = (a_1,\ldots,a_{q+1})$ *and* $\mathbf{s}' = (a'_1,\ldots,a'_{q+2})$ *corresponding to capacities* q *and* $q+1$ *respectively:*

1. *if* $a_k = a_0$ *for some* $k \in [1,q+1]$*, then either* $a'_k = a_0$ *or* $a'_{k+1} = a_0$*;*
2. (a) *if* $a_k = a'_{k+1} = a_0$*, then* $a'_i < a_i < a'_{i+1}$ *for any* $i \in [1,k-1]$*,*

   (b) *if* $a_k = a'_k = a_0$*, then* $a'_i < a_i < a'_{i+1}$ *for any* $i \in [k+1,q+1]$*.*

To better appreciate the significance of proposition 3.4, consider a weight function that assigns high weights to a small neighborhood of attributes and low weights to all other attributes: for instance, such a weight function could be the pdf of a low-variance normal distribution, such as the pdf of $\mathcal{N}(1/2,\epsilon)$ for $\epsilon > 0$ small. In this example, the agent cares mostly about attributes around $1/2$. Suppose $a_0 = 0$. Proposition 3.4 establishes that despite the strong interest the agent has in a neighborhood of attributes around $1/2$, the two outermost sampled attributes move further away from $1/2$ as capacity increases.[17] Rather than exclusively using the extra capacity to learn more intensively about the same neighborhood around $1/2$, the agent expands the span of sampling to attributes he cares less about.

3.2 Equivalence of sequential and simultaneous sampling

Proposition 3.2 highlighted the irrelevance of the ex-ante expected value v_0 of the project for optimal simultaneous sampling. This observation extends to sequential sampling: the posterior expected value that the agent holds after having sampled any k attributes does not inform optimal sampling of the remaining $q-k$ attributes. Any two such subsamples can be ranked solely based on the variance they induce on the posterior expected value, which depends only on the attributes known thus far. As a result, the agent does not benefit from the flexibility of sampling attributes sequentially; the full optimal sample can be identified since the beginning of sampling. Put differently, it is not crucial for the agent to observe realizations in real time in order to be able to carry out optimal sampling.

Proposition 3.5. *The optimal sequence of attributes acquired through sequential sampling coincides with the optimal simultaneous sample.*

3.3 Shifts in the initial attribute and in the weight function

This subsection seeks to understand how the optimal sample changes with shifts of the initially known attribute a_0 and of the weight function ω. First, keeping ω fixed, as a_0 increases, the agent becomes more interested in sampling in $[0,a_0]$ and less interested in sampling in $[a_0,1]$. That is, for a given capacity, the greater a_0 is, the lower is the share of that capacity that the agent spends in sampling attributes smaller than a_0. In the context of the policymaking example, despite how the policymaker weighs the outcomes of different communities, a

[17] For a sufficiently small $\epsilon > 0$, the sampled attribute under $q = 1$ is arbitrarily close to $1/2$, whereas under $q = 2$ the two sampled attributes a_1 and a_2 are such that $a_1 < 1/2 < a_2$.

richer initial community encourages him to accord a higher share of the budget to the testing of communities poorer than the initial community. Moreover, for any two initial communities of arbitrarily close income levels a_0 and $a_0+\varepsilon$, the respective optimal samples that they generate cannot differ too much. More specifically, the number of trials allocated to communities richer than the initial community is either the same under both a_0 and $a_0+\varepsilon$, or it differs by one (i.e. x communities sampled in $[a_0,1]$ under a_0 and $x+1$ communities sampled in $[a_0+\varepsilon,1]$ under $a_0+\varepsilon$).[18] These two observations are summarized in proposition A.1 in appendix A.

Let any two weight functions ω and $\tilde{\omega}$ be such that $\omega(a_0)=\tilde{\omega}(a_0)$ and for any $a_2>a_1$, the ratio of the weight functions is nondecreasing in the attribute, i.e.

$$\frac{\omega(a_2)}{\omega(a_1)} \geq \frac{\tilde{\omega}(a_2)}{\tilde{\omega}(a_1)}. \tag{MLR}$$

The weight functions ω and $\tilde{\omega}$ are said to satisfy the monotone likelihood ratio property: ω is a monotone likelihood ratio shift of $\tilde{\omega}$ around a_0.[19] Loosely speaking, this captures a shift of the agent's interest toward higher-indexed attributes. The more distant an attribute is from a_0, the more pronounced the increase in its weight relative to its original weight. Intuitively, this shift should translate into more intensive sampling of higher-indexed attributes, both in terms of the number of attributes allocated to $[a_0,1]$ and in terms of higher-indexed sampled attributes within each interval $[0,a_0]$ and $[a_0,1]$.

The first part of proposition 3.6 establishes that an MLR shift of the weight function induces weakly more sampling in $[a_0,1]$. The weight of any attribute greater (resp., smaller) than a_0 increases (resp., decreases) as a result of the MLR shift. The variance induced by any subsample in $[a_0,1]$ increases, thus making sampling in $[a_0,1]$ more attractive for the agent. The second part of proposition 3.6 states that, conditional on allocating the same number of attributes to $[0,a_0]$ and $[a_0,1]$, each sampled attribute in each interval shifts further to the right. Sampling uniformly shifts towards higher-indexed attributes. This result is subject to the following single-crossing condition on the weight function.

Assumption 2 (Single-crossing condition)**.** *For any $a_1,a_2\in[0,1]$ such that $a_1<a_2$, the functions*

$$r(x)=\int_x^1 \omega(a)da-\int_{a_2}^x \frac{a-a_2}{x-a_2}\omega(a)da, \qquad x\in[a_2,1]$$
$$\ell(x)=\int_x^{a_1} \frac{a_1-a}{a_1-x}\omega(a)da-\int_0^x \omega(a)da, \qquad x\in[0,a_1]$$
$$m(x)=\int_x^{a_2} \frac{a_2-a}{a_2-x}\omega(a)da-\int_{a_1}^x \frac{a-a_1}{x-a_1}\omega(a)da, \qquad x\in[a_1,a_2]$$

[18] If the number of attributes sampled to the right of the initial attribute differs by one, the optimal sample within this interval is quite different for a_0 and $a_0+\varepsilon$. Sampling in this interval expands discontinuously as the initial attribute increases from a_0 to $a_0+\varepsilon$. This feature was illustrated earlier in the context of example 1.

[19] Because any weight function is normalized so that $\int_0^1\omega(a)da=1$, such a shift is akin to an MLR shift of a probability density function on $[0,1]$.

cross zero only once from above in their respective domains.

This single-crossing condition essentially requires that the weight function does not drop or peak too abruptly within a small neighborhood of attributes. In the context of the example, this means that the policymaker does not assign radically different weights to communities of similar median income. Lemma A.5 in appendix A provides stronger –but easier to interpret– conditions on the weight function that guarantee that assumption 2 holds. Within the context of $q = 1$, assumption 2 guarantees that the variance $\sigma^2(\mathbf{s})$ is single-peaked on $[0, a_0]$ and $[a_0, 1]$. We return to this observation in the principal-agent analysis in section 4.

Proposition 3.6 (MLR shift of the weight function)**.** *Suppose ω is a MLR shift of $\tilde{\omega}$ around a_0. Let $\mathbf{s} = (a_1, \ldots, a_{q+1})$ and $\tilde{\mathbf{s}} = (\tilde{a}_1, \ldots, \tilde{a}_{q+1})$ be two optimal samples corresponding to ω and $\tilde{\omega}$ respectively, and let $a_k = \tilde{a}_{k'} = a_0$.*

(i) $\mathbf{s}$ *features more sampled attributes in $[a_0, 1]$ than $\tilde{\mathbf{s}}$, i.e. $k' \geq k$.*

(ii) Suppose condition 2 holds, and $k = k'$. Then, $a_i > \tilde{a}_i$ for any $i \neq k$.

3.4 One-directional sampling: attributes as time

This subsection briefly remarks on a special case of single-agent evaluation. Let us reinterpet the domain $[0, 1]$ as the lifespan of a (single) pilot program, by the end of which the agent decides whether to roll out the program at large scale. The agent knows the initial state of the program at $t = 0$; its state over time follows a Brownian process. The pilot program is too insignificant to affect the agent's utility if rejected for large-scale implementation. If approved, the progress (i.e. realized state path) from the large-scale program is assumed to be identical to that of the pilot program. The agent cares about the entire progress of the program; the weight function captures the intrinsic interest the agent has in the state at different stages of the program. The agent has limited opportunities to inspect its progress in real time.

If constrained to inspect the program only at q particular points in time, what is the optimal timing of inspections? Our single-agent analysis offers the following insights:

1. The expected value of the program at $t = 0$ does not affect the optimal inspection schedule. By all means, an initially unpromising pilot program has a lower chance of being implemented at large scale at $t = 1$, but all pilot programs are inspected identically. The program is never inspected at its end at $t = 1$.

2. Observed progress does not affect the optimal timing of future inspections. The player knows the optimal timing of inspections since $t = 0$.

3. With more opportunities to inspect the program, the agent takes the first inspection earlier and the last one later. Moreover, if given the opportunity to take one additional inspection (i.e. as capacity increases by exactly one unit), the agent schedules the k^{th} inspection earlier than before, but later than the time of the $(k-1)^{th}$ inspection under the smaller capacity.

4. If the player takes an equal interest in all stages of the program (i.e. $\omega(t) = 1$ for all $t \in [0,1]$), he schedules periodic inspections. That is, he inspects the program at equally spaced intervals.

Tangentially, it is interesting to compare the optimal behavior described above with the sequence of inspections chosen by a myopic player who decides on the next inspection as if capacity is equal to one. It follows from our analysis that the myopic player waits too long until the next inspection. He fails to recognize the benefit of an earlier inspection for the informativeness of future inspections. As a result, she crams inspections towards the end of the program horizon. Figure 4 contrasts the myopic and optimal schedules for $\omega(t) = 1\ \forall t$ and $q = 4$.

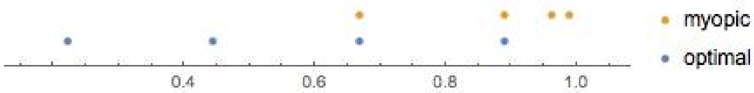

Figure 4: A myopic player schedules delayed inspections.

4 Principal-agent evaluation game

This section turns attention to the strategic interaction between an agent endowed with the authority to sample attributes and a principal endowed with the authority to adopt. The agent decides which attributes are to be sampled simultaneously and publicly, subject to sampling capacity $q \in \mathbb{N}$. The principal observes the sampled realizations and decides whether to adopt the project. The initial disagreement is captured by the ex-ante expected values for the two players: $v_0^{\mathscr{P}}$ and $v_0^{\mathscr{A}}$. We distinguish between two types of disagreement. If the two ex-ante expected values hold opposite signs, the players disagree on the adoption decision absent any sampling: we refer to this as *drastic disagreement.* If they hold the same sign, the players agree on whether to adopt but differ in the expected payoff from this decision: we call this *mild disagreement.* After sample $\mathbf{s}$ is discovered, the players update to the posterior expected values $v^{\mathscr{P}}(\mathbf{s}) := v^{\mathscr{P}}(\mathbf{s}, \mathbf{B}(\mathbf{s}))$ and $v^{\mathscr{A}}(\mathbf{s}) := v^{\mathscr{A}}(\mathbf{s}, \mathbf{B}(\mathbf{s}))$.

4.1 Characterization of the agent's problem

After the agent samples $\mathbf{s}$ publicly, the principal approves if and only if $v^{\mathscr{P}}(\mathbf{s}) \geq 0$. Hence, the agent's sampling problem is:

$$\begin{aligned} \max_{\mathbf{s}} \quad & \underbrace{\Pr\left(v^{\mathscr{P}}(\mathbf{s}) \geq 0\right)}_{\text{adoption probability}} \cdot \underbrace{\mathbb{E}\left(v^{\mathscr{A}}(\mathbf{s}) \mid v^{\mathscr{P}}(\mathbf{s}) \geq 0\right)}_{\text{expected value of an adopted project for} \mathscr{A}} \qquad (\mathscr{A}\text{'s problem}) \\ \text{subject to} \quad & \mathbf{s} \in \bigcup_{1 \leq k \leq q+1} \mathscr{S}_k(a_0). \end{aligned}$$

This formulation allows for the possibility that the agent does not exhaust q. Due to the conflict of interest inherent in the interaction, she might find it beneficial to restrict the number of attributes that inform the decision of the principal.

A sample $\mathbf{s} = (a^{\ell}_{n_\ell}, \dots, a^{\ell}_1, a_0, a^r_1, \dots, a^r_{n_r})$, with attributes relabeled in increasing distance from a_0, induces a distribution over the pair of random variables $\left(v^{\mathscr{P}}(\mathbf{s}), v^{\mathscr{A}}(\mathbf{s})\right)$. The two posterior expected values are jointly normal, with covariance and variance respectively:

$$\operatorname{cov}\left(v^{\mathscr{P}}(\mathbf{s}), v^{\mathscr{A}}(\mathbf{s})\right) = \sum_{k\in\{\ell,r\}} \sum_{i=1}^{n_k} \tau_{\mathscr{P}}(a^k_i; (a^k_i, a^k_{i-1}))\tau_{\mathscr{A}}(a^k_i; (a^k_i, a^k_{i-1}))|a^k_i - a^k_{i-1}| := \operatorname{cov}(\mathbf{s}),$$

$$\sigma^2_j\left(v^j(\mathbf{s})\right) = \sum_{k\in\{\ell,r\}} \sum_{i=1}^{n_k} \tau^2_j(a^k_i; (a^k_i, a^k_{i-1}))|a^k_i - a^k_{i-1}| := \sigma^2_j(\mathbf{s}), \text{ for } j = \mathscr{P}, \mathscr{A}.$$

Note that the three sample statistics, $\operatorname{cov}(\mathbf{s}), \sigma^2_{\mathscr{A}}(\mathbf{s})$, and $\sigma^2_{\mathscr{P}}(\mathbf{s})$ are independent of the initial observation B_0. The first one captures the joint variability of the expected values $v^{\mathscr{A}}(\mathbf{s})$ and $v^{\mathscr{P}}(\mathbf{s})$: the higher the covariance is, the more aligned the interests of the two players are expected to be in the adoption stage. Because higher attribute realizations are always good news for both players, covariance is positive for any feasible sample. The second and the third statistic capture the informativeness of the sample $\mathbf{s}$ for each player. The analysis in section 3 showed that in the absence of the other player, each player discovers the sample that induces the most variable posterior expected value, i.e. the one that maximizes $\sigma^2_j(\mathbf{s})$ for $j = \mathscr{P}, \mathscr{A}$.

Given (v_0, a_0) and a sample $\mathbf{s}$, let $P(\mathbf{s}; v^{\mathscr{P}}_0, v^{\mathscr{A}}_0)$ denote the expected payoff of the agent from discovering this sample. The following lemma shows that for the purpose of the agent's problem, each sample can be summarized by two sufficient statistics: $\operatorname{cov}(\mathbf{s})$ and $\sigma^{\mathscr{P}}(\mathbf{s})$. The adoption decision is made with respect to the induced posterior of the principal. The agent translates the consequences of an adoption (i.e. $v^{\mathscr{P}}(\mathbf{s}) \geq 0$) into what that implies for her own posterior expected value. But observe that the two posteriors are jointly normal, hence the conditional distribution $v^{\mathscr{A}}(\mathbf{s}) \mid v^{\mathscr{P}}(\mathbf{s})$ is normal as well. This distribution is fully described by the respective variances of $v^{\mathscr{A}}(\mathbf{s})$ and $v^{\mathscr{P}}(\mathbf{s})$, their correlation coefficient, and the ex-ante values $(v^{\mathscr{A}}_0, v^{\mathscr{P}}_0)$.

Proposition 4.1 (Sufficient statistics for $\mathscr{A}$'s problem)**.** *Given a sample* $\mathbf{s}$*, the payoff of the agent depends on* $\mathbf{s}$ *only through two sample statistics:* $\operatorname{cov}(\mathbf{s})$ *and* $\sigma^{\mathscr{P}}(\mathbf{s})$*, i.e.,*

$$P(\mathbf{s}; v^{\mathscr{P}}_0, v^{\mathscr{A}}_0) = v^{\mathscr{A}}_0 \Phi\left(\frac{v^{\mathscr{P}}_0}{\sigma_{\mathscr{P}}(\mathbf{s})}\right) + \frac{\operatorname{cov}(\mathbf{s})}{\sigma_{\mathscr{P}}(\mathbf{s})} \phi\left(\frac{v^{\mathscr{P}}_0}{\sigma_{\mathscr{P}}(\mathbf{s})}\right). \tag{9}$$

Writing the agent's payoff in terms of the correlation coefficient $\rho(\mathbf{s})$ and $\sigma_{\mathscr{A}}(\mathbf{s})$ rather than $\operatorname{cov}(\mathbf{s})$ offers a natural interpretation of the strategic considerations at play here. The correlation coefficient between the two posteriors $v_{\mathscr{P}}(\mathbf{s})$ and $v_{\mathscr{A}}(\mathbf{s})$ is given by

$$\rho(\mathbf{s}) = \frac{\operatorname{cov}(\mathbf{s})}{\sigma_{\mathscr{P}}(\mathbf{s})\sigma_{\mathscr{A}}(\mathbf{s})}.$$

The correlation ρ captures how much the players agree on the relative importance of the realizations of particular attributes within the sample. The more they agree on this, the more strongly the posterior expected values of the two players are related. The payoff of the agent is given by:

$$P(\mathbf{s}; v_0^{\mathscr{P}}, v_0^{\mathscr{A}}) = v_0^{\mathscr{A}} \Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right) + \underbrace{\rho(\mathbf{s})\sigma_{\mathscr{A}}(\mathbf{s})}_{\text{variation of } v^{\mathscr{A}}(\mathbf{s}) \text{ explained by } v^{\mathscr{P}}(\mathbf{s})} \phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right).$$

For the jointly normal distribution of $(v^{\mathscr{P}}(\mathbf{s}), v^{\mathscr{A}}(\mathbf{s}))$, the correlation $\rho(\mathbf{s})$ captures the proportion of the dispersion in one player's posterior that can be explained by the other player's posterior. That is, $\rho^2(\mathbf{s})$ measures the proportion of the variance of $v^{\mathscr{A}}(\mathbf{s})$ that is accounted for by $v^{\mathscr{P}}(\mathbf{s})$: if the agent were to predict her posterior expected value $v^{\mathscr{A}}(\mathbf{s})$ from observing only the realized posterior of the principal $v^{\mathscr{P}}(\mathbf{s})$, how good would that prediction be? Correspondingly, $\rho^2(\mathbf{s})\sigma^2_{\mathscr{A}}(\mathbf{s})$ measures the explained variation in $v^{\mathscr{A}}(\mathbf{s})$. The higher is this explained variation, the better the adoption decision made by the principal reflects the interests of the agent, therefore *ceteris paribus* the agent prefers samples with higher explained variance. An ideal sample would induce perfect correlation between the two posterior expected values, in which case the explained variance of $v^{\mathscr{A}}(\mathbf{s})$ would equal its actual variance $\sigma^2_{\mathscr{A}}(\mathbf{s})$. The conditional expectation for the agent given that the principal adopts is

$$\mathbb{E}[v^{\mathscr{A}}(\mathbf{s}) \mid v^{\mathscr{P}}(\mathbf{s}) \geq 0] = v_0^{\mathscr{A}} + \rho(\mathbf{s})\sigma_{\mathscr{A}}(\mathbf{s}) \frac{\phi\left(-\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right)}{1 - \Phi\left(-\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right)}.$$

The higher the explained variance, the higher is the expected quality of an adopted project. Therefore, to summarize this discussion, the sample choice enters the payoff of the agent only through (i) its informativeness for the principal, i.e. $\sigma_{\mathscr{P}}(\mathbf{s})$, and (ii) the variation in the agent's posterior expected value that is explained by the principal's posterior, i.e. $\rho(\mathbf{s})\sigma_{\mathscr{A}}(\mathbf{s})$.

The sampling of $\mathbf{s}$ leads to a probability of adoption by the principal given by

$$\Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right).$$

A more informative sample for the principal – captured by a higher $\sigma_{\mathscr{P}}(\mathbf{s})$ – leads to a lower adoption probability if the principal finds the project initially promising (i.e. $v_0^{\mathscr{P}} > 0$) and to a higher adoption probability if she finds it unpromising (i.e. $v_0^{\mathscr{P}} < 0$). Hence, providing more information to the principal makes her ex-ante preferred decision less likely. As it will be argued in this section, the agent compares the principal's adoption probability to the following adoption probability:

$$\Phi\left(\frac{v_0^{\mathscr{A}}}{\rho(\mathbf{s})\sigma_{\mathscr{A}}(\mathbf{s})}\right),$$

rather than to the probability he would adopt in his single-agent problem, namely $\Phi\left(\frac{v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}(\mathbf{s})}\right)$.

Let us interpret this expression. Suppose that the agent does not observe the particular attribute realizations within the sample, but only observes the realized posterior of the principal $v^{\mathscr{P}}(\mathbf{s})$. What is the predicted value of $v^{\mathscr{A}}(\mathbf{s})$ based on $v^{\mathscr{P}}(\mathbf{s})$? By the joint normality, the predicted value – let that be denoted by $\hat{v}^{\mathscr{A}}(\mathbf{s})$ – is

$$\hat{v}^{\mathscr{A}}(\mathbf{s}) = \frac{\sigma_{\mathscr{A}}(\mathbf{s})\rho(\mathbf{s})}{\sigma_{\mathscr{P}}(\mathbf{s})}\left(v^{\mathscr{P}}(\mathbf{s}) - v_0^{\mathscr{P}}\right) + v_0^{\mathscr{A}}.$$

Before $v^{\mathscr{P}}(\mathbf{s})$ is realized, the distribution of $\hat{v}^{\mathscr{A}}(\mathbf{s})$ is given by $\hat{v}^{\mathscr{A}}(\mathbf{s}) \sim \mathcal{N}\left(v_0^{\mathscr{A}}, \rho^2(\mathbf{s})\sigma^2_{\mathscr{A}}(\mathbf{s})\right)$. The agent would adopt if $\hat{v}^{\mathscr{A}}(\mathbf{s}) \geq 0$, hence $\Phi\left(\frac{v_0^{\mathscr{A}}}{\rho(\mathbf{s})\sigma_{\mathscr{A}}(\mathbf{s})}\right)$ gives the probability of an adoption by the agent if he were constrained to decide only based on the observed posterior of the principal. To summarize, rather than simply comparing the adoption rates of the respective single-player problems, i.e. $\Phi\left(\frac{v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}(\mathbf{s})}\right)$ and $\Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right)$, the agent compares $\Phi\left(\frac{v_0^{\mathscr{A}}}{\rho(\mathbf{s})\sigma_{\mathscr{A}}(\mathbf{s})}\right)$ and $\Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right)$ in order to account for the informational overlap among the players, captured by ρ.

Proposition 4.2 examines the effect of the initial realization B_0 on the optimal sample choice. With $\mu = 0$, the two players fully agree on the ex-ante expected value of the project. If $B_0 > 0$, the principal initially prefers adoption, which guarantees both players a payoff of B_0. If instead the expected value is $-B_0 < 0$, both players obtain a zero payoff from the principal's rejection. The added value of sampling in the two cases is equal: a realization $B_0 > 0$ is as suggestive that adoption is desirable as $(-B_0)$ is suggestive that rejection is desirable. Hence, optimal sampling depends on the absolute value of B_0 rather than its sign.

If the drift is nonzero, the optimal samples are generically different for B_0 and $(-B_0)$. A project with B_0 need not be as promising as a project with $(-B_0)$ is unpromising. Both observations might suggest that adoption is desirable, or the former might be highly suggestive of an adoption while the latter only slightly suggestive of a rejection. Yet, when comparing across environments, an initial observation B_0 coupled with drift μ induces the same optimal sampling as an initial observation $(-B_0)$ with drift $(-\mu)$. Both players are as optimistic in one case as they are pessimistic in the other.

Proposition 4.2 (Conclusiveness of initial evidence)**.** *For any $(B_0, \mu) \in \mathbb{R}^2$, if a sample is optimal for (B_0, μ), it is also optimal for $(-B_0, -\mu)$.*

4.2 Minimal capacity: $q = 1$

Here we analyze optimal attribute choice when the agent is limited to $q = 1$. A minimal capacity captures instances in which the agent is severely time-constrained or materially constrained in researching the uncertain project: at most one attribute of the project can be fully learned before a decision is made. Alternatively, this is the problem of an agent who persuades a principal with very limited attention and is constrained to disclose to this principal all the realizations that he gathers.

The first part (subsection 4.2.1) considers optimal sampling *conditional* on the agent sampling an attribute. The discussion in this subsection builds up towards two key results:

propositions 4.3 (the dependence of compromise sampling on the form of ex-ante disagreement) and 4.4 (sufficient conditions for non-compromise sampling). The issue of whether the agent prefers to sample this attribute or forgo sampling at all is treated separately in subsection 4.2.2.

4.2.1 Optimality of compromise sampling

We seek to understand how the optimal attribute in the principal-agent game compares to the optimal attributes in the respective single-player problems. For such a comparison, it is useful to impose the condition that each player has a single local optimum on either side of a_0. So, we assume that for each player, the variance $\sigma_i^2(a_0, a) = |a - a_0|\tau_i^2(a;(a_0, a))$ and $\sigma_i^2(a, a_0)$ are single-peaked in a in $[a_0, 1]$ and $[0, a_0]$ respectively.[20] This condition is not particularly restrictive, and it is closely connected to assumption 2.

Assumption 3 (Single-peakedness of σ_i^2)**.** *For each player, a_0 and ω_i need to be such that the functions $\ell_i(\cdot)$ and $r_i(\cdot)$ as defined in assumption 2 with domains $[0, a_0]$ and $[a_0, 1]$ respectively:*

$$r_i(a) = \int_a^1 \omega_i(s)ds - \int_{a_0}^a \frac{s - a_0}{a - a_0}\omega_i(s)ds$$
$$\ell_i(a) = \int_a^{a_0} \frac{a_0 - s}{a_0 - a}\omega_i(s)ds - \int_0^a \omega_i(s)ds$$

cross zero only once from above on $[a_0, 1]$ and $[0, a_0]$ respectively.

Hereafter, we refer to the local optima of player i, i.e. the attributes chosen in the single-player problem by i when sampling is constrained to $[0, a_0]$ and $[a_0, 1]$, as (a_i^ℓ, a_i^r) respectively, and to her global optimal attribute as a_i.

An attribute is a local compromise in $[0, a_0]$ if it is between the players' local optima in this interval; a local compromise in $[a_0, 1]$ is defined analogously. Attributes that are further away from (resp., closer to) a_0 than both players' local optima are referred to as locally overshooting (resp., undershooting) attributes. We occasionally refer to overshooting and undershooting with the common term of *no-compromise* sampling. Overshooting attributes are more uncertain than what both players prefer, while undershooting attributes are less uncertain.

Definition 1 (Local compromise)**.** *An attribute $\tilde{a}$ is a **local compromise** in $[a_0, 1]$ if it is strictly between the single-player local optima, i.e.,*

$$\min\{a_{\mathscr{P}}^r, a_{\mathscr{A}}^r\} \le \tilde{a} \le \max\{a_{\mathscr{P}}^r, a_{\mathscr{A}}^r\}. \tag{10}$$

*An attribute $\tilde{a}$ **undershoots** in $[a_0, 1]$ if $\tilde{a} < \min\{a_{\mathscr{P}}^r, a_{\mathscr{A}}^r\}$ and **overshoots** in $[a_0, 1]$ if $\tilde{a} > \max\{a_{\mathscr{P}}^r, a_{\mathscr{A}}^r\}$.*

Figure 5a illustrates the definition for attributes in $[a_0, 1]$, whereas figure 5b shows the regions of local compromises, overshooting, and undershooting for the entire attribute domain. Note that by the single-peakedness assumption, at any compromise attribute one

[20] Only the proof of lemma 4.5 requires strict concavity of σ_i^2 rather than single-peakedness.

player prefers sampling further to the right but the other player prefers it further to the left if each were acting in isolation of the other (i.e. σ_i^2 is decreasing whereas σ_{-i}^2 is increasing). In contrast, at any attribute that is not a compromise, the players agree whether sampling further in one direction leads to more informative sampling or not. E.g., in figure 5a, moving further to the right of a_{-i}^r leads to less informative sampling about the posterior expected value of both the principal and the agent.

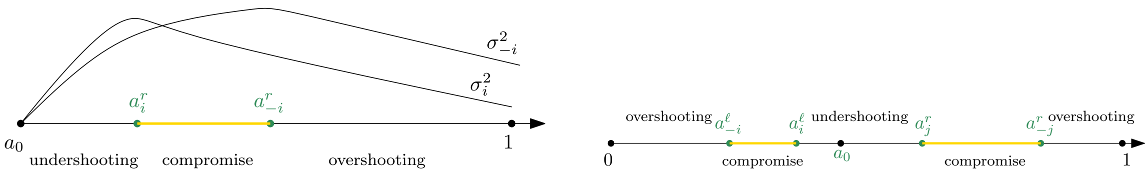

(a) Local compromise in $[a_0, 1]$ (b) Regions of local compromise in $[0, 1]$

Figure 5

Note, also, that definition 1 invokes only local optima of the single-player problems. This highlights the conflict among players about whether a local modification of the attribute would be desirable in their single-player problems. For a stark illustration of this definition, consider a setting in which (i) the two players have identical local optima: $a_{\mathscr{A}}^r = a_{\mathscr{P}}^r = a^r$ and $a_{\mathscr{A}}^\ell = a_{\mathscr{P}}^\ell = a^\ell$, but (ii) the global optimum for the agent is a^ℓ, whereas for the principal it is a^r.According to definition 1 there are no local compromises in this setting. If we were to instead define a compromise as any attribute between the global optima, i.e., any $a \in [a^\ell, a^r]$, at any such attribute both single-player objective would increase from moving closer to a_0. No player would have to make a concession from such a local modification.

For the rest of the discussion, let $v^i(a)$ denote the posterior expected value of player i after a is sampled, and let $\sigma_i^2(a)$ denote the variance of this posterior induced by sampling a. Under $q = 1$, for any $a \in [0, 1]$ the two posterior expected values $v^{\mathscr{P}}(a)$ and $v^{\mathscr{A}}(a)$ are perfectly correlated. Put differently, an outside observer that observes $v^{\mathscr{P}}(a)$ and B_0 but not $v^{\mathscr{A}}(a)$ or $B(a)$ can fully infer $v^{\mathscr{A}}(a)$.

Claim 1 (Perfect correlation of posteriors)**.** *Let $q = 1$. Then $\rho(\mathbf{s}) = 1$ for any $\mathbf{s} \in \mathscr{S}_2(a_0)$.*

From the agent's perspective, each attribute is fully described by the pair $(\sigma^{\mathscr{P}}(a), \sigma^{\mathscr{A}}(a))$. The agent's objective (9) from sampling attribute a simplifies to:

$$P(a; v_0^{\mathscr{P}}, v_0^{\mathscr{A}}) = v_0^{\mathscr{A}} \Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a)}\right) + \sigma_{\mathscr{A}}(a) \phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a)}\right). \tag{11}$$

If both sampling and adoption authority were perfomed by player i, the ex-ante expected value v_0^i would affect the adoption decision but not sampling. However, under separate authorities, the agent uses $(v_0^{\mathscr{P}}, v_0^{\mathscr{A}})$ to decide on optimal sampling as well. The ex-ante expected values are of persuasive value to the agent: they guide him in how to best persuade the principal to take an adoption decision well-aligned with his interest. Even in the case of prior agreement, i.e. $v_0^{\mathscr{P}} = v_0^{\mathscr{A}}$, the common ex-ante expected value informs the sampling choice.

So it is not merely the ratio of these two values, – a measure of the prior disagreement between the two players, –but also the magnitudes of the ex-ante expected values that matters for optimal sampling.

We first examine optimal sampling when $v_0^i = 0$ for some i. Player i is ex-ante indifferent between adoption and rejection: the initial evidence looks neutral from her perspective. If this indifferent player is the principal, the optimal attribute chosen by the agent is the agent's ideal attribute, irrespective of the agent's belief about the project. That is, the agent discovers the attribute she would discover in the absence of the principal as well. If $v_0^{\mathscr{P}} = 0$, the principal is expected to adopt half of all projects despite which attribute is sampled. Confronted with the inability to affect the principal's decision through attribute choice, the agent chooses the attribute that is most informative to him. The expected value of an adopted project for the agent is highest when the attribute with the highest $\sigma^2_{\mathscr{A}}$ is sampled.

Lemma 4.1 (Neutral initial evidence)**.**

(i) If $v_0^{\mathscr{P}} = 0$, the optimal attribute is equal to the agent's single-player optimum $a_{\mathscr{A}}$.

(ii) If $v_0^{\mathscr{A}} = 0$, the optimal attribute is a strict compromise. It tends to a local optimum of the principal as $v_0^{\mathscr{P}}$ approaches $\pm\infty$. The payoff of the agent strictly decreases in $|v_0^{\mathscr{P}}|$.

(iii) The agent weakly prefers ($v_0^{\mathscr{P}} = 0$ and $v_0^{\mathscr{A}} = x$) to ($v_0^{\mathscr{A}} = 0$ and $v_0^{\mathscr{P}} = x$) if $x > 0$. That is, given that one player perceives the initial evidence as neutral and the other as good news, the agent is better off when the principal perceives the initial evidence as neutral rather than when he himself does.

If the agent is ex-ante indifferent between adoption and rejection but the principal has a strict preference between the two, the optimal attribute is strictly between the two local optima, i.e. it is a strict compromise. When the initial evidence is very conclusive for the principal, an indifferent agent samples an attribute arbitrarily close to one of the two local optima of the principal. The conflict over adoption is sharp, as the principal either adopts or rejects almost surely. The agent resorts to an attribute that the principal finds sufficiently informative in hope of ameliorating this conflict.

A natural next question is whether, when all else is kept constant, the agent prefers to face an indifferent principal or be indifferent herself. The last part of proposition 4.1 resolves this question in favor of an indifferent principal when the agent herself is optimistic. This is due to her optimism about the project and the benefit she obtains from having the flexibility to choose her ideal attribute. This need not be the case when one of the players is pessimistic and the other is indifferent. If the degree of pessimism is very small or very large, the agent prefers to be herself the pessimistic player. The consideration that the ex-ante probability of approval of an indifferent principal cannot be affected by the sample choice trumps the disutility the agent obtains from a compromise attribute if the principal is pessimistic. For a moderate degree of pessimism, the agent still prefers the principal to be the indifferent player.

The next result establishes conditions for the optimal attribute to be a compromise. If the ex-ante expected values lead the two players to disagree on the adoption decision before sampling, the optimal attribute is always a strict compromise. When the players disagree initially, the agent unequivocally prefers a more informative attribute for the principal, all else kept constant. As long as the agent can increase both $\sigma_{\mathscr{A}}$ and $\sigma_{\mathscr{P}}$, she will do so. Hence, the optimal attribute is a compromise: the agent cannot alter it without decreasing the variance of some player and increasing the variance of the other. Therefore, compromise is an unavoidable consequence of the initial disagreement on approval. Alternatively, for the optimal attribute to undershoot or overshoot the single-player optima, it is necessary that the players agree on what the right decision is based on the initial evidence.

Yet, compromise can arise when the players initially agree on the adoption decision as well. The second part of proposition 4.3 identifies a condition for the optimal attribute to be a compromise. Essentially, the condition requires that at any optimal attribute that is a compromise the principal reacts more strongly than the agent: if the two players agree ex-ante on adoption, the principal adopts a larger expected share of projects than the agent, whereas if both agree on rejection, the principal adopts a smaller share than the agent.

Proposition 4.3 (Form of disagreement and compromise)**.**

(i) Under drastic disagreement ($v_0^i < 0 < v_0^{-i}$) any optimal attribute is a strict compromise.

(ii) Under mild disagreement ($sgn(v_0^{\mathscr{P}}) = sgn(v_0^{\mathscr{A}})$), an interior optimal attribute a^* is a compromise if and only if

$$\left|\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a^*)}\right| > \left|\frac{v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}(a^*)}\right|. \tag{12}$$

A natural implication of proposition 4.3 is that undershooting or overshooting optimally arises only if the initial evidence induces some minimal agreement among players, that is, only if the principal and the agent at least agree on whether the project is promising. Initial agreement among players about the right decision might bring about the sampling of an attribute that is more variable than both individual optima, or less variable than both individual optima. In the first case, sampling entails excessive variability, while in the second, sampling is too conservative compared to the single-player benchmarks.

Supposing that the two players agree initially on whether the project is promising, under what conditions is it that the optimal attribute is a compromise? The following result establishes a sufficient condition that follows immediately from part (ii) of proposition 4.3. The optimal attribute is guaranteed to be a compromise if the discovery of any unsampled attributes leads the principal to react more strongly than the agent. Lemma 4.2 simplifies condition (12) by using the fact that $\sigma_i(a) = |a - a_0|\tau_i(a)$, where $\tau_i(a)$ is the coefficient corresponding to the sampled realization $B(a)$ in the posterior expected value of player i.[21] Compromise sampling is guaranteed to exist if, for any attribute other than a_0, the principal starts

[21] The notation surpresses the dependence of τ on the entire sample, which is of the form (a_0, a) or (a, a_0).

with a sufficiently stronger ex-ante belief: she is sufficiently more optimistic (pessimistic) than the agent if the project is initially promising (unpromising) to the two players.

Lemma 4.2. *Suppose that there is mild disagreement between the players.*

(i) Compromise sampling emerges if for any $a \neq a_0$,

$$\frac{v_0^{\mathscr{P}}}{v_0^{\mathscr{A}}} > \frac{\tau_{\mathscr{P}}(a)}{\tau_{\mathscr{A}}(a)}. \tag{13}$$

(ii) Non-compromise (undershooting or overshooting) sampling emerges if the optimal attribute is interior and for any $a \neq a_0$,

$$\frac{v_0^{\mathscr{P}}}{v_0^{\mathscr{A}}} < \frac{\tau_{\mathscr{P}}(a)}{\tau_{\mathscr{A}}(a)}. \tag{14}$$

Lemma 4.2 is instrumental in pinning down sufficient conditions for the optimal attribute to undershoot or overshoot the single-player optima. Let us summarize the conditions in terms of the initial observation B_0 for any given pair of weight functions $(\omega_{\mathscr{P}}, \omega_{\mathscr{A}})$. Remember that the ex-ante value of the project for player i is

$$v_0^i = B_0 + \mu \underbrace{\int_0^1 \omega_i(a)(a-a_0)\,da}_{:=K_i}. \tag{15}$$

Player i is indifferent between adoption and rejection if $B_0 = -\mu K_i$. Hence, the two players disagree on the adoption decision prior to any sampling if B_0 is between $-\mu K_{\mathscr{P}}$ and $-\mu K_{\mathscr{A}}$.[22] By part (i) of proposition 4.3 any such intermediate B_0 leads to the sampling of a compromise. In the regions outside this interval, the two players agree on whether the initial evidence suggests adoption. Lemma 4.2 implies that a compromise is guaranteed in a small neighborhood around $B_0 = -\mu K_{\mathscr{A}}$: in such a neighborhood, the ratio $v_0^{\mathscr{P}}/v_0^{\mathscr{A}}$ is very high because $v_0^{\mathscr{A}}$ is close to zero. The opposite is true in a neighborhood around $B_0 = -\mu K_{\mathscr{P}}$ for which the players ex-ante agree on rejection. In this neighborhood, the ratio is very close to zero: the agent has much stronger views than the principal. Hence, any emerging interior optimal attribute either undershoots or overshoots the individual optima. There is a non-degenerate interval of values of B_0 for which it is optimal to sample a non-compromise attribute (i.e. undershooting or overshooting). Figure 6 summarizes this discussion.

Why is the agent willing to offer non-compromise sampling when the two players are in mild disagreement? In figure 6, consider a B_0 in the green region. Initial evidence suggests to both players that the project should be rejected; yet the principal is only barely convinced that B_0 is bad news, while the agent is more strongly convinced. Hence, the agent seeks to induce the principal to react more strongly in rejecting projects. She does so by decreasing

[22] If $-\mu K_{\mathscr{A}} > -\mu K_{\mathscr{P}}$, the agent is more demanding than the principal because she requests a higher B_0 in order to be indifferent between adoption and rejection.

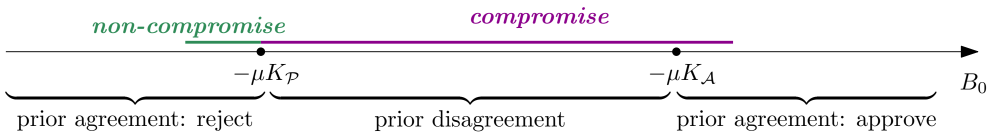

Figure 6: The figure assumes that the agent is ex-ante more demanding than the principal, i.e. $-\mu K_{\mathscr{A}} > -\mu K_{\mathscr{P}}$.

the informativeness of the sampled attribute for the principal. Remember that for $v_0^{\mathscr{P}} = 0$, the optimal attribute is the agent's ideal attribute; so for $v_0^{\mathscr{P}}$ in a sufficiently small neighborhood around zero, the optimal attribute moves further away from both local optima. The expected share of approved projects $\Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a)}\right)$ goes down as a result. For sure, choosing a lower $\sigma_{\mathscr{P}}$ comes at the cost of a lower $\sigma_{\mathscr{A}}$ for the agent as well. Yet, the agent is willing to bear this cost for a sufficiently uncertain principal, i.e. sufficiently small $|v_0^{\mathscr{P}}|$.

The following result provides sufficient conditions for the optimal attribute to overshoot or undershoot the agent's ideal attribute and the principal's local optimum in that corresponding interval. The optimal attribute is more variable than these two single-player optima if (i) B_0 induces the players to agree on the promise, but leaves the principal sufficiently skeptic, and (ii) $a_{\mathscr{A}}$, which is the optimal attribute when the principal is exactly indifferent, is further away from a_0 than the principal's single-player optimum in the same region. The latter condition is akin to the agent seeking more breadth of sampling than the principal in that side of a_0. She is willing to expand sampling even further than what is ideal to her when facing a principal who is sympathetic to her interpretation of the promise, but nonetheless too skeptical.

Proposition 4.4 (Sufficient conditions for overshooting/undershooting)**.** *The optimal attribute overshoots (resp., undershoots) the two single-player local optima if the following conditions hold simultaneously:*

(a) players have mild disagreement over the project: $sgn(v_0^{\mathscr{P}}) = sgn(v_0^{\mathscr{A}})$;

(b) $v_0^{\mathscr{P}}$ is sufficiently close to zero;

(c) $a_{\mathscr{A}}$ is more distant (resp., less distant) to a_0 than the principal's local optimum on the same side of a_0.

The rest of this discussion uses two examples to illustrate our results so far and other interesting distortions that might arise in optimal sampling. It moreover illustrates that the conditions of proposition 4.4 are not necessary: non-compromise sampling arises even in the case in which $v_0^{\mathscr{A}} = v_0^{\mathscr{P}}$. We elaborate on two classes of environments: (a) linear ω_i for $i = \mathscr{P}, \mathscr{A}$, and (b) a narrow-interest player whose ω_i is centered around a single attribute, interacting with a broad-interest player whose $\omega_j(a) = 1$ for all a. Both examples are in the context of $\mu = 0$ and $\sigma = 1$: the players are initially on full agreement about the promise, that is, $v_0^{\mathscr{P}} = v_0^{\mathscr{A}} = v_0$. More importantly, both examples highlight instances in which: 1) the optimal attribute is more variable (excessive sampling) or less variable (conservative sampling) than both single-agent optima, 2) optimal sampling occurs in the cell of least interest to both players, 3) more conclusive initial promise encourages sampling increasingly close to one of $\mathscr{P}$'s local optima.

Example 2 (Linear weight functions)**.** *Let $\mu = 0$. Suppose linear weight functions $\omega_i(a) = b_i + k_i a$ such that $b_i > 0$ and $b_i + k_i > 0$ for $i = \mathscr{P}, \mathscr{A}$. The two players start with the same ex-ante expected value for the project v_0.*

1. ***Clear intervals of compromise and non-compromise**: For any initial attribute a_0, it can be shown that the principal is more sensitive than the agent (i.e. $\tau_{\mathscr{P}}(a) > \tau_{\mathscr{A}}(a)$) to realizations of attributes in one side of a_0, and the agent is more sensitive than the principal to realizations of attributes in the other side of a_0.*[23] *That is, $\tau^i(a;(a,a_0)) > \tau^{-i}(a;(a,a_0))$ for $a \in [0,a_0]$ iff*
$$\frac{k_i}{b_i} < \frac{k_{-i}}{b_{-i}}.$$
The more sensitive player on $[0,a_0]$ is the one for which the ratio of weights of extreme attributes $\omega_i(1)/\omega_i(0)$ is the smallest: this player cares more about lower-index attributes. For the rest of this example, let the agent (resp., the principal) be the more sensitive player on $[0,a_0]$ (resp., $[a_0,1]$).[24] *By lemma 4.2, the optimum in $[0,a_0]$ is a compromise and the optimum in $[a_0,1]$ either undershoots or overshoots. For this reason, we refer to $[0,a_0]$ and $[a_0,1]$ as compromise and non-compromise regions respectively.*

2. ***Ranking of local optima**: The local optima of the two players can be ranked easily. In the compromise region $[0,a_0]$, $a^\ell_{\mathscr{A}} < a^\ell_{\mathscr{P}}$. The opposite holds in the region of non-compromise $[a_0,1]$: i.e., $a^r_{\mathscr{A}} < a^r_{\mathscr{P}}$, hence the local optimum of the agent is closer to a_0 in this region.*

3. ***Sampling in the compromise interval**: For any such two linear weight functions $(\omega_{\mathscr{P}}, \omega_{\mathscr{A}})$, the local optimum a^*_ℓ monotonically increases from $a^{\mathscr{A}}_\ell$ to $a^{\mathscr{P}}_\ell$ as the initial promise becomes more conclusive (that is, as $|v_0|$ increases). Although the added value of sampling an attribute is smaller as $|v_0|$ increases, –in fact, $\mathscr{A}$'s payoff approaches v_0 as v_0 becomes more increasingly conclusive, – the local optimal attribute in the compromise region remains in the interiority of $[a^{\mathscr{A}}_\ell, a^{\mathscr{P}}_\ell]$.*

4. ***Undershooting under the same ex-ante expected value** v_0: Note that under initial agreement, i.e. $v^i_0 = v_0$ for $i = \mathscr{A}, \mathscr{P}$, equation (13) simplifies to*
$$1 > \frac{\tau_{\mathscr{P}}(a)}{\tau_{\mathscr{A}}(a)},$$
which is indeed the case for any $a \in [a_0,1]$; hence, non-compromise sampling arises in this interval. In fact, only undershooting can arise; this occurs when $a^r_{\mathscr{A}}$ is the (global)

[23] All proofs for results presented in this example are gathered in appendix C.2.

[24] This can also be interpreted in terms of the adoption threshold that the agent follows $\bar{B}(a)$ (as defined in section 3). For any $a \in [0,a_0]$, the agent is more demanding than the principal towards unpromising projects and less demanding than him towards promising projects. Roles are reversed for $a \in [a_0,1]$. For instance, if $v_0 < 0$ and either player could veto the final approval decision, the agent would be the effective decisionmaker for $a \in [0,a_0]$ and the principal for $a \in [a_0,1]$.

single-agent optimum for the agent. From lemma 4.1, the optimum is $a^r_{\mathscr{A}}$ for $v_0 = 0$. As v_0 becomes more conclusive, the local optimal attribute decreases further away from $a^{\mathscr{A}}_r$ and towards a_0. More conclusive projects promote more conservative sampling in this region. Yet, there exists a critical value, such that for $|v_0|$ greater than this critical value the local optimal attribute jumps to a_0. Sampling switches to the compromise region.

Example 3 (Narrow interests, broad interests)**.** *Suppose an agent with narrow interests centered around $a = 1/3$ interacts with a principal that deems all attributes to be of equal importance. The weight functions are depicted and described in figure 7. The initially known attribute is $a_0 = 1/2$. Let $\mu = 0$.*

1. ***Optimal sampling is a compromise for any** v_0: The single-player optima are $a^{\mathscr{P}} \in \{1/6, 5/6\}$, and $a^{\mathscr{A}} = 0.280488$. The principal is indifferent between exploring the two sides of a_0, given that $a_0 = 1/2$ and her weight function is symmetric around a_0. The agent strictly prefers to sample in $[0, 1/2)$, as the attributes of greatest interest to her are in this region. Similarly to example 2, a_0 divides the attribute space into a compromise region $[0, a_0]$ and a non-compromise region $[a_0, 1]$. The agent samples to the left of a_0 for any initial promise v_0, which is (weakly) the area of most interest to both players. The optimal attribute is a compromise for any initial promise: it monotonically decreases from $a^{\mathscr{A}}$ to $a^{\mathscr{P}} = 1/6$ as the initial promise becomes more conclusive.*

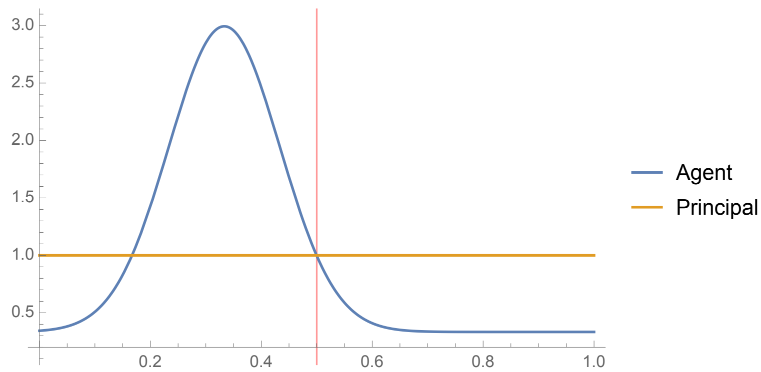

Figure 7: The weight function for $\mathscr{P}$ (in blue) is $\omega_{\mathscr{A}}(a) = \frac{1}{c}\left(\frac{1}{2} + \frac{1}{10}\phi\left(\frac{a-1/3}{1/10}\right)\right)$ where ϕ is the standard normal pdf and c is a normalizing constant. The weight function for $\mathscr{P}$ (in yellow) is $\omega_{\mathscr{P}}(a) = 1$. The vertical line denotes the initial attribute $a_0 = 1/2$.

2. ***Overshooting**: For the rest of this example suppose that the roles get reversed: the agent is the player with broad interests, and the principal is the one primarily interested in attributes close to $1/3$. Consider a moderately promising project: $v_0 = 1/3$. The optimal attribute to be discovered is $a^* = 0.0847334 < \min\{a^\ell_{\mathscr{A}}, a^\ell_{\mathscr{P}}\}$. The optimal attribute is more distant from a_0, and hence more variable, than both individually preferred attributes. The agent has a stronger preference for breadth of sampling, hence she is willing to expand sampling to more distant attributes.*

3. ***Sampling in the least interesting region for both players**: For $a_0 = 1/2 + 1/100$, both players' single-player optima are in the region $[0, a_0]$. Yet, for sufficiently large $|v_0|$, optimal*

attribute switches to the least interesting region $[a_0, 1]$.

4. ***Switching sampling regions**: The switching of sampling regions can happen more than once as $|v_0|$ increases. For $a_0 = 1/2 - 1/100$, the principal prefers to sample on $[0, a_0]$ and the agent on $[a_0, 1]$. With $v_0 = 0$, agent samples in her preferred region; as $|v_0|$ increases, sampling switches to the principal's preferred region and it takes the form of overshooting. As $|v_0|$ goes up even more, sampling reverts back to the agent's preferred region (which is also the compromise region).*

Finally, we seek to understand the tendency of optimal sampling as B_0 becomes very conclusive. Let us restrict attention to pairs of weight functions for which all attributes in one side of a_0 are more informative for one player, and all attributes in the other side are more informative for the other player. The two examples we discussed obviously fit into this class of weights. As B_0 becomes sufficiently conclusive, $v_0^{\mathscr{P}}$ and $v_0^{\mathscr{A}}$ get sufficiently close. There exists a sufficiently conclusive B_0 beyond which the agent samples exclusively in the cell that is more informative for her: the optimal attribute in this case is a compromise arbitrarily close to the principal's preferred attribute in this cell.

Proposition 4.5 (Eventual compromise)**.** *Suppose that $\tau_{\mathscr{P}}(a) > \tau_{\mathscr{A}}(a)$ for $a \in [0, a_0)$ and $\tau_{\mathscr{P}}(a) < \tau_{\mathscr{A}}(a)$ for $a \in (a_0, 1]$.*

1. *There exists a sufficiently high $\bar{b}$ such that for sufficiently conclusive initial evidence $|B_0| > \bar{b}$ the optimal attribute is a compromise in $(a_0, 1]$.*
2. *The optimal attribute gets arbitrarily close to $a^r_{\mathscr{P}} \in (a_0, 1]$ for an arbitrarily conclusive B_0.*

In particular, this result highlights that optimal sampling settles on a compromise for sufficiently conclusive initial evidence. This is the case even if both players are more interested in sampling in $[0, a_0]$ in the absence of the other player.

4.2.2 Optimal restriction of information

We next turn to whether the agent ever finds it beneficial to supply less information to the principal by not exhausting the capacity q. The following result shows that this is never the case under a minimal capacity $q = 1$ and $\mu = 0$. For a driftless process –for which the two ex-ante expected values are equal,– there is always some attribute that the agent strictly prefers to sample. An immediate implication of this statement is that for any given capacity q, $\mathscr{A}$ always mandates some sampling as long as the process is driftless. Under full agreement on the ex-ante value of the project, the agent imposes some minimal sampling standards.

Assumption 4 (Divisive initial attribute)**.** *The initial attribute a_0 is such that*

$$\int_0^{a_0} \omega_{\mathscr{P}}(a)\,da \neq \int_0^{a_0} \omega_{\mathscr{A}}(a)\,da.$$

Proposition 4.6 (Minimal provision of information)**.** *Suppose $\mu = 0$ and assumption 4 holds. For any (a_0, B_0), the agent strictly prefers sampling an additional attribute beyond a_0.*

We briefly explain the intuition behind the proof for proposition 4.6. For the agent to prefer not to sample for some initial promise B_0, it is necessary that all attributes be strictly less informative for the agent than for the principal. That is, at any sampled attribute the agent should prefer a lower variance $\sigma_{\mathscr{P}}$ for the principal. Assumption 4 guarantees that this cannot happen: there always exists an attribute close to a_0 that is more informative to the agent than to the principal.[25]

Next, we show that with different ex-ante expected values, the agent might find it beneficial to drastically curtail information. That is, the agent might forgo the opportunity to sample even when subject to a very limited capacity ($q = 1$). This is never the case if the two players disagree on the initial promise of the project, that is, if v_0^i and v_0^j have opposite signs. In such cases, it is always better for the agent to sample an attribute. So a necessary condition for the agent to prefer forgoing sampling is for the players to agree on whether the initial evidence suggests a promising project.

Example 4 (Optimal curtailment of information)**.** *Consider $\omega_{\mathscr{A}}(a) = \frac{2}{9}(4+a)$ and $\omega_{\mathscr{P}}(a) = \frac{3}{13}(1+10a^2)$. Let $\mu = 3, \sigma = 1, a_0 = 2/3$, and $B_0 = -1/10$. The players agree that the project is unpromising before any sampling:*

$$(v_0^{\mathscr{A}}, v_0^{\mathscr{P}}) = \left(-\frac{49}{90}, -\frac{3}{130}\right).$$

The agent is more convinced than the principal that the project is unpromising. The agent finds it beneficial to forgo the sampling of any new attributes. The payoff from sampling any attribute $a \neq 2/3$ grants her a strictly negative payoff, while in the absence of any further information, the principal rejects the project, granting to both himself and the agent a zero payoff. (If $\mu = 1$ instead of $\mu = 3$, while all else is kept constant, the agent prefers to sample a non-compromise rather than forgo discovery).

To summarize, if the initial evidence is such that (i) the two players agree about whether the project is promising, but (ii) the principal is sufficiently skeptic, i.e. $v_0^{\mathscr{P}}$ is sufficiently close to zero, the agent either forgoes all sampling, or resorts to one of the two forms of non-compromise sampling: excessive sampling (the optimal attribute overshoots both single-player optima), or conservative sampling (the optimal attribute undershoots both single-player optima). The agent discovers a non-compromise attribute if $v_0^{\mathscr{A}}$ is sufficiently inconclusive. If the initial evidence is sufficiently conclusive for her but the principal is very skeptic, she prefers to forgo all sampling. In the absence of further information, the principal has to make the decision that the agent strongly believes is the right decision in light of the initial evidence.

[25] It is straightforward to construct examples for which the agent refrains from sampling for sufficiently conclusive initial evidence when assumption 4 is violated. For instance, let $\omega_{\mathscr{A}}(a) = 1/10$, $\omega_{\mathscr{P}}(a) = (a-1/2)^2$ and $a_0 = 1/2$. For high enough $|v_0|$ (e.g. $v_0 = 3/2$), the optimal attribute is $a^* = a_0$.

4.3 Non-minimal capacities: $q > 1$

The problem of the agent becomes more complex for $q > 1$ as samples of greater size differ not only in the variabilities of the posterior expected values $v^{\mathscr{A}}(\mathbf{s})$ and $v^{\mathscr{P}}(\mathbf{s})$ that they induce, but also in the correlation $\rho(\mathbf{s})$. Let us briefly remark on how to interpret the correlation induced by a sample. The following lemma shows that a sample displays perfect correlation if for any two attributes in that sample, the relative importance of the two attributes is the same for both players.

Lemma 4.3 (Perfect correlation)**.** *Given $q \geq 2$, a sample $\mathbf{s} = (a^{\ell}_{n_\ell}, \ldots, a^{\ell}_1, a_0, a^r_1, \ldots, a^r_{n_r})$ displays perfect correlation if for any two attributes $a^k_i, a^{k'}_j$ in the sample:*

$$\frac{\tau_{\mathscr{P}}(a^k_i; (a^k_i, a^k_{i-1}))}{\tau_{\mathscr{P}}(a^{k'}_j; (a^{k'}_j, a^{k'}_{j-1}))} = \frac{\tau_{\mathscr{A}}(a^k_i; (a^k_i, a^k_{i-1}))}{\tau_{\mathscr{A}}(a^{k'}_j; (a^{k'}_j, a^{k'}_{j-1}))}. \tag{16}$$

Larger capacities introduce the possibility of $\rho(\mathbf{s}) \neq 1$. Yet, the next result argues that given any capacity, the optimal sample is straightforward to characterize if initial evidence is neutral for the principal ($v^{\mathscr{P}}_0 = 0$). Similarly to the analogous case of $q = 1$, the principal approves is expected to adopt half of the projects, and the agent cannot influence this adopted share through her sampling choice. Whereas for $q = 1$ the agent chose the sample with the highest variance for $v^{\mathscr{A}}(\mathbf{s})$, for $q > 1$ the agent chooses the sample with the highest explained variation $\rho^2(\mathbf{s})\sigma^2_{\mathscr{A}}(\mathbf{s})$. Unlike in the case of a minimal capacity, the agent might have to distort her most preferred sample if the correlation it induces is too low: the explained variation scales the informativeness of the sample for the agent's posterior expected value by the (squared) correlation that the sample induces. The optimal sample trades off distance from the agent's ideal sample for higher correlation.

Proposition 4.7 (Neutral evidence for the principal)**.** *Suppose $v^{\mathscr{P}}_0 = 0$. Given capacity $q > 1$, the optimal sample solves the following:*

$$\mathbf{s}^* \in \underset{\mathbf{s} \in \cup_{1 \leq k \leq q+1} \mathscr{S}_k(a_0)}{\arg\max} \quad \rho(\mathbf{s})\sigma_{\mathscr{A}}(\mathbf{s}). \tag{17}$$

The next simple example illustrates such deviations from the agent's sample preferred sample in the case of linear weights for both players and $q = 2$. Considerations about correlation might induce the agent to either bring both attributes closer to a_0 than her preferred sample, or to move them both further away from a_0, or to narrow the sampling scope by increasing one attribute and decreasing the other relative to the agent's ideal sample.

Example 5 (Linear interests with $q = 2$)**.** *Let $\omega_{\mathscr{P}}(a) = 3/2 - a$, $\mu = 0$, and $a_0 = 1/20$. We alter $\omega_{\mathscr{A}}$ and show three different patterns of the optimal sample as compared with $\mathscr{A}$'s ideal sample.*

1. $\omega_{\mathscr{A}}(a) = 1/4(3 + 2a)$*: the two attributes of the optimal sample are further away from a_0 than the attributes in $\mathscr{P}$'s ideal sample;*

2. $\omega_{\mathscr{A}}(a) = 2/5(3-a)$*: the two attributes of the optimal sample are closer together than the attributes in* $\mathscr{P}$*'s ideal sample. This means that one attribute is further away from* a_0 *and the other is closer to* a_0*;*

3. $\omega_{\mathscr{A}}(a) = 1/6(10-8a)$*: the two attributes of the optimal sample are both closer* a_0 *than the attributes in* $\mathscr{A}$*'s ideal sample.*

If the slope of $\omega_{\mathscr{A}}$ *is negative, so that the agent just like the principal has less of an interest in higher-indexed attributes, the optimal sample, which is located in* $[a_0, 1]$*, is either uniformly closer to* a_0 *or uniformly closer together. The opposite seems to be the case if* $\omega_{\mathscr{A}}$ *is positively sloped: the optimal sample is either uniformly further away from* a_0*, or the two attributes are closer together. With linear weights and* $v_0^{\mathscr{P}} = v_0^{\mathscr{A}} = 0$*, the optimal attributes are never farther apart from each other than in the agent's optimal sample.*

Proposition 4.7 is a first step in building an analogy between the case of $q = 1$ and the case of any greater capacity $q > 1$. Whereas with a minimal capacity, the agent compares his adoption probability with that of the principal, the discussion in subsection 4.1 explained that with $q > 1$ the agent compares

$$\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})} \quad \text{with} \quad \frac{v_0^{\mathscr{A}}}{\rho(\mathbf{s})\sigma_{\mathscr{A}}(\mathbf{s})}.$$

The first is the adoption probability of the principal after $\mathbf{s}$ is sampled, and the second is the adoption probability of the agent if he were constrained to decide on adoption after observing only the realized posterior of the principal. For instance, if both $v_0^{\mathscr{P}}$ and $v_0^{\mathscr{A}}$ are positive,

Does the agent ever optimally select an optimal sample that can be modified locally to generate both higher informativeness for the principal $\sigma^2_{\mathscr{P}}$ and higher explained variation $\rho^2\sigma^2_{\mathscr{A}}$? This is akin to undershooting/overshooting in sampling under a minimal capacity. The answer is positive, but such optimal sampling can only arise if there is mild disagreement between players. Based on an intuition similar to that for $q = 1$, such an optimal sample is guaranteed to arise if $v_0^{\mathscr{P}}$ is sufficiently close to zero compared to $v_0^{\mathscr{A}}$.

Proposition 4.8. *Given capacity* q*, let* $\mathbf{s}^* = (a_1^*, \ldots, a_n^*)$ *denote an optimal sample.*

1. *Under drastic disagreement, either* $\sigma^2_{\mathscr{P}}(\mathbf{s}^*)$ *is decreasing and* $\rho^2(\mathbf{s}^*)\sigma^2_{\mathscr{A}}(\mathbf{s}^*)$ *is increasing in* a_k^* *for any* $k \leq n$ *such that* $a_k^* \neq a_0$*, or vice versa.*

2. *Under mild disagreement, for any* $v_0^{\mathscr{A}}$ *there exists a sufficiently small* $|v_0^{\mathscr{P}}|$ *such that* $\sigma^2_{\mathscr{P}}(\mathbf{s}^*)$ *and* $\rho^2(\mathbf{s}^*)\sigma^2_{\mathscr{A}}(\mathbf{s}^*)$ *are both increasing in* $a_k^* \neq a_0$ *for any* $k \leq n$*.*

5 Extensions

5.1 Collective adoption

The main analysis studied the case of fully separate authorities on sampling and adoption. A natural variation on the configuration of authorities is the case in which only one player has the expertise or authority to sample, but both players decide collectively on the adoption decision. In particular, this subsection explores unanimous adoption: the project is adopted if both players are in favor of adoption after sampling takes place. Let us briefly revisit the example of an investigative journalist sampling a vast primary source, mentioned in the introduction. The editor, being unable to access the primary source himself, relies on the investigation techniques chosen by the journalist. The journalist, on the other hand, has a responsibility to report all the gathered hard evidence to the editor. Yet, when it comes to the decision of publishing a report on this primary source, both the editor and the journalist have to agree on whether publication is the right decision.

Upon sampling, each player favors adoption if and only if her posterior expected value is greater than zero. The agent seeks to maximize the expected value of the project conditional on unanimous adoption. The objective of the agent is:

$$\max_{\mathbf{s}\in\cup_{k=1}^{q+1}\mathscr{S}_k(a_0)} \quad \Pr(v^{\mathscr{P}}(\mathbf{s})\geq 0, v^{\mathscr{A}}(\mathbf{s})\geq 0)\mathbb{E}\left[v^{\mathscr{A}}(\mathbf{s}) \mid v^{\mathscr{P}}(\mathbf{s})\geq 0, v^{\mathscr{A}}(\mathbf{s})\geq 0\right] \tag{18}$$

Lemma 5.1. *The objective of the agent simplifies to:* [26]

$$\max_{\mathbf{s}\in\cup_{k=1}^{q+1}\mathscr{S}_k(a_0)} \quad v_0^{\mathscr{A}} BvN\left(c_{\mathscr{A}}, c_{\mathscr{P}}; \rho\right) + \sigma_{\mathscr{A}}\left(\rho\phi(c_{\mathscr{P}})\Phi\left(\frac{c_{\mathscr{A}} - c_{\mathscr{P}}\rho}{\sqrt{1-\rho^2}}\right) + \phi(c_{\mathscr{A}})\Phi\left(\frac{c_{\mathscr{P}} - \rho c_{\mathscr{A}}}{\sqrt{1-\rho^2}}\right)\right), \tag{19}$$

where

$$c_i(\mathbf{s}) := \frac{v_0^i}{\sigma_i(\mathbf{s})}.$$

When the principal is the only player with adoption authority, the agent prefers samples that induce higher correlation, because they are more likely to guarantee that an adoption by the principal (i.e., $v^{\mathscr{P}}(\mathbf{s})\geq 0$) is the right adoption decision for the agent too (i.e. $v^{\mathscr{A}}(\mathbf{s})\geq 0$). Higher correlation implies higher expected value for the agent conditional on adoption. When the adoption authority is shared, higher correlation implies higher probability of unanimous adoption and higher expected value of the project for the agent conditional on unanimous approval. Hence, all else constant, in both cases the agent prefers more highly correlated samples. But, in the case of collective adoption the effect is more pronounced due to correlation entering the adoption probability as well.

Corollary 5.2. *Suppose $q=1$. The expected payoff of the agent from sampling attribute a is:*

$$v_0^{\mathscr{A}}\Phi\left(\min\left\{\frac{v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}(a)}, \frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a)}\right\}\right) + \sigma_{\mathscr{A}}\phi\left(\min\left\{\frac{v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}(a)}, \frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a)}\right\}\right). \tag{20}$$

It is instructive to compare this objective to the agent's objective (11) when the principal is the only one with adoption authority. Under unanimous adoption, the decisive or pivotal "voter" for any sampled attribute a is the one with the smallest $v_0^i/\sigma_i(a)$. This is the player with a smaller probability of adoption of the project before $B(a)$ is realized: she is the more

[26] The argument $\mathbf{s}$ has been surpressed in $\rho\equiv\rho(\mathbf{s}), \sigma^i\equiv\sigma^i(\mathbf{s})$, $c_i\equiv c_i(\mathbf{s})$.

demanding decisionmaker towards realizations of attribute a. As a consequence, the problem of collective adoption divides attributes into two subsets: (i) attributes the sampling of which gives to the agent her single-player payoff, and (ii) attributes the sampling of which gives to the agent her decentralized evaluation payoff.

Proposition 5.1 (Agent's ideal sampling)**.** *If the agent's ideal sample $a_{\mathscr{A}}$ is such that the agent is the decisive player when $a_{\mathscr{A}}$ is sampled, i.e.,*

$$\frac{v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}((a_0, a_{\mathscr{A}}))} < \frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}((a_0, a_{\mathscr{A}}))},$$

the optimal attribute is $a^* = a_{\mathscr{A}}$.

We need to distinguish two important instances of the premise of proposition 5.1. First, the agent samples her ideal attribute if $v_0^{\mathscr{A}} < 0 < v_0^{\mathscr{P}}$. When the players drastically disagree on the promise of the project, and the agent believes that the project is unpromising, optimal sampling consists of $a_{\mathscr{A}}$. Despite the form the weight functions might take, an agent who considers the project unpromising samples her ideal attribute whenever the principal thinks the project is promising. Second, when the two players agree on the sign of the promise, $a_{\mathscr{A}}$ arises as optimal whenever the ratio of sensitivities to the realization $B(a_{\mathscr{A}})$, denoted by $\tau_{\mathscr{P}}(a_{\mathscr{A}};(a_0, a_{\mathscr{A}}))/\tau_{\mathscr{A}}(a_{\mathscr{A}};(a_0, a_{\mathscr{A}}))$, is (i) sufficiently small if the project seems promising to the players, (ii) sufficiently large if the project seems unpromising to the agents. Put differently, $a_{\mathscr{A}}$ is optimal if $v_0^{\mathscr{A}}$ is sufficiently smaller than $v_0^{\mathscr{P}}$. Under these conditions, the agent is the decisive decisionmaker.

The following result establishes that unlike in decentralized evaluation, $a_{\mathscr{A}}$ arises as an optimal solution for a non-degenerate interval of values for the initial observation B_0. Lemma 4.1 established that under decentralized evaluation $a_{\mathscr{A}}$ arises as optimal only for

$$B_0 = -\mu \int_0^1 (a - a_0)\omega_{\mathscr{P}}(a)\,da,$$

the value of B_0 for which the principal is indifferent.

Proposition 5.2. *For any pair of weights $(\omega_{\mathscr{A}}, \omega_{\mathscr{P}})$, there exists a non-degenerate interval of values of B_0 for which optimal sampling consists of the agent's ideal attribute $a_{\mathscr{A}}$.*

5.2 Preemptive sampling by an uncertain principal

This subsection relaxes the assumption that the principal does not possess any sampling authority in a simple, yet instructive, way. Players start by observing attribute a_0 with realization

$$B_0 = -\mu \int_0^1 \omega_{\mathscr{P}}(a)\,a\,da.$$

This observation leaves the principal indifferent towards adoption or rejection: he is highly uncertain of the quality of the project. Because of this initial condition, the process across attributes is a standard Brownian motion. Before the agent decides on sampling, the principal can ask her to include some attribute a_1 in the sample. This requirement to sample a_1 contextualizes how the agent views other attributes. Neither the agent nor the principal observes $B(a_1)$ before any additional sampling: the agent is aware of having to discover a_1, but she samples attributes simultaneously, as in the main setup.

By lemma 4.1, whenever the principal is highly uncertain, i.e. $v_0^{\mathscr{P}} = 0$, the agent samples a set of attributes that, given a_0 and q, solves

$$\max_{\mathbf{s}\in\cup_{1\le k\le q+1}\mathscr{S}_k(a_0)} \quad \rho(\mathbf{s})\sigma_{\mathscr{A}}(\mathbf{s}).$$

The principal's choice of a_0 shapes the feasible set of samples $\bigcup_{1\le k\le q+1}\mathscr{S}_{q+1}(a_0)$. Given a selection $\mathbf{s}^*(a_0)$ from the solution set of the agent's problem, the principal seeks to maximize

$$\max_{a_0} \quad \sigma_{\mathscr{P}}(\mathbf{s}^*(a_0)).$$

The principal chooses a_0 that encourages the discovery of a sample that induces the highest variability in her posterior $v^{\mathscr{P}}(\mathbf{s})$. Note that the sample chosen jointly is not informed by the promise of the project as perceived by the agent, i.e. $v_0^{\mathscr{A}}$.

We illustrate preemptive sampling with the following example of a broad-interest agent and an extreme-interest principal. Each player's sampling is restricted by a minimal quota $q = 1$.

Example 6 (Preemptive sampling)**.** *Let $\omega_{\mathscr{A}}(a) = 1$ and $\omega_{\mathscr{P}}(a) = 12(a-1/2)^2$. Given the choice of the principal a_1, the optimal choice of the agent is*

$$a_2^*(a_1) = \begin{cases} a_1 + 2/3(1-a_1) & \text{if } a_1 < 0.63 \\ a_1/2 & \text{if } a_1 > 0.63. \end{cases}$$

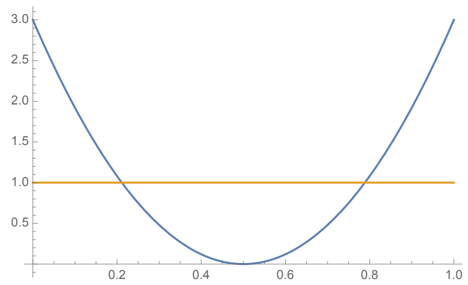

Figure 8: The agent cares equally about all attributes. The principal's interests are centered in the extremes of the attribute spectrum.

If the agent samples exclusively to the right of what the principal requests, the principal prefers to sample an attribute as high as 0.74*, so as to push agent's sampling to higher-indexed attributes. Yet, if the attribute requested by the principal is higher than* 0.63*, the agent switches to the exploration of lower-indexed attributes in* $[0, a_1]$*. If the agent samples exclusively in*

$[0, a_1]$, *the principal has an interior optimal attribute at* $a_1 = 0.91$. *In this way, the principal guarantees that sampling has shifted towards higher-indexed attributes while still being sufficiently informative of lower-indexed attributes. Thus, the principal compares sampling* $a_1 = 0.63$, *which encourages* $a_2 = 0.87$ *to sampling* $a_1 = 0.91$ *which induces* 0.451. *The latter is more informative to the principal. In the absence of preemptive sampling by the principal, the agent would choose* $a_1 = 2/3$ *if restricted to* $q = 1$, *and* $(a_1, a_2) = (0.4, 0.8)$ *if restricted to* $q = 2$. *Compared to the latter, sampling has shifted towards higher-indexed attributes due to preemptive sampling.*

6 Discussion: A finite-attribute model

This section briefly presents a finite-attribute model that exhibits common features with the model of section 2. While the model with a continuum of attributes is both more tractable and allows for a richer analysis, this alternative model is instructive in identifying what modeling choices yield similar results.

A project is described by four attributes $\{a_1, a_2, a_3, a_4\}$, each taking a binary realization $\theta_k \in \{b, g\}$. The attributes are correlated according to a binary Markov chain, with transition probability λ. That is, for any $k \geq 2$, λ denotes the persistence of the realization of an attribute to the next attribute:

$$\Pr(\theta_k = \theta_{k-1} | \theta_{k-1}) = \lambda \in (1/2, 1].$$

The value of the project to each player $i = 1, 2$ is a sum of the values from each individual attribute. The attribute-specific payoff to player i is given by:

$$v_i(a_k) = \begin{cases} 1 & \text{if } \theta_k = g \\ -\ell_k^i & \text{if } \theta_k = b, \end{cases}$$

where $\ell_k^i > 0$ for any i, k. An attribute with a greater loss ℓ_k^i is more important to the players. The value of the project to each player is assumed to be additive:

$$V_i = \sum_{k=1}^{4} v_i(a_k).$$

The players jointly consider whether to adopt the project. In the case of a rejection, the outside option is normalized to zero. Suppose the players perfectly observe θ_1, and the players can discover one additional attribute beyond this known attribute. Player 1 samples this additional attribute publicly, and player 2 decides whether to adopt the project. Three features of this simple sampling problem presented below echo key results in our main analysis in sections 3 and 4.

Irrelevance of θ_1 for optimal sampling: The optimal attribute is the same despite the realization of θ_1. First, only attributes the realization of which is consequential for the adoption decision are worth sampling: that is, an attribute a_k is sampled only if $\theta_k = g$ leads to an adoption and $\theta_k = b$ leads to a rejection. By the assumed symmetry of persistence and the law of iterated expectations, any attribute that maximizes $\Pr(\theta_k = g \mid \theta_1 = g)\mathbb{E}[V \mid \theta_1 = \theta_k = g]$ also maximizes

$$\Pr(\theta_k = g \mid \theta_1 = b)\mathbb{E}[V \mid \theta_1 = b, \theta_k = g].$$

Hence this attribute is weakly optimal for $\theta_1 = b$ as well.

Relative unimportance and interiority of optimal attribute: In the single-player problem in which the same player decides on both sampling and adoption, the optimal attribute is not necessarily the one with the highest loss. The player might decide to sample an attribute that is central rather than the one incurring the greatest loss upon a b realization. For instance, suppose $(\ell_2, \ell_3, \ell_4) = (2, 9/5, 2)$ and $\lambda = 2/3$.[27] The player optimally samples a_3, even though its loss is $9/5 < 2$. By sampling this middle attribute, the player learns more valuable information about a_2 and a_4 as well: as long as ℓ_3 is not too small relative to other attributes, she is willing to sample it for this reason.

Moreover, given θ_1, the most variable attribute is a_4. Indeed, the variance of θ_4 is equal to $\Pr(\theta_k = g \mid \theta_1)\Pr(\theta_k = b \mid \theta_1)$: due to the correlation structure, more distant attributes from a_1 are more variable. Yet, the player might not optimally sample the most variable attributes. Interior attributes might be more appealing due to their greater informativeness about sorrounding attributes. That is, interior attributes have stronger extrapolating power.

Overshooting by the optimal attribute: Suppose that player 1, who decides on sampling, has loss profile $(\ell_1^1, \ell_2^1, \ell_3^1, \ell_4^1) = (2, 2, 2, 2)$ while player 2, who is the adopter, has profile $(\ell_1^2, \ell_2^2, \ell_3^2, \ell_4^2) = (1, 4, 3/2, 7/2)$. Let $\lambda = 35/64$ and suppose that the initial evidence suggests a promising project, i.e. $\theta_1 = g$. Despite $\theta_1 = g$, in the absence of additional information both players reject the project because the respective expected payoffs from adoption are -0.35 and -1.76. Player 1, with uniform loss over attributes, ideally discovers a_3. While other attributes a_2 and a_4 are consequential to adoption as well, a_3 is the one yielding the most useful information for player 1. Player 2, on the other hand, prefers to discover a_2. Any realization of θ_3 would lead her to reject the project, because ℓ_2 and ℓ_4 are so much larger that even $\theta_3 = g$ is not sufficiently strong evidence to convince her to adopt the project. In contrast, her adoption decision would be responsive to θ_4, but sampling a_4 is not as valuable as sampling a_2. In the joint problem, from the perspective of player 1:

(i) sampling a_3 leads to a sure rejection by player 2, from which both players obtain a zero payoff;

(ii) sampling a_2 gives player 1 an expected payoff of 0.64, while sampling a_4 gives her an expected payoff of 0.65. Hence, she prefers to sample a_4. The optimal sampling choice overshoots the two single-player optimal attributes.[28]

[27] The value of ℓ_1 is immaterial for the sampling decision.

[28] If $\theta_1 = b$ instead, both players reject the project despite any realizations that additional sampling might reveal.

7 Related literature

The central problem of this paper tangentially relates to a number of distinct literatures: optimal discovery of multi-attribute objects, experimentation with potentially correlated alternatives with and without shared control, as well as persuasion through hard evidence.

First, the paper contributes to a recent but growing literature on the gradual discovery and adoption of multi-attribute objects, albeit interest on this issue has been longstanding in economics at least since **?**. **?** analyze optimal sequential and simultaneous discovery of a single object characterized by finitely many independent attributes. Although the present paper models a continuum of correlated attributes, the optimality criterion in its single-player benchmark is reminiscent of Proposition 1 in **?**. They show that under equal discovery cost, attributes that are dominated in the second-order stochastic dominance sense (SOSD) are discovered first. In the present setting, although attributes are not ordered according to SOSD, the distributions of the posterior expected values that they generate are. The player optimally samples attributes that are dominated in this sense.

Other existing work combines optimal search across several multi-attribute objects with optimal sampling of attributes within an object. All such work assumes independent attributes and a single decision-maker. **?** considers sequential search of ex-ante identical objects with limited sampling of a single attribute. **?** consider search among two-attribute objects by a mass of identical searchers who check at most one attribute. **?** identifies necessary conditions for optimal sequential search across a small number of objects and attributes. **?** compare two search methods: sampling of all attributes of a single object, and sampling of a single attribute across all objects. What distinguishes the present paper from this line of work is its focus on a single-attribute object, the presence of correlated attributes, and the separation of sampling and adoption authorities.

? embed attribute discovery into the bilateral trade of a multi-attribute good. The seller and the buyer disagree on the relative importance of attributes, but only the seller can sample attributes privately and make offers. The current paper also studies a setting in which only one party can attributes, but she does so publicly in order to influence the adoption decision. The case in **?** that is closer to my setup is that in which the sampled attributes are observable, but not their respective realizations (rather than both objects being unobservable).

My analysis makes use of the Brownian process construction first introduced in optimal search models by **?**. They use it to model the uncertain and complex nature of technological discovery, in order to analyze the link between inventive intensity and productivity growth. A sequence of myopic agents gradually learn about an infinite family of independent Brownian motions. **?** introduces the Brownian process in the strategic experimentation literature. It is the unknown mapping from policies to outcomes: a sequence of identical myopic players experiment with policies hoping to identify an outcome close to their ideal outcome. **?** study the optimal path of technological experimentation by a sequence of forward-looking agents, where possible technologies are correlated according to a single Brownian process. In these three papers, the players search over an unknown mapping to identify a point that maximizes their reward. In contrast to the present paper, the players' utilities do not depend on the entire realized mapping, but rather just on the sampled realizations. A conceptual precursor to these problems is **?**.

The present model shares with the literature on collective experimentation (**?**, **?**) the dependence of each player on the other regarding learning about an object of common interest and its eventual adoption. In **?**, voters decide collectively whether to continue experimenting with a project, but each voter learns about the project's value for her as long as the group experiments. **?** adopt the Brownian framework to study two-player sequential experimentation: each player experiments with a single policy but they both reap the payoff from the other's experimentation as well. In the current setup, the principal depends on the agent in order to learn about the project, and the agent depends on the principal for adoption.

? establish the possibility of the first player experimenting with a policy that overshoots both players' ideal policies. The reason for overshooting in their setup is different from ours: in their problem, overshooting preempts more distant experimentation by the second player and potentially causes her to settle for a tried compromise policy. In this paper, the agent uses overshooting so as to better align the principal's adoption standard with her own.

The paper is also related to strategic communication with verifiable information. **?** study cheap-talk communication about a two-dimensional state between an informed speaker and an uninformed listener. The latter perfectly verifies at most one dimension. The current setting also features a binary adoption decision, but the players are symmetrically informed and all accumulated information is verifiable by both. Moreover, the acquired information is endogeneous. **?** extends **?** to any number of verifications and dimensions. More recently, Carroll and Egorov (2017) study a similar problem, but their analysis focuses on the possibility of fully-revealing communication with general payoff functions, an arbitrary decision set for the listener, and multi-dimensional states.

Another related literature is that of Bayesian sequential testing (**?**). But unlike in Wald's framework in which the agent sequentially draws iid signals about an unknown payoff state, the agent in the current model samples attributes, the mean and variance of which varies. Moreover, current sampling affects the informativeness of remaining unsampled attributes, and therefore future sampling. Somewhat related, **?** study the problem of an agent who, besides deciding whether to continue sampling, can also vary the intensity of sampling by choosing the number of iid signals she draws at once.[29] **?** studied a similar problem, but with exogeneous intensity. In the management science literature, **?**, **?**, **?**, and **?** use sequential testing to study optimal technology adoption and search across different uncertain technologies. This paper also shares features with Wald persuasion games between an evaluator and an approver, as introduced in **?**. The two players have separate authorities, as in the current paper, but the sampling authority chooses the type of evidence to be disclosed to the approver through her choice of stopping time. In contrast, in my setup the sampling authority has to disclose publicly all sampled evidence. Moreover, the state space is rich, the prefer-

[29] **?** combines sequential testing with sequential search in a context in which the agent determines the intensity of sampling of alternatives in each round.

ences of both players are state-dependent, and they disagree on the relative importance of attributes.

8 Concluding remarks

This paper studied the optimal sampling of a multi-attribute project of unknown quality. We explored both the case of single-agent evaluation, in which sampling and adoption are in the hands of a single player, and that of shared evaluation, in which a principal and an agent have separate authorities on these decisions. The single-player optimal sample takes a simple form: it balances the variability of the sampled attributes, on the one hand, with their usefulness for inferring the realizations of unsampled attributes, on the other. Optimal sampling treats projects with the same initially known attribute in the same way, despite their ex-ante expected values. Moreover, the optimal sample can be implemented simultaneously. Principal-agent evaluation gives rise to optimal sampling that is biased by the ex-ante expected value of the project for the two players. Sampling need not be a compromise between the samples the principal and the agent would choose in their respective single-player problems: we characterized conditions under which the agent optimally restricts the informativeness of the sample for the principal and for himself.

The current analysis restricts attention to simultaneous sampling for principal-agent evaluation: the agent determines which attributes to sample prior to observing any of their realizations. Naturally, the equivalence between simultaneous and sequential sampling, which was a key feature of the single-agent problem, breaks down in the presence of separate authorities. An immediate extension is to study sequential sampling in principal-agent evaluation.

Moreover, the model assumed positive weights across all attributes. This guarantees that the players agree that higher attribute realizations are preferable. Allowing for negative weights would introduce new distortions into principal-agent evaluation. Exploring such distortions is a natural extension. Also, we intentionally assumed away all sources of asymmetric information between players, so as to focus the analysis on distortions due to dissimilar interests on attributes. Extending the analysis to scenarios in which the agent holds private information over a_0 and/or B_0 is another interesting direction for future work. This direction opens up the possibility of studying optimal delegation of sampling to a privately informed agent.

A Proofs for sections 2 and 3

Proof for lemma 2.1. Let $\mathbf{s} = (a_1, \dots, a_n)$, where $0 \leq a_1 < a_2 < \dots < a_n \leq 1$. The posterior expected value given $(\mathbf{s}, \mathbf{B}(\mathbf{s}))$ is:

$$\begin{aligned}
v(\mathbf{s},\mathbf{B}(\mathbf{s})) &= \mathbb{E}\left[\int_0^{a_1} B(a)\omega(a)da + \sum_{i=1}^{n-1}\int_{a_i}^{a_{i+1}} B(a)\omega(a)da + \int_{a_n}^{1}\omega(a)B(a)da \middle| \mathbf{s},\mathbf{B}(\mathbf{s})\right] \\
&= B(a_1)\int_0^{a_1}\omega(a)da - \mu\int_0^{a_1}(a_1-a)\omega(a)da + \sum_{i=1}^{n-1}\left\{B(a_i)\left(\int_{a_i}^{a_{i+1}}\frac{a_{i+1}-a}{a_{i+1}-a_i}\omega(a)da\right)\right. \\
&\quad \left. + B(a_{i+1})\left(\int_{a_i}^{a_{i+1}}\frac{a-a_i}{a_{i+1}-a_i}\omega(a)da\right)\right\} + B(a_n)\int_{a_n}^{1}\omega(a)da + \mu\int_{a_n}^{1}\omega(a)(a_n-a)da \\
&= B(a_1)\underbrace{\left(\int_0^{a_1}\omega(a)da + \int_{a_1}^{a_2}\frac{a_2-a}{a_2-a_1}\omega(a)da\right)}_{\tau(a_1;\mathbf{s})} + \sum_{i=2}^{n-1}B(a_i)\underbrace{\left(\int_{a_{i-1}}^{a_i}\frac{a-a_{i-1}}{a_i-a_{i-1}}\omega(a)da + \int_{a_i}^{a_{i+1}}\frac{a_{i+1}-a}{a_{i+1}-a_i}\omega(a)da\right)}_{\tau(a_i;\mathbf{s})} \\
&\quad + B(a_n)\underbrace{\left(\int_{a_{n-1}}^{a_n}\frac{a-a_{n-1}}{a_n-a_{n-1}}\omega(a)da + \int_{a_n}^{1}\omega(a)da\right)}_{\tau(a_n;\mathbf{s})} + \mu\left(\int_{a_n}^{1}\omega(a)(a_n-a)da - \int_0^{a_1}\omega(a)(a_1-a)da\right).
\end{aligned}$$

□

Proof for lemma 2.2. (i) Up to t, agent has sampled $\mathbf{s}^t := \mathbf{s}_0 \cup \mathbf{s}_1 \cup \dots \cup \mathbf{s}_t$; hence $\mathbf{s}^t$ is the profile of all known attributes up to t. Given $(\mathbf{s^t}, \mathbf{B}(\mathbf{s^t}))$, v_t is a linear combination of realizations in $B(\mathbf{s}^t)$. Hence, given $(\mathbf{s^k}, \mathbf{B}(\mathbf{s^k}))$ at some $k < t$, realizations of attributes in $\mathbf{s}_{k+1} \cup \dots \cup \mathbf{s}_t$ are normally distributed. v_t is a sum of normally distributed random variables, hence it is normally distributed for any $k < t$.

(ii) Fix an arbitrary t and let $\mathbf{s}_t = \{a_t\}$ be a singleton (the case of larger cardinality follows from identical reasoning). Suppose $\mathbf{s}^{t-1} = \mathbf{s}_0 \cup \mathbf{s}_1 \cup \dots \cup \mathbf{s}_{t-1}$ consists of N attributes. Index them in increasing order as $0 \leq a_{(1)} < \dots < a_{(N)} \leq 1$. From lemma 2.1,

$$v_{t-1}\left(\mathbf{s}^{t-1}, \mathbf{B}\left(\mathbf{s}^{t-1}\right)\right) = \sum_{i=1}^{N}\tau(a_{(i)};\mathbf{s}^{t-1})B(a_{(i)}) + \mu\left(\int\limits_{a_{(N)}}^{1}(s-a_{(t)})\omega(s)ds - \int\limits_{0}^{a_{(1)}}(a_{(1)}-s)\omega(s)ds\right).$$

We need to distinguish two cases.

Case I: Suppose first that $a_{(k)} < a_t < a_{(k+1)}$ for some $k \leq N$. Then,

$$\mathbb{E}[B(a_t)] = \frac{a_{(k+1)}-a_t}{a_{(k+1)}-a_{(k)}}B(a_{(k)}) + \frac{a_t-a_{(k)}}{a_{(k+1)}-a_{(k)}}B(a_{(k+1)}).$$

When taking the expectation of v_t with respect to the realization $B(a_t)$, the coefficient in front of $B(a_{(k)})$ is:

$$\begin{aligned}
\tau(a_{(k)};\mathbf{s}^t) + \frac{a_{(k+1)}-a_t}{a_{(k+1)}-a_{(k)}}\tau(a_t;\mathbf{s}^t) &= \left(\int_{a_{(k-1)}}^{a_{(k)}}\frac{s-a_{(k-1)}}{a_{(k)}-a_{(k-1)}}\omega(s)ds + \int_{a_{(k)}}^{a_t}\frac{a_t-s}{a_t-a_{(k)}}\omega(s)ds\right) \\
&\quad + \frac{a_{(k+1)}-a_t}{a_{(k+1)}-a_{(k)}}\left(\int_{a_{(k)}}^{a_t}\frac{s-a_{(k)}}{a_t-a_{(k)}}\omega(s)ds + \int_{a_t}^{a_{(k+1)}}\frac{a_{(k+1)}-s}{a_{(k+1)}-a_{(k)}}\omega(s)ds\right) \\
&= \int_{a_{(k-1)}}^{a_{(k)}}\frac{s-a_{(k-1)}}{a_{(k)}-a_{(k-1)}}\omega(s)ds + \int_{a_{(k)}}^{a_{(k+1)}}\frac{a_{(k+1)}-s}{a_{(k+1)}-a_{(k)}}\omega(s)ds \\
&= \tau\left(a_{(k)};\mathbf{s}^{t-1}\right).
\end{aligned}$$

Similarly, it can also be shown that: $\tau(a_{(k+1)};\mathbf{s}^t) + \frac{a_t - a_{(k)}}{a_{(k+1)} - a_{(k)}}\tau(a_t;\mathbf{s}^t) = \tau(a_{(k+1)};\mathbf{s}^{t-1})$. Hence, taking the expectation of v_t with respect to realization $B(a_t)$ yields exactly v_{t-1}.

Case II: Suppose that $a_t \in [0, a_{(1)}] \cup [a_{(N)}, 1]$. Consider $a_t < a_{(1)}$. The expectation of its realization is:

$$\mathbb{E}[B(a_t)] = B(a_{(1)}) - \mu(a_{(1)} - a_t).$$

Hence, when evaluating $\mathbb{E}[v_t]$, the coefficient in front of $B(a_{(1)})$ is $\tau(a_{(1)};\mathbf{s}^t) + \tau(a_t;\mathbf{s}^t)$. It is straightforward to verify that

$$\tau(a_{(1)};\mathbf{s}^t) + \tau(a_t;\mathbf{s}^t) = \tau(a_{(1)};\mathbf{s}^{t-1}).$$

Moreover, v_t and v_{t-1} differ in the last additive term featuring μ. When taking $\mathbb{E}[v_t]$, this last term collects:

$$-\mu(a_{(1)} - a_t)\tau(a_t;\mathbf{s}^t) - \mu\int_0^{a_t}(a_t - s)\omega(s)ds$$

which straightforwardly simplifies to $-\mu\int_0^{a_{(1)}}(a_{(1)} - s)\omega(s)ds$. This is precisely the analagous term in v_{t-1}. Hence, $\mathbb{E}[v_t] = v_{t-1}$. Identical reasoning shows that $\mathbb{E}[v_t] = v_{t-1}$ for $a_t > a_{(N)}$ as well.

(iii) From parts (i) and (ii), for any t, $v_t \sim \mathcal{N}(v_{t-1}, \sigma_t)$, where $\sigma_t := var(v_t \mid \mathbf{s}^{t-1}, \mathbf{B}(\mathbf{s}^{t-1}))$. This variance term does not depend on $(v_{t-1}, \ldots, v_0)$, because the coefficients in front of $B(a)$ for any $a \in \mathbf{s}_t$, i.e. $\tau(a;\mathbf{s}^t)$, do not depend on any realizations. Therefore, the density function for v_t depends only on v_{t-1}.

□

Proof for lemma 3.1. (i) Let $\mathbf{s} = (a_1, \ldots, a_{q+1})$ denote the sample of known attributes. Suppose $B(\mathbf{s}) < 0$. Then,

$$\sum_{i=1}^{q+1} B(a_i)\tau(a_i;\mathbf{s}) < 0.$$

If extrapolation to peripheral attributes is such that:

$$\int_{a_{q+1}}^1 \omega(a)(a - a_{q+1})da - \int_0^{a_1}(a_1 - a)\omega(a)da > (<)0$$

there exists $\mu > 0$ sufficiently large (resp., $\mu < 0$ sufficiently small) for which $v_1 > 0$. The proof for $B(\mathbf{s}) > 0$ is similar and ommitted.

(ii) Let $\mathbf{s} = (a_1, \ldots, a_{q+1})$ denote the sample of known attributes. We show that for any arbitrary k, the adoption threshold $\bar{B}(a)$ is discontinuous at $a = a_k$. Consider attributes $a_k - \varepsilon$ and $a_k + \varepsilon$, for $\varepsilon > 0$ sufficiently small so that $a_k - \varepsilon > a_{k-1}$ and $a_k + \varepsilon < a_{k+1}$. Suppose first that $1 < k < q+1$, hence $a_{k-1} > 0$ and $a_{k+1} < 1$. Note that the coefficient τ corresponding to $a_k - \varepsilon$ converges to

$$\begin{aligned}\lim_{\varepsilon\to 0} \quad \tau(a_k - \varepsilon;\mathbf{s}\cup a_k - \varepsilon) &= \lim_{\varepsilon\to 0} \quad \int_{a_{k-1}}^{a_k-\varepsilon}\frac{s - a_{k-1}}{a_k - a_{k-1}}\omega(s)ds + \int_{a_k-\varepsilon}^{a_k}\frac{a_k - s}{\varepsilon}\omega(s)ds \\ &= \int_{a_{k-1}}^{a_k}\frac{s - a_{k-1}}{a_k - a_{k-1}}\omega(s)ds.\end{aligned}$$

Similarly,

$$\lim_{\varepsilon\to 0} \quad \tau(a_{k-1};\mathbf{s}\cup a_k - \varepsilon) = \tau(a_{k-1};\mathbf{s}), \quad \lim_{\varepsilon\to 0} \quad \tau(a_k + \varepsilon;\mathbf{s}\cup a_k + \varepsilon) = \int_{a_k}^{a_{k+1}}\frac{a_{k+1} - s}{a_{k+1} - a_k}\omega(s)ds.$$

Therefore, as ε shrinks to zero, the adoption threshold approaches

$$\lim_{\varepsilon\to 0}\bar{B}(a_k-\varepsilon)=\frac{-B(a_k)\int_{a_k}^{a_{k+1}}\frac{a_{k+1}-s}{a_{k+1}-a_k}\omega(s)ds-\sum_{i\neq k}\tau(a_i;\mathbf{s})B(a_i)-\mu\Big(\int_{a_{q+1}}^{1}\omega(a)(a-a_{q+1})da-\int_0^{a_1}(a_1-a)\omega(a)da\Big)}{\int_{a_{k-1}}^{a_k}\frac{s-a_{k-1}}{a_k-a_{k-1}}\omega(s)ds},$$

where a_1 and a_{q+1} are the leftmost and the rightmost attributes in $\mathbf{s}$. By similar reasoning,

$$\lim_{\varepsilon\to 0}\bar{B}(a_k+\varepsilon)=\frac{-B(a_k)\int_{a_{k-1}}^{a_k}\frac{s-a_{k-1}}{a_k-a_{k-1}}\omega(s)ds-\sum_{i\neq k}\tau(a_i;\mathbf{s})B(a_i)-\mu\Big(\int_{a_{q+1}}^{1}\omega(a)(a-a_{q+1})da-\int_0^{a_1}(a_1-a)\omega(a)da\Big)}{\int_{a_k}^{a_{k+1}}\frac{a_{k+1}-s}{a_{k+1}-a_k}\omega(s)ds}.$$

Generically,

$$\int_{a_k}^{a_{k+1}}\frac{a_{k+1}-s}{a_{k+1}-a_k}\omega(s)ds\neq\int_{a_{k-1}}^{a_k}\frac{s-a_{k-1}}{a_k-a_{k-1}}\omega(s)ds,$$

hence $\lim_{\varepsilon\to 0}\bar{B}(a_k-\varepsilon)\neq\lim_{\varepsilon\to 0}\bar{B}(a_k+\varepsilon)$. The adoption threshold is discontinuous at $a=a_k$.

The reasoning for the other two cases (i) $k=q+1$, and (ii) $k=1$ is similar and ommited. The genericity conditions for these two cases respectively are:

$$\int_{a_{q+1}}^{1}\omega(s)ds\neq\int_{a_q}^{a_{q+1}}\frac{s-a_q}{a_{q+1}-a_q}\omega(s)ds,\qquad \int_0^{a_1}\omega(s)ds\neq\int_{a_1}^{a_2}\frac{a_2-s}{a_2-a_1}\omega(s)ds.$$

□

Lemma A.1 (The variance of $v(\mathbf{s},\mathbf{B}(\mathbf{s}))$)**.** *Consider a sample $\mathbf{s}\in\mathscr{S}_{q+1}(a_0)$ such that*

$$\mathbf{s}=\left(a_m^{\ell},\ldots,a_1^{\ell},a_0,a_1^r,\ldots,a_{q-m}^r\right).$$

The posterior expected value $v(\mathbf{s},\mathbf{B}(\mathbf{s}))$ is normally distributed:

$$v(\mathbf{s},\mathbf{B}(\mathbf{s}))\sim\mathcal{N}\left(v_0,\ \sum_{i=1}^{m}\tau^2\left(a_i^{\ell};(a_i^{\ell},a_{i-1}^{\ell})\right)(a_{i-1}^{\ell}-a_i^{\ell})+\sum_{i=1}^{q-m}\tau^2\left(a_i^r;(a_{i-1}^r,a_i^r)\right)(a_i^r-a_{i-1}^r)\right).$$

Proof for lemma A.1. Attributes $a_1^{\ell}>\ldots>a_m^{\ell}$ and $a_1^r<\ldots<a_{q-m}^r$ are sampled to the left and to the right of a_0 respectively. Note that for any $t\in\{\ell,r\}$ and any $k_\ell\in\{1,\ldots,m\}$, $k_r\in\{1,\ldots,q-m\}$, by the property of independent increments of the Brownian motion, each random variable $B(a_{k_t}^t)$ can be expressed as $B_0+\mu|a_{k_t}^t-a_0|+\sum_{i=1}^{k_t}Z_i^t$, where $Z_i^t\sim\mathcal{N}(0,|a_i^t-a_{i-1}^t|)$ are independent across i and t. Hence,

$$\begin{aligned}v(\mathbf{s},\mathbf{B}(s))&=\tau(a_0;\mathbf{s})B_0+\sum_{i=1}^{m}\tau(a_i^{\ell};\mathbf{s})B(a_i^{\ell})+\sum_{j=1}^{q-m}\tau(a_j^r;\mathbf{s})B(a_j^r)+\mu\left(\int_{a_{q-m}^r}^{1}(a-a_{q-m}^r)\omega(a)da-\int_0^{a_m^{\ell}}(a_m^{\ell}-a)\omega(a)da\right)\\&=B_0+\sum_{k=1}^{m}\left(\sum_{i=k}^{m}\tau(a_i^{\ell};\mathbf{s})\right)Z_k^{\ell}+\sum_{k=1}^{q-m}\left(\sum_{i=k}^{q-m}\tau(a_i^r;\mathbf{s})\right)Z_k^r\mu\left(\int_{a_{q-m}^r}^{1}(a-a_{q-m}^r)\omega(a)da-\int_0^{a_m^{\ell}}(a_m^{\ell}-a)\omega(a)da\right)\end{aligned}$$

The manipulation of $v(s)$ into this new form has made use of the fact that the weights sum up to 1:

$$\tau(a_0;\mathbf{s})+\sum_{i=1}^{m}\tau(a_i^{\ell};\mathbf{s})+\sum_{j=1}^{q-m}\tau(a_j^r;\mathbf{s})=1.$$

First, $v(\mathbf{s},\mathbf{B}(\mathbf{s}))$ is a weighted sum of normally distributed random variables, so it is itself distributed normally. Using the fact that $\mathbb{E}[Z_i^t]=0$ for any i,t, we have that the mean of $v(\mathbf{s},\mathbf{B}(\mathbf{s}))$ is $\mathbb{E}[v(\mathbf{s},\mathbf{B}(\mathbf{s}))\mid B_0,\mathbf{s}]=B_0$. Let us now consider the variance of $v(\mathbf{s},\mathbf{B}(\mathbf{s}))$. Remember that Z_i^t and $Z_k^{t'}$ are independent for any t, t', i, k:

$$var\left[v(\mathbf{s}, \mathbf{B}(\mathbf{s})) \mid B_0, \mathbf{s}\right] = \sum_{k=1}^{m} \left(\sum_{i=k}^{m} \tau(a_i^\ell; \mathbf{s}) \right)^2 (a_{k-1}^\ell - a_k^\ell) + \sum_{k=1}^{q-m} \left(\sum_{i=k}^{q-m} \tau(a_i^r; \mathbf{s}) \right)^2 (a_k^r - a_{k-1}^r).$$

Note that

$$\begin{aligned}
\sum_{i=k}^{m} \tau(a_i^\ell; \mathbf{s}) = & \int_0^{a_m^\ell} \omega(a)da + \int_{a_m^\ell}^{a_{m-1}^\ell} \frac{a_{m-1}^\ell - a}{a_{m-1}^\ell - a_m^\ell} \omega(a)da + \int_{a_m^\ell}^{a_{m-1}^\ell} \frac{a - a_m^\ell}{a_{m-1}^\ell - a_m^\ell} \omega(a)da + \\
& \int_{a_{m-1}^\ell}^{a_{m-2}^\ell} \frac{a_{m-2}^\ell - a}{a_{m-2}^\ell - a_{m-1}^\ell} \omega(a)da + \ldots + \int_{a_{k-1}^\ell}^{a_k^\ell} \frac{a - a_{k-1}^\ell}{a_k^\ell - a_{k-1}^\ell} \omega(a)da + \int_{a_k^\ell}^{a_0} \frac{a_0 - a}{a_0 - a_k^\ell} \omega(a)da \\
= & \int_0^{a_k^\ell} \omega(a)da + \int_{a_k^\ell}^{a_0} \frac{a_0 - a}{a_0 - a_k^\ell} \omega(a)da \\
= & \tau\left(a_i^\ell; (a_i^\ell, a_{i-1}^\ell)\right).
\end{aligned}$$

A similar argument shows that

$$\sum_{i=k}^{q-m} \tau(a_i^r; \mathbf{s}) = \int_{a_0}^{a_k^r} \frac{a - a_0}{a_k^r - a_0} \omega(a)da + \int_{a_k^r}^{1} \omega(a)da = \tau\left(a_i^r; (a_i^r, a_{i-1}^r)\right).$$

This yields the desired expression for the variance. □

Proof for proposition 3.2. (i) The objective of the agent is:

$$\max_{\mathbf{s} \in \mathscr{S}_{q+1}(a_0)} \Pr(v(\mathbf{s}, \mathbf{B}(\mathbf{s})) \geq 0) \mathbb{E}[v(\mathbf{s}, \mathbf{B}(\mathbf{s})) \mid v(\mathbf{s}, \mathbf{B}(\mathbf{s})) \geq 0].$$

Using the distribution of $v(\mathbf{s}, \mathbf{B}(\mathbf{s}))$ derived in lemma A.1,

$$\Pr(v(\mathbf{s}, \mathbf{B}(\mathbf{s}))) = 1 - \Phi\left(\frac{-v_0}{\sigma(\mathbf{s})}\right) = \Phi\left(\frac{v_0}{\sigma(\mathbf{s})}\right),$$

$$\mathbb{E}[v(\mathbf{s}, \mathbf{B}(\mathbf{s})) \mid v(\mathbf{s}, \mathbf{B}(\mathbf{s})) \geq 0] = v_0 + \sigma(\mathbf{s}) \frac{\phi\left(\frac{v_0}{\sigma(\mathbf{s})}\right)}{\Phi\left(\frac{v_0}{\sigma(\mathbf{s})}\right)}.$$

Therefore, the objective of the agent simplifies to

$$\max_{\mathbf{s}} \quad v_0 \Phi\left(\frac{v_0}{\sigma(\mathbf{s})}\right) + \sigma(\mathbf{s}) \phi\left(\frac{v_0}{\sigma(\mathbf{s})}\right).$$

The objective is strictly increasing in $\sigma(\mathbf{s})$, hence an optimal sample maximizes $\sigma(\mathbf{s})$.

Suppose that an optimal sample $\mathbf{s}^*$ consists of strictly less than $q+1$ distinct attributes; let $\mathbf{s}^* = (a_m^\ell, \ldots, a_1^\ell, a_0, a_1^r, \ldots, a_{m'}^r)$ such that $m' + m < q$. By the argument in part (ii) of this proof, $a_m^\ell > 0$ and $a_{m'}^r < 1$. Consider the following modification of $\mathbf{s}^*$ to a larger sample $\tilde{\mathbf{s}}$: $\tilde{a}_i^r = a_i^r$ for all $i = 1, \ldots, m'$, $\tilde{a}_{m'+1}^r = a_{m'}^r + \varepsilon < 1$, and $\tilde{a}_i^\ell = a_i^\ell$ for all $i = 1, \ldots, m$. That is, $\tilde{\mathbf{s}}$ differs from $\mathbf{s}^*$ in that it samples one more attribute between $a_{m'}^r$ and 1. This strictly increases the sample variance by $\varepsilon\tau(a_{m'}^r + \varepsilon; (a_{m'}^r + \varepsilon, a_{m'}^r)) > 0$. This contradicts the supposed optimality of $\mathbf{s}^*$.

(ii) Let $\mathbf{s}^* = (a_1^*, \ldots, a_{q+1}^*)$. Suppose, towards a contradiction, $a_0 = a_k$ for some $k < q+1$ and that

$a^*_{q+1} = 1$ in the optimal sample. The term

$$(a^*_{q+1} - a^*_q)\tau^2(a^*_{q+1}; (a^*_q, a^*_{q+1}))$$

is strictly decreasing at $a^*_{q+1} = 1$ because its derivative with respect to a^*_{q+1} is

$$\tau(a^*_{q+1}; (a^*_q, a^*_{q+1}))\left(\int_{a^*_{q+1}}^1 \omega(a)da - \int_{a^*_q}^{a^*_{q+1}} \frac{a - a^*_q}{a^*_{q+1} - a^*_q}\omega(a)da\right)\Big|_{a^*_{q+1}=1} < 0.$$

Because a^*_{q+1} appears in $\sigma^2(\mathbf{s})$ only through this term, $\sigma^2(\mathbf{s})$ is strictly decreasing at $a^*_{q+1} = 1$. Hence, $\sigma^2(\mathbf{s})$ strictly increases if a^*_{q+1} decreases by a sufficiently small ϵ, so that $1 - \varepsilon > a^*_q$. Hence, the agent's payoff is strictly improved by this modification. This contradicts the optimality of the original sample. A similar argument shows that if $k > 1$, then $a^*_1 \neq 0$ as well. □

Proof for proposition 3.3. Let $\sigma^2_q := \max_{\mathbf{s}\in\mathcal{S}_{q+1}(a_0)} \sigma^2(\mathbf{s})$ denote the maximum attained variance of the posterior expected value if q attributes are sampled beyond a_0. The proof proceeds in two steps: (i) we show that

$$\sigma^2_{q+1} - \sigma^2_q < \sigma^2_q - \sigma^2_{q-1},$$

and (ii) we show that $V_{q+1} - V_q < V_q - V_{q-1}$.

(i) We show this property for an arbitrary sample under capacity q constrained to be in $[a_0, 1]$. The reasoning is identical for samples constrained in $[0, a_0]$; taken together, these two observations establish the property for unconstrained samples in $[0, 1]$.

Let $R_q := \{(a_0, a_1, \ldots, a_q) \in [a_0, 1]^{q+1} : a_0 < a_1 < a_2 < \ldots < a_q\}$. A sample $\mathbf{s} = (a_0, a_1, a_2, \ldots, a_q) \in R_q$ induces variance on the posterior expected value equal to $\sigma^2(\mathbf{s}) = \sum_{k=0}^{q-1} v(a_k, a_{k+1})$, where

$$v(a_k, a_{k+1}) = (a_{k+1} - a_k)\tau^2(a_{k+1}; (a_k, a_{k+1})).$$

Lemma A.2. *The function $v(a_{i-1}, a_i) := (a_i - a_{i-1})\tau^2(a_i; (a_i, a_{i-1}))$ is strictly supermodular in $(a_{i-1}, a_i) \in R_2$.*

Proof. Given (a_{i-1}, a_i), consider an arbitrary $a' \in [a_{i-1}, a_i]$. The added value from splitting $[a_{i-1}, a_i]$ into two subcells $[a_{i-1}, a']$ and $[a', a_i]$ is $\Delta v(a', a_{i-1}, a_i) := v(a_{i-1}, a') + v(a', a_i) - v(a_{i-1}, a_i)$. This added value is strictly increasing in a_i and strictly decreasing in a_{i-1}:

$$\frac{\partial \Delta v}{\partial a_{i-1}} = \tau^2(a_i; (a_{i-1}, a_i)) - \tau^2(a'; (a_{i-1}, a')) < 0,$$

$$\begin{aligned}\frac{\partial \Delta v}{\partial a_i} &= \tau(a_i; (a_{i-1}, a_i))\left(\int_{a_{i-1}}^{a_i} \frac{s - a_{i-1}}{a_i - a_{i-1}}\omega(s)ds - \int_{a_i}^1 \omega(s)ds\right) \\ &\quad - \tau(a_i; (a', a_i))\left(\int_{a'}^{a_i} \frac{s - a'}{a_i - a'}\omega(s)ds - \int_{a_i}^1 \omega(s)ds\right) \\ &= \left(\int_{a_{i-1}}^{a_i} \frac{s - a_{i-1}}{a_i - a_{i-1}}\omega(s)ds\right)^2 - \left(\int_{a'}^{a_i} \frac{s - a'}{a_i - a'}\omega(s)ds\right)^2 > 0.\end{aligned}$$

The first inequality uses the observation that because $a' < a_i$, $\tau^2(a'; (a_{i-1}, a')) > \tau^2(a_i; (a_{i-1}, a_i))$. The second inequality uses the fact that $\int_{a_{i-1}}^{a_i} \frac{s - a_{i-1}}{a_i - a_{i-1}}\omega(s)ds$ is increasing in a_{i-1}.

Take $(\tilde{a}_{i-1}, \tilde{a}_i)$ such that $a_{i-1} < \tilde{a}_{i-1} < \tilde{a}_i < a_i$. Applying the property we just established for $a' = \tilde{a}_i$,

$\Delta v(\tilde{a}_i, \tilde{a}_{i-1}, a_i) < \Delta v(\tilde{a}_i, a_{i-1}, a_i)$. Simplifying this inequality, it yields

$$v(a_{i-1}, \tilde{a}_i) + v(\tilde{a}_{i-1}, a_i) > v(\tilde{a}_{i-1}, \tilde{a}_i) + v(a_{i-1}, a_i).$$

Letting $a := (a_{i-1}, a_i)$ and $\tilde{a} := (\tilde{a}_{i-1}, \tilde{a}_i)$, this translates into the familiar strict supermodularity inequality $v(a \wedge \tilde{a}) + v(a \vee \tilde{a}) > v(a) + v(\tilde{a})$, where $\wedge$ and $\vee$ denote the componentwise minimum and maximum, respectively. Hence, v is strictly supermodular in its pair of arguments. □

Lemma A.3. *The variance* $\sigma^2(a_0, a_1, \dots, a_q) := \sigma^2(\mathbf{s})$ *is strictly supermodular in* $(a_1, \dots, a_q)$.

Proof. By lemma A.2, σ^2 is a sum of strictly supermodular functions of the form $v(a_k, a_{k+1})$, hence it is strictly supermodular. □

By lemma A.3, σ_q^2 is supermodular for any q. We want to show that for any $q \geq 1$,

$$\sigma_{q+1}^2 - \sigma_q^2 < \sigma_q^2 - \sigma_{q-1}^2 \quad \Leftrightarrow \quad \sigma_{q+1}^2 + \sigma_{q-1}^2 < 2\sigma_q^2.$$

Let $\mathbf{s} = (a_0, a_1, \dots, a_{q-1}) \in R_{q-1}$ and $\tilde{\mathbf{s}} = (a_0, \tilde{a}_1, \dots, \tilde{a}_{q+1}) \in R_{q+1}$ be two arbitrary samples of sizes $q-1$ and $q+1$ respectively. The sample $\mathbf{s}$ is equivalent to the extended sample $(a_0, \mathbf{s}, a_{q-1})$ (where both a_0 and a_{q-1} is repeated twice). Hence,

$$\begin{aligned}
\sigma^2(\tilde{\mathbf{s}}) + \sigma^2(\mathbf{s}) &= \sigma^2(\tilde{\mathbf{s}}) + \sigma^2(a_0, \mathbf{s}, 1) - (1 - a_{q-1})\tau^2(1; (a_{q-1}, 1)) && (21)\\
&< \sigma^2(\tilde{\mathbf{s}} \wedge (a_0, \mathbf{s}, 1)) + \sigma^2(\tilde{\mathbf{s}} \vee (a_0, \mathbf{s}, 1)) - (1 - a_{q-1})\tau^2(1; (a_{q-1}, 1)) && (22)\\
&\leq 2\sigma_q^2. && (23)
\end{aligned}$$

The first inequality follows from the fact that the sample $(a_0, \mathbf{s}, 1)$ is strictly more variable than $(a_0, \mathbf{s}, a_{q-1})$ (in which a_{q-1} is repeated twice). The second inequality uses strict supermodularity. Notice that the first two attributes in sample $\tilde{\mathbf{s}} \wedge (a_0, \mathbf{s}, 1)$ are necessarily a_0, hence $\tilde{\mathbf{s}} \wedge (a_0, \mathbf{s}, 1) \in R_q$. Sample $\tilde{\mathbf{s}} \vee (a_0, \mathbf{s}, 1)$ starts with a_0 and necessarily ends with attribute 1. If $a_{q-1} \geq \tilde{a}_q$,

$$\sigma^2(\tilde{\mathbf{s}} \vee (a_0, \mathbf{s}, 1)) = \sigma^2(\tilde{\mathbf{s}} \setminus \tilde{a}_{q+1} \vee (a_0, \mathbf{s})) + (1 - a_{q-1})\tau^2(1; (a_{q-1}, 1)),$$

where $\tilde{\mathbf{s}} \setminus \tilde{a}_{q+1} \vee (a_0, \mathbf{s})$ ends with a_{q-1}. This reduces the second line (22) to a sum of variances of two samples of size q. By optimality, it must be that this sum is less than $2\sigma_q^2$. If, on the other hand, $a_{q-1} < \tilde{a}_q$,

$$\begin{aligned}
\sigma^2(\tilde{\mathbf{s}} \vee (a_0, \mathbf{s}, 1)) &= \sigma^2(\tilde{\mathbf{s}} \setminus \tilde{a}_{q+1} \vee (a_0, \mathbf{s})) + (1 - \tilde{a}_{q+1})\tau^2(1; (\tilde{a}_{q+1}, 1))\\
&< \sigma^2(\tilde{\mathbf{s}} \setminus \tilde{a}_{q+1} \vee (a_0, \mathbf{s})) + (1 - a_{q-1})\tau^2(1; (a_{q-1}, 1)),
\end{aligned}$$

because $a_{q-1} < \tilde{a}_{q+1}$, hence $\tau(1; (\tilde{a}_{q+1}, 1)) < \tau(1; (a_{q-1}, 1))$. This again gives the last inequality in (23).

(ii) Consider

$$V(q) := \nu_0 \Phi\left(\frac{\nu_0}{\sigma_q}\right) + \sigma_q \phi\left(\frac{\nu_0}{\sigma_q}\right).$$

V is strictly increasing and strictly convex in σ. Moreover, $\sigma_{q-1}^2 < \sigma_q^2 < \sigma_{q+1}^2$, and by part (i) of this proof, $\sigma_q^2 > \frac{1}{2}\left(\sigma_{q+1}^2 + \sigma_{q-1}^2\right)$. Therefore,

$$\begin{aligned} V_q = V(\sigma_q^2) &> V\left(\frac{1}{2}\left(\sigma_{q+1}^2 + \sigma_{q-1}^2\right)\right) \\ &> \frac{1}{2}V(q+1) + \frac{1}{2}V(q-1), \end{aligned}$$

where the first line follows from strict monotonicity and the second follows from strict convexity. □

Proof for proposition 3.4. (i) Consider samples $\mathbf{s}$ and $\tilde{\mathbf{s}}$ respectively optimal for capacities q and $q+1$. Suppose that $\mathbf{s}$ samples $m \le q$ attributes in $[a_0, 1]$ while $\tilde{\mathbf{s}}$ samples $\tilde{m} \le q+1$ attributes in $[a_0, 1]$. Without loss, let $\tilde{m} < m$. The allocation of attributes in $[0, a_0]$ and $[a_0, 1]$ is $(q-m, m)$ and $(q+1-\tilde{m}, \tilde{m})$ respectively. Let $(\sigma_k^i)^2$, where $i \in \{\ell, r\}$, denote the variance attained by the k attributes on side i (left or right) of a_0. By optimality of $\mathbf{s}$,

$$(\sigma_{q-m}^\ell)^2 + (\sigma_m^r)^2 \ge (\sigma_{q-\tilde{m}}^\ell)^2 + (\sigma_{\tilde{m}}^r)^2$$

Also by optimality of $\tilde{\mathbf{s}}$,

$$(\sigma_{q+1-\tilde{m}}^\ell)^2 + (\sigma_{\tilde{m}}^r)^2 \ge (\sigma_{q+1-m}^\ell)^2 + (\sigma_m^r)^2.$$

Combining these two,

$$(\sigma_{q-\tilde{m}}^\ell)^2 - (\sigma_{q-m}^\ell)^2 \le (\sigma_{q+1-\tilde{m}}^\ell)^2 - (\sigma_{q+1-m}^\ell)^2.$$

Because $q - m < q - \tilde{m}$, this contradicts proposition 3.3. Hence, it must be that $\tilde{m} \ge m$ and $q+1-\tilde{m} \ge q-m$. Therefore, $m \le \tilde{m} \le m+1$.

(ii) We show the property for constrained samples in $[0, a_0]$ and $[a_0, 1]$. Without loss, consider

$$\mathbf{s} = (a_0, a_1^r, \ldots, a_m^r), \quad \tilde{\mathbf{s}} = (a_0, \tilde{a}_1^r, \ldots, \tilde{a}_{m+1}^r)$$

for $m \in [1, q]$. We want to show that $\tilde{a}_1^r < a_1^r < \tilde{a}_2^r < \ldots < \tilde{a}_m^r < a_m^r < \tilde{a}_{m+1}^r$. Let $\tilde{\mathbf{s}}_m = (a_0, \tilde{a}_1^r, \ldots, \tilde{a}_m^r)$. By optimality of $\mathbf{s}$, $\sigma^2(\tilde{\mathbf{s}}) \ge \sigma^2(\mathbf{s} \wedge \tilde{\mathbf{s}}_m, \tilde{a}_{m+1}^r)$. Let $a > \tilde{a}_{m+1}^r$. By strict supermodularity of σ^2 established in lemma A.3, if $\tilde{\mathbf{s}}$ and $(\mathbf{s}, a)$ for $a \in (a_m^r, 1]$ are unordered,

$$\sigma^2(\mathbf{s} \vee \tilde{\mathbf{s}}_m, a) - \sigma^2(\mathbf{s}, a) > \sigma^2(\tilde{\mathbf{s}}) - \sigma^2(\mathbf{s} \wedge \tilde{\mathbf{s}}_m, a_{m+1}^r) \ge 0.$$

But $a - \max\{a_m^r, \tilde{a}_m^r\} < a - a_m^r$ and $\tau(a; (\max\{a_m^r, \tilde{a}_m^r\}, a)) < \tau(a; (a_m^r, a))$. Therefore, if $\sigma^2(\mathbf{s} \vee \tilde{\mathbf{s}}_m, a) > \sigma^2(\mathbf{s}, a)$, then $\sigma^2(\mathbf{s} \vee \tilde{\mathbf{s}}_m) > \sigma^2(\mathbf{s})$. This contradicts the optimality of $\mathbf{s}$ for capacity q. Hence $\tilde{\mathbf{s}}$ and $(\mathbf{s}, a)$ for $a > \tilde{a}_{m+1}^r$ are weakly ordered. That is, for any i, $\tilde{a}_i^r \le a_i^r$ for $i \le q$. Repeating this argument with $\tilde{\mathbf{s}}_1 = (a_0, \tilde{a}_2^r, \ldots, \tilde{a}_{m+1}^r)$ and $a \in [a_0, \tilde{a}_1^r)$, we obtain that $\tilde{a}_{1+i}^r \ge a_i^r$ for $i = 1, \ldots, q$. Hence, $\tilde{a}_i^r \le a_i^r \le \tilde{a}_{i+1}^r$ for all $i \le q$.

To show strict monotonicity, we invoke Theorem 3 in **?**(pg. 212). By proposition 3.2(ii), $\tilde{a}_1^r > a_0$. Let $a \in \{a_0, a_1^r\}$, where a_1^r is the second attribute in $\mathbf{s}$. We consider the choice of an optimal sample of capacity m with initial attribute a, constrained in $[a, 1]$. It is already established that the variance $\sigma^2(\mathbf{s}')$ is continuous and strictly supermodular in the sample $\mathbf{s}' = (a, a_1', \ldots, a_m')$ and in the m-tuple of discovered attributes $(a_1', \ldots, a_m')$. Moreover, the marginal returns to a_1' is increasing in a because

$$\begin{aligned} \frac{\partial \sigma^2}{\partial a_1'} &= \tau(a_1'; (a, a_1'))\left(\int_{a_1'}^1 \omega(s)ds - \int_a^{a_1'} \frac{s-a}{a_1'-a}\omega(s)ds\right) - \tau(a_2'; (a_1', a_2'))\left(\int_{a_1'}^1 \omega(s)ds + \int_{a_1'}^{a_2'} \frac{a_2'-s}{a_2'-a_1'}\omega(s)ds\right) \\ &= \left(\int_{a_1'}^1 \omega(s)ds\right)^2 - \left(\int_a^{a_1'} \frac{s-a}{a_1'-a}\omega(s)ds\right)^2 - \tau(a_2'; (a_1', a_2'))\left(\int_{a_1'}^1 \omega(s)ds + \int_{a_1'}^{a_2'} \frac{a_2'-s}{a_2'-a_1'}\omega(s)ds\right). \end{aligned}$$

The term $\int_a^{a_1'} \frac{s-a}{a_1'-a}\omega(s)ds$ is decreasing in a. Hence, this derivative is increasing in a. Moreover, the optimal samples for $a = a_0$ and $a = \tilde{a}_1^r$ are interior, as established in proposition 3.2. Hence, by Theorem 3 in **?**, $(\tilde{a}_2^r, \dots, \tilde{a}_{m+1}^r) > (a_1^r, \dots, a_m^r)$. Also, by the optimality of $\mathbf{s}$ and $\tilde{\mathbf{s}}$, it must be that $\tilde{a}_1^r < a_1^r$. Hence, these strict inequalities combined with the weak inequalities derived above imply that for any $i = 1, \dots, m$,

$$\tilde{a}_i^r < a_i^r < \tilde{a}_{i+1}^r.$$

□

Lemma A.4 (Strictly ordered optimal samples)**.** *Given capacity q, any two distinct optimal samples* $\mathbf{s}$ *and* $\mathbf{s}'$ *such that for some $i \le q$ $a_i \ne a_i'$ are strictly ordered, i.e. $a_i > a_i'$ for all i or $a_i < a_i'$ for all i.*

Proof. By theorem 2.7.5. in **?**, the strict supermodularity of $\sigma^2(\mathbf{s})$ implies that $\arg\max_s \sigma^2(\mathbf{s})$ is a chain, i.e., for any $\mathbf{s}', \mathbf{s}'' \in \arg\max_{\mathbf{s}} \sigma^2(\mathbf{s})$, either $a_i' \le a_i''$ for all i or $a_i'' \le a_i'$ for all i. Without loss, suppose that for some i, ${a^r}'_i = {a^r}''_i$. The FOC with respect to a_i is

$$\tau(a_i;(a_{i-1},a_i))\left(\int_{a_i}^1 \omega(s)ds - \int_{a_{i-1}}^{a_i} \frac{s-a_{i-1}}{a_i - a_{i-1}}\omega(s)ds\right) \\ - \tau(a_{i+1};(a_i,a_{i+1}))\left(\int_{a_i}^1 \omega(s)ds + \int_{a_i}^{a_{i+1}} \frac{a_{i+1}-s}{a_{i+1}-a_i}\omega(s)ds\right) = 0$$

The first term is increasing in a_{i-1} and the second term is decreasing in a_{i+1}. Hence, the RHS is increasing in both a_{i-1} and a_{i+1}. By the fact that $\mathbf{s}'$ and $\mathbf{s}''$ are ordered, it must be that ${a^r}'_{i-1} = {a^r}''_{i-1}$ and ${a^r}'_{i+1} = {a^r}''_{i+1}$. Repeating this argument, we obtain that $\mathbf{s}'$ and $\mathbf{s}''$ coincide in all attributes. Therefore, for any two optimal samples $\mathbf{s}'$ and $\mathbf{s}''$ either they are identical in all attributes or they are strictly ordered: ${a^r}'_i < {a^r}''_i$ for all i. □

Proof for proposition 3.5. (i) We establish the result by recursive reasoning, starting from the last sampling step q. Suppose attributes $\mathbf{s}_{q-1} = (a_0, a_1, \dots, a_{q-1})$ is sampled and the corresponding realizations are $\mathbf{B_{q-1}} = (B_0, B(a_1), \dots, B(a_{q-1}))$.[30] Let v_q denote the posterior expected value of the project. Therefore, for any unsampled a_q,

$$v_q \sim \mathcal{N}\left(v_{q-1}, \sigma^2(a_q;\mathbf{s}_{q-1})\right),$$

where $\sigma^2(\cdot;\mathbf{s})$ gives the variance in the posterior induced by sampling a after observing the realizations of $\mathbf{s}$. The optimal attribute is $a_q^* \in \arg\max_a \sigma^2(a;\mathbf{s}_{q-1})$. Consider sampling step $(q-k)$, for $q > k > 1$. Suppose the agent has sampled $\mathbf{s}_{q-k-1}$, with respective realizations $\mathbf{B_{q-k-1}}$ and posterior expectation v_{q-k}, and the optimal choice of the next attributes $\left(a_{q-k}^*, a_{q-k+1}^*, \dots, a_q^*\right)$ does not depend on $(v_{q-t})_{t=0}^k$. Consider now stage $q-k-1$, with sampled attributes $\mathbf{s}_{q-k-2}$ and posterior expected value v_{q-k-1}. For any a_{q-k-1},

$$v_{q-k-1} \sim \mathcal{N}\left(v_{q-k-2}, \sigma^2(a_{q-k-1};\mathbf{s}_{q-k-2})\right).$$

Because by hypothesis $(v_{q-t})_{t=0}^k$ does not inform the choice of subsequent attributes, the optimal choice of a_{q-k-1}^* is such that

$$(a_{q-k-1}^*, a_{q-k}^*, \dots, a_q^*) \in \underset{(a_{q-k-1}, a_{q-k}, \dots, a_q) \in [0,1]^k}{\arg\max} \sigma^2\left((a_{q-k-1}, a_{q-k}, \dots, a_q);\mathbf{s}_{q-k-2}\right).$$

This concludes the proof.

[30] In this proof, in order to more clearly denote sequentiality of sampling, the subindices of attributes denote timing of sampling rather than rank order.

(ii) By part (i), at any stage k the choice of $(a_k^*,\dots,a_q^*)$ is not informed by $(\nu_t)_{t=k}^q$. Hence, at stage $k=1$, the agent chooses a sample $(a_1^*,\dots,a_q^*)$ that maximizes $\sigma((a_1^*,\dots,a_q^*);a_0)$. This is the same problem as that faced in simultaneous sampling with capacity q. It remains to establish the sequential rationality of $(a_1^*,\dots,a_q^*)$ implemented in *some* order.

Relabel the optimal sample $(a_1^*,\dots,a_q^*)$ as $(a_1^\ell,\dots,a_m^\ell,a_1^r,\dots,a_{q-m}^r)$ in terms of distance from a_0. Consider a sequential strategy $\tilde{a}_t$ that takes attributes in increasing distance from a_0 in either side of a_0: let $\tilde{a}_t = a_t^\ell$ for $t \le m$ and $\tilde{a}_t = a_{t-m}^r$ for $m < t \le q$. Consider $\tilde{a}_k$ at step $k < m$.

1. It is suboptimal to pick $m' \neq m$ attributes in $[0,a_0)$ and $q-m'-k+1$ attributes in $[a_0,1]$. Hence, the player continues to sample only $(q-m-k+1)$ attributes in $[a_0,1]$.
2. Conditional on all $(q-m-k+1)$ attributes being on $[a_{k-1}^r,1]$, $\tilde{a}$ continues to be optimal. This follows from the sample variance being pairwise separable in adjacent attributes.
3. Consider an alternative $\tilde{a}_k' \in [0,a_{k-1}^r]$ and its continuation play $(\tilde{a}_t')_{t=k}^q$. The variance of this subsample at stage k is strictly lower than what it was before sampling $\tilde{a}_{k-1}$ (due to strictly lower variance for $\tilde{a}_k'$), which in turn was strictly lower than the variance of continuation play $(\tilde{a}_t)_{t=k}^q$. Hence it is suboptimal to deviate to $\tilde{a}_k'$.

□

Proposition A.1 (Optimal sample as a function of a_0)**.** *Fix capacity q. There exist thresholds*

$$0 = a^0 < a^1 < \ldots < a^q < a^{q+1} = 1$$

such that if $a_0 \in (a^k, a^{k+1})$, where $k \in [0,q]$, any optimal sample corresponding to a_0 includes exactly k attributes smaller than a_0.

Proof for proposition A.1. Let $m(a_0)$ denote the number of attributes sampled in $[0,a_0]$. The proof proceeds in two steps. We first show that for any $a_0', a_0'' \in [0,1]$ such that $a_0' < a_0''$:

$$m(a_0'') \ge m(a_0').$$

Second, we show that for any $x \in \{0,1,\dots,q\}$, there exists $\tilde{a}_0$ such that $m(\tilde{a}_0) = x$.

(i) Suppose $m(a_0'') < m(a_0')$. Let $s^{\text{side}}(a_0,m) := (\sigma^{\text{side}}(a_0,m))^2$ denote the maximum variance of one of the two cells $[0,a_0]$ (with side$=\ell$), or $[a_0,1]$ (with side$=r$) when m attributes are allocated optimally to this interval. By the Envelope theorem,

$$\frac{\partial s^r(a_0,m)}{\partial a_0} = -\tau(a_1;(a_0,a_1))\left(\int_{a_0}^1 \omega(a)da + \int_{a_0}^{a_1} \frac{a_1-a}{a_1-a_0}\omega(a)da\right).$$

The term $-\tau(a_1;(a_0,a_1))$ is increasing in a_0; so is the term $\int_{a_0}^{a_1} \frac{a_1-a}{a_1-a_0}\omega(a)da$. Hence, $s^r(a_0,m)$ is convex in a_0. Moreover, because for any a_0 the optimal a_1 for capacity m is smaller than for $m-1$, for any $m' < m$,

$$\frac{\partial(s^r(a_0,m) - s^r(a_0,m'))}{\partial a_0} < 0.$$

By similar reasoning, in $[0,a_0]$,

$$\frac{\partial(s^\ell(a_0,m) - s^\ell(a_0,m'))}{\partial a_0} > 0.$$

By optimality of $m(a_0')$, $s^\ell(a_0', m(a_0')) - s^\ell(a_0', m(a_0'')) \geq s^r(a_0', q - m(a_0'')) - s^r(a_0', q - m(a_0'))$. But, by the monotonicity in a_0,

$$s^\ell(a_0'', m(a_0')) - s^\ell(a_0'', m(a_0'')) > s^\ell(a_0', m(a_0')) - s^\ell(a_0', m(a_0''))$$

$$s^r(a_0'', q - m(a_0'')) - s^r(a_0'', q - m(a_0')) < s^r(a_0', q - m(a_0'')) - s^r(a_0', q - m(a_0')).$$

This contradicts the optimality of $m(a_0'')$ for a_0''. Hence, $m(a_0'') \geq m(a_0')$.

(ii) Suppose, towards a contradiction, that there exists a_0 such that (i) the agent is indifferent between m and $m' > m+1$, (ii) the agent prefers m and m' to any other $m'' \neq m, m'$. Surpressing the argument a_0, the indifference condition is:

$$s^\ell(m) + s^r(q - m) = s^\ell(m') + s^r(q - m'),$$

whereas the optimality conditions for $m'' = m + 1 \in (m, m')$ give,

$$s^\ell(m+1) - s^\ell(m) < s^r(q - m) - s^r(q - m - 1),$$

$$s^\ell(m') - s^\ell(m+1) > s^r(q - m - 1) - s^r(q - m').$$

But by part (i) of the proof for proposition 3.3,

$$\frac{1}{m' - (m+1)} \left(s^\ell(m') - s^\ell(m+1) \right) < s^\ell(m+1) - s^\ell(m).$$

Combining this inequality with the optimality conditions above, we obtain that

$$\frac{1}{m' - (m+1)} \left(s^r(q - m - 1) - s^r(q - m') \right) < s^r(q - m) - s^r(q - m - 1).$$

But $q - m' < q - m - 1 < q - m$. For this to hold, it must be that

$$s^r(q - m) - s^r(q - m - 1) > s^r(q - m - k) - s^r(q - m - k - 1)$$

for $k = 1, \ldots, m' - m - 1$. This contradicts proposition 3.3. □

Lemma A.5 (Sufficient conditions for single-peakedness of $\sigma(a; \{a_0, a\})$. [31] *Suppose that for each $i = \mathscr{P}, \mathscr{A}$, the initial attribute a_0 is such that for any $a'' < a_0 < a'$,*

$$\omega_i(a') > \frac{1}{2(a' - a_0)} \int_{a_0}^{a'} \frac{s - a_0}{a' - a_0} \omega_i(s) ds,$$

$$\omega_i(a'') > \frac{1}{2(a_0 - a'')} \int_{a''}^{a_0} \frac{a_0 - s}{a_0 - a''} \omega_i(s) ds.$$

Then, $\sigma(a; \{a_0, a\})$ is single peaked on $[0, a_0]$ and $[a_0, 1]$.

Proof for lemma A.5. Suppose $a \in [a_0, 1]$. If a_0 is such that for any $a' > a_0$,

$$\omega_i(a') > \frac{1}{2(a' - a_0)} \int_{a_0}^{a'} \frac{s - a_0}{a' - a_0} \omega_i(s) ds,$$

[31] An example of a weight function that *violates* the single-peakedness condition for some value of a_0 is $\omega_i(a) = \frac{1}{10 + \frac{\sin^2(10)}{10}} (\sin(20a) + 10)$. The left and right single-crossing conditions fail for $a_0 = 0.63$ and $a_0 = 0.8$ respectively.

then the second derivative of $\sigma_i(a;\{a_0,a\})$ is strictly decreasing in a. Hence, σ_i is single-peaked. The proof for $a\in[0,a_0]$ follows a similar reasoning. □

Proof for proposition 3.6. (i) By the premise $\omega(a_0)=\tilde{\omega}(a_0)$, for any $a_1>a_0$, $\omega(a_1)\ge\tilde{\omega}(a_1)$ and for any $a_1'<a_0$, $\omega(a_1')\le\tilde{\omega}(a_1')$. It is immediate that for any sample of attributes $\mathbf{s}$ constrained in $[a_0,1]$, its variance $\sigma^2(\mathbf{s})$ is greater under ω than under $\tilde{\omega}$. For any sample of attributes on $[0,a_0]$, its variance $\sigma^2(\mathbf{s})$ is lower under ω than under $\tilde{\omega}$. As a result, sampling in $[a_0,1]$ becomes more attractive.

(ii) Let $\mathbf{s}=(a_1,\dots,a_{q+1})$ and $\tilde{\mathbf{s}}=(\tilde{a}_1,\dots,\tilde{a}_{q+1})$. Suppose that $1<k<q+1$; the following argument can be easily adapted for the cases when $k=1$ or $k=q+1$. Consider first a_1 and $\tilde{a}_1$. Dividing the FOC by $\omega(a_1)$ and using the MLR property, we obtain

$$0=\int_0^{a_1}\frac{\omega(a)}{\omega(a_1)}da-\int_{a_1}^{a_2}\frac{a_2-a}{a_2-a_1}\frac{\omega(a)}{\omega(a_1)}da\le\int_0^{a_1}\frac{\tilde{\omega}(a)}{\tilde{\omega}(a_1)}da-\int_{a_1}^{a_2}\frac{a_2-a}{a_2-a_1}\frac{\tilde{\omega}(a)}{\tilde{\omega}(a_1)}da.$$

By assumption 2, these functions cross zero once from below. But $\tilde{\omega}(a_1)>0$, and $\tilde{a}_2>a_2$. this means that $\int_0^{a_1}\tilde{\omega}(a)da-\int_{a_1}^{a_2}\frac{a_2-a}{a_2-a_1}\tilde{\omega}(a)da\ge 0$. Therefore, supposing that the second attribute is equal across the two samples $\tilde{a}_2=a_2$, the optimum $\tilde{a}_1$ on $[0,a_2]$ for $\tilde{\omega}$ is weakly smaller than a_1. The argument to follow shows that $\tilde{a}_2<a_2$, hence $\tilde{a}_1$ is even smaller than what it would be under $\tilde{a}_2=a_2$. The proof that $\tilde{a}_{q+1}<a_{q+1}$ in $[a_0,1]$ is similar and ommited.

Consider a_q and $\tilde{a}_q$. Dividing the FOC with respect to $\omega(a_q)$, we obtain

$$0=\int_{a_q}^{a_{q+1}}\frac{a_{q+1}-a}{a_{q+1}-a_q}\frac{\omega(a)}{\omega(a_q)}da-\int_{a_{q-1}}^{a_q}\frac{a-a_{q-1}}{a_q-a_{q-1}}\frac{\omega(a)}{\omega(a_q)}da\ge\int_{a_q}^{a_{q+1}}\frac{a_{q+1}-a}{a_{q+1}-a_q}\frac{\tilde{\omega}(a)}{\tilde{\omega}(a_q)}da-\int_{a_{q-1}}^{a_q}\frac{a-a_{q-1}}{a_q-a_{q-1}}\frac{\tilde{\omega}(a)}{\tilde{\omega}(a_q)}da.$$

The inequality follows from the MLR property, keeping a_{q-1} and a_{q+1} fixed in this comparison. By assumption 2, this FOC (as a function of a_q) crosses zero only once from above. Hence, $\tilde{a}_q<a_q$, supposing that samples $\tilde{\mathbf{s}}$ and $\mathbf{s}$ have the same a_{q-1} and a_{q+1}. But the argument above established that for any $\tilde{a}_q=a_q$, $\tilde{a}_{q+1}<a_{q+1}$; hence

$$\int_{a_q}^{a_{q+1}}\frac{a_{q+1}-a}{a_{q+1}-a_q}\frac{\tilde{\omega}(a)}{\tilde{\omega}(a_q)}da>\int_{a_q}^{\tilde{a}_{q+1}}\frac{\tilde{a}_{q+1}-a}{\tilde{a}_{q+1}-a_q}\frac{\tilde{\omega}(a)}{\tilde{\omega}(a_q)}da.$$

Hence, $\tilde{a}_{q+1}>a_{q+1}$ and $\tilde{a}_q>a_q$, keeping a_{q-1} fixed across the two samples.

By following a recursive argument, it follows that for any $q+1>m>k$, $\tilde{a}_m<a_m$. Similarly, for any $1<m<k$, $\tilde{a}_m<a_m$. This concludes the proof. □

B Proofs for sections 4 and 5

Proof for lemma 4.1. First, note that $v^i(\mathbf{s})\sim\mathcal{N}\left(v_0^i,{\sigma_i}^2(\mathbf{s})\right)$ are each normal random variables. Moreover, $v^{\mathcal{P}}(\mathbf{s})$ and $v^{\mathcal{A}}(\mathbf{s})$ are jointly normal, because for any $i\in\{\mathcal{P},\mathcal{A}\}$, $v^i(\mathbf{s})$ can be expressed as a linear combination of the form

$$c_0^iB_0+\sum_{j=1}^{n_\ell}c_j^iZ_j^\ell+\sum_{j=1}^{n_r}d_j^iZ_j^r+\gamma_i,$$

where $c_0^i,(c_j^i),(d_j^i)$ are player-specific constants and Z_i^ℓ,Z_i^r are independent normal random shocks. Constant γ_i captures the drift term, independent of any realizations. Here, we are invoking the fact that for any $a_j^k\in\mathbf{s}$, the random variable $B(a_j^k)$ can be decomposed into a weighted sum of B_0 and a sequence of independent normal random variables $Z_1^k,Z_2^k,\dots,Z_j^k$, $k\in\{r,\ell\}$. The covariance between

$v^{\mathscr{P}}(\mathbf{s})$ and $v^{\mathscr{A}}(\mathbf{s})$ is:

$$\begin{aligned}
\operatorname{cov}(v^{\mathscr{P}}(\mathbf{s}), v^{\mathscr{A}}(\mathbf{s})) &= \operatorname{cov}\left(\sum_{i=1}^{n_\ell}\tau_{\mathscr{P}}(a_i^\ell;\mathbf{s})B(a_i^\ell) + \sum_{i=1}^{n_r}\tau_{\mathscr{P}}(a_i^r;\mathbf{s})B(a_i^r), \sum_{i=1}^{n_\ell}\tau_{\mathscr{A}}(a_i^\ell;\mathbf{s})B(a_i^\ell) + \sum_{i=1}^{n_r}\tau_{\mathscr{A}}(a_i^r;\mathbf{s})B(a_i^r)\right) \\
&= \operatorname{cov}\left(\sum_{i=1}^{n_\ell}\tau_{\mathscr{P}}(a_i^\ell;\mathbf{s})B(a_i^\ell), \sum_{i=1}^{n_\ell}\tau_{\mathscr{A}}(a_i^\ell;\mathbf{s})B(a_i^\ell)\right) + \operatorname{cov}\left(\sum_{i=1}^{n_r}\tau_{\mathscr{P}}(a_i^r;\mathbf{s})B(a_i^r), \sum_{i=1}^{n_r}\tau_{\mathscr{A}}(a_i^r;\mathbf{s})B(a_i^r)\right) \\
&= \sum_{k=1}^{n_\ell}\left(\sum_{i=k}^{n_\ell}\tau_{\mathscr{P}}(a_k^\ell;\mathbf{s})\right)\left(\sum_{i=k}^{n_\ell}\tau_{\mathscr{A}}(a_k^\ell;\mathbf{s})\right)(a_{k-1}^\ell - a_k^\ell) + \sum_{k=1}^{n_r}\left(\sum_{i=k}^{n_r}\tau_{\mathscr{P}}(a_k^r;\mathbf{s})\right)\left(\sum_{i=k}^{n_r}\tau_{\mathscr{A}}(a_k^r;\mathbf{s})\right)(a_k^r - a_{k-1}^r) \\
&= \sum_{k=1}^{n_\ell}\tau_{\mathscr{P}}(a_k^\ell;(a_k^\ell,a_{k-1}^\ell))\tau_{\mathscr{A}}(a_k^\ell;(a_k^\ell,a_{k-1}^\ell))(a_{k-1}^\ell - a_k^\ell) + \\
&\qquad \sum_{k=1}^{n_r}\tau_{\mathscr{P}}(a_k^r;(a_k^r,a_{k-1}^r))\tau_{\mathscr{A}}(a_k^r;(a_k^r,a_{k-1}^r))(a_k^r - a_{k-1}^r).
\end{aligned}$$

The first line has ommitted the terms $\tau_i(a_0;\mathbf{s})B_0$ by using the fact that $\operatorname{cov}(c+X, c+Y) = \operatorname{cov}(X,Y)$ for any two random variables X, Y and constant c. The second line has used the fact that for any j,k, $\operatorname{cov}(B(a_j^\ell), B(a_k^r)) = 0$. The third equality has used a number of observations. First, for any $j < k$ and $s \in \{\ell, r\}$, $\operatorname{cov}(B(a_j^s), B(a_k^s)) = a_j - a_0$. Second, it is useful to break down terms of the form: $a_j^s - a_0 = \sum_{i=1}^{j}(a_i^s - a_{i-1}^s)$. After recombining terms, we get

$$\begin{aligned}
\operatorname{cov}\left(\sum_{i=1}^{n_\ell}\tau_{\mathscr{P}}(a_i^\ell;\mathbf{s})B(a_i^\ell), \sum_{i=1}^{n_\ell}\tau_{\mathscr{A}}(a_i^\ell;\mathbf{s})B(a_i^\ell)\right) &= \left(\sum_{i=1}^{n_\ell}\tau_{\mathscr{P}}(a_1^\ell;\mathbf{s})\right)\left(\sum_{i=1}^{n_\ell}\tau_{\mathscr{A}}(a_1^\ell;\mathbf{s})\right)(a_0 - a_1^\ell) + \\
&\left(\sum_{i=2}^{n_\ell}\tau_{\mathscr{P}}(a_2^\ell;\mathbf{s})\right)\left(\sum_{i=2}^{n_\ell}\tau_{\mathscr{A}}(a_2^\ell;\mathbf{s})\right)(a_1^\ell - a_2^\ell) + \ldots + \tau_{\mathscr{P}}(a_{n_\ell}^\ell;\mathbf{s})\tau_{\mathscr{A}}(a_{n_\ell}^\ell;\mathbf{s}).
\end{aligned}$$

The fourth equality uses the fact that $\sum_{i=k}^{n_s}\tau_i(a_k^s;\mathbf{s}) = \tau_i\left(a_k^s;(a_k^s,a_{k-1}^s)\right)$.

Claim 2.

$$f(v^{\mathscr{A}}(\mathbf{s}) \mid v^{\mathscr{P}}(\mathbf{s}) \geq 0) = \frac{\phi\left(\frac{v_{\mathscr{A}}(\mathbf{s}) - v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}(\mathbf{s})}\right)}{\sigma_{\mathscr{A}}(\mathbf{s})\Phi\left(\frac{v_{\mathscr{P}}(\mathbf{s}) - v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right)}\Phi\left(\frac{v_0^{\mathscr{P}}(\mathbf{s}) + \rho(\mathbf{s})\frac{\sigma_{\mathscr{P}}(\mathbf{s})}{\sigma_{\mathscr{A}}(\mathbf{s})}\left(v_{\mathscr{A}}(\mathbf{s}) - v_0^{\mathscr{A}}\right)}{\sigma_{\mathscr{P}}(\mathbf{s})\sqrt{1-\rho(\mathbf{s})^2}}\right),$$

where $\sigma_i \equiv \sigma_i(\mathbf{s})$ and $\rho(\mathbf{s}) \equiv \rho(v^{\mathscr{P}}(\mathbf{s}), v^{\mathscr{A}}(\mathbf{s}))$ denotes the correlation between $v^{\mathscr{P}}(\mathbf{s})$ and $v^{\mathscr{A}}(\mathbf{s})$.

Proof. Let x_1, x_2 be jointly normal with means μ_1, μ_2, variances σ_1^2, σ_2^2 and covariance σ_{12}. Let f_1, f_2 and F_1, F_2 denote the pdf and cdf, respectively, of x_1 and x_2. Let $f(\cdot,\cdot)$ and $f(\cdot \mid \cdot)$ denote the joint pdf and the conditional pdf respectively. Then,

$$\begin{aligned}
f(x_1 \mid x_2 \geq 0) &= \frac{1}{1 - F_2(0)}\Pr(x_2 \geq 0) f(x_1 \mid x_2 \geq 0) \\
&= \frac{1}{1 - F_2(0)}\int_0^\infty f(x_1, x_2)dx_2 \\
&= \frac{1}{1 - F_2(0)}\int_0^\infty f(x_2 \mid x_1) f_1(x_1)dx_2 \\
&= \frac{f_1(x_1)}{1 - F_2(0)}\int_0^\infty f(x_2 \mid x_1)dx_2 \\
&= \frac{f_1(x_1)}{1 - F_2(0)}(1 - F_{x_2|x_1}(0)).
\end{aligned}$$

The first line multiplies and divides by $\Pr(x_2 \geq 0)$. The second line rewrites $\Pr(x_2 \geq 0) f(x_1 \mid x_2 \geq 0)$ using the joint density. The third line uses the fact that $f(x_1, x_2) = f(x_2 \mid x_1) f_1(x_1)$. The last two lines use the conditional distribution of $x_2 \mid x_1$. But,

$$x_2 \mid x_1 \sim \mathcal{N}\left(\mu_2 + \rho\frac{\sigma_2}{\sigma_1}(x_1 - \mu_1), (1-\rho^2)\sigma_2^2\right)$$

and $\rho = \frac{\sigma_{12}}{\sigma_1\sigma_2}$. Therefore, we can insert the expression for $F_{x_2|x_1}$ to obtain:

$$f(x_1 \mid x_2 \geq 0) = \frac{f_1(x_1)}{1-F_2(0)}\left(1 - \Phi\left(-\frac{\mu_2 + \rho\frac{\sigma_2}{\sigma_1}(x_1-\mu_1)}{\sigma_2\sqrt{1-\rho^2}}\right)\right)$$

Switching back to our variables of interest, let $x_1 := v^{\mathscr{A}}(\mathbf{s}) \sim \mathcal{N}(v_0^{\mathscr{A}}, \sigma_{\mathscr{A}}(\mathbf{s}))$ and $x_2 := v^{\mathscr{P}}(\mathbf{s}) \sim \mathcal{N}(v_0^{\mathscr{P}}, \sigma_{\mathscr{P}}(\mathbf{s}))$. Therefore,

$$\begin{aligned} f(v^{\mathscr{A}}(\mathbf{s})|v^{\mathscr{P}}(\mathbf{s}) \geq 0) &= \frac{\phi\left(\frac{v_{\mathscr{A}}(\mathbf{s}) - v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}(\mathbf{s})}\right)}{\sigma_{\mathscr{A}}(\mathbf{s})\left(1-\Phi\left(-\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right)\right)}\left(1-\Phi\left(-\frac{v_0^{\mathscr{P}} + \rho(\mathbf{s})\frac{\sigma_{\mathscr{P}}(\mathbf{s})}{\sigma_{\mathscr{A}}(\mathbf{s})}(v^{\mathscr{A}}(\mathbf{s}) - v_0^{\mathscr{A}})}{\sigma_{\mathscr{P}}(\mathbf{s})\sqrt{1-\rho(\mathbf{s})^2}}\right)\right) \\ &= \frac{\phi\left(\frac{v^{\mathscr{A}}(\mathbf{s}) - v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}(\mathbf{s})}\right)}{\sigma_{\mathscr{A}}(\mathbf{s})\Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right)}\Phi\left(\frac{v_0^{\mathscr{A}} + \rho(\mathbf{s})\frac{\sigma_{\mathscr{P}}(\mathbf{s})}{\sigma_{\mathscr{A}}(\mathbf{s})}(v^{\mathscr{A}}(\mathbf{s}) - v_0^{\mathscr{A}})}{\sigma_{\mathscr{P}}(\mathbf{s})\sqrt{1-\rho(\mathbf{s})^2}}\right). \end{aligned}$$

□

Using the claim, the objective simplifies to the following integral:

$$\begin{aligned} \Pr(v^{\mathscr{P}}(\mathbf{s}) \geq 0)\mathbb{E}[v^{\mathscr{A}}(\mathbf{s}) \mid v^{\mathscr{P}}(\mathbf{s}) \geq 0] &= \Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right)\int_{-\infty}^{\infty} v^{\mathscr{A}}(\mathbf{s}) f(v^{\mathscr{A}}(\mathbf{s}) \mid v^{\mathscr{P}}(\mathbf{s}) \geq 0) dv^{\mathscr{A}}(\mathbf{s}) \\ &= \int_{-\infty}^{\infty} v^{\mathscr{A}}(\mathbf{s})\frac{1}{\sigma_{\mathscr{A}}(\mathbf{s})}\phi\left(\frac{v^{\mathscr{A}}(\mathbf{s}) - v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}(\mathbf{s})}\right)\Phi\left(\frac{v_0^{\mathscr{P}} + \rho(\mathbf{s})\frac{\sigma_{\mathscr{P}}(\mathbf{s})}{\sigma_{\mathscr{A}}(\mathbf{s})}(v^{\mathscr{A}}(\mathbf{s}) - v_0^{\mathscr{A}})}{\sqrt{(1-\rho(\mathbf{s})^2)(\sigma_{\mathscr{P}}(\mathbf{s}))^2}}\right) dv^{\mathscr{A}}(\mathbf{s}). \end{aligned}$$

Let $x \equiv \frac{v^{\mathscr{A}}(\mathbf{s}) - v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}(\mathbf{s})}$; then, $v^{\mathscr{A}}(\mathbf{s}) = x\sigma_{\mathscr{A}}(\mathbf{s}) + v_0^{\mathscr{A}}$. Substituting this into the expression above, we get:

$$\Pr(v^{\mathscr{P}}(\mathbf{s}) \geq 0)\mathbb{E}[v^{\mathscr{A}}(\mathbf{s}) \mid v^{\mathscr{P}}(\mathbf{s}) \geq 0] = \int_{-\infty}^{\infty}(x\sigma_{\mathscr{A}}(\mathbf{s}) + v_0^{\mathscr{A}})\phi(x)\Phi\left(\frac{v_0^{\mathscr{P}} + \rho(\mathbf{s})\sigma_{\mathscr{P}}(\mathbf{s})x}{\sqrt{(1-\rho(\mathbf{s})^2)(\sigma_{\mathscr{P}}(\mathbf{s}))^2}}\right)dx \tag{24}$$

From **?**, we have the following identities (respectively, numbered 10,010.8 and 10,011.1 therein):

$$\int_{-\infty}^{\infty}\phi(x)\Phi(a+bx)dx = \Phi\left(\frac{a}{\sqrt{1+b^2}}\right)$$

$$\int_{-\infty}^{\infty}x\phi(x)\Phi(a+bx)dx = \frac{b}{\sqrt{1+b^2}}\phi\left(\frac{a}{\sqrt{1+b^2}}\right).$$

In equation 24, $a = \frac{v_0^{\mathscr{A}}}{\sigma_{\mathscr{P}}(\mathbf{s})\sqrt{1-\rho(\mathbf{s})^2}}$ and $b = \frac{\rho(\mathbf{s})}{\sqrt{1-\rho(\mathbf{s})^2}}$. Hence, expression 24 simplifies to:

$$\Pr(v^{\mathscr{P}}(\mathbf{s}) \geq 0)\mathbb{E}[v^{\mathscr{A}}(\mathbf{s}) \mid v^{\mathscr{P}}(\mathbf{s}) \geq 0] = v_0^{\mathscr{A}}\Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right) + \rho(\mathbf{s})\sigma_{\mathscr{A}}(s)\phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right).$$

Using the fact that $\rho = \frac{\text{cov}(\mathbf{s})}{\sigma_{\mathscr{P}}(\mathbf{s})\sigma_{\mathscr{A}}(\mathbf{s})}$,

$$\Pr(v^{\mathscr{P}}(\mathbf{s}) \geq 0)\mathbb{E}[v^{\mathscr{A}}(\mathbf{s}) \mid v^{\mathscr{P}}(\mathbf{s}) \geq 0] = v_0^{\mathscr{A}}\Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right) + \frac{\text{cov}(\mathbf{s})}{\sigma_{\mathscr{P}}(\mathbf{s})}\phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right).$$

□

Proof for proposition 4.2.

Let $P(\mathbf{s}; v_0^{\mathscr{P}}, v_0^{\mathscr{A}})$ denote the payoff of $\mathscr{A}$ from sampling $\mathbf{s}$ given $(v_0^{\mathscr{P}}, v_0^{\mathscr{A}}, a_0)$. Let $\mathbf{s}^*$ be such that

$$\mathbf{s}^* \in \arg\max_{\mathbf{s}\in\mathscr{S}_{q+1}(a_0)} P(\mathbf{s}; v_0^{\mathscr{P}}, v_0^{\mathscr{A}}).$$

We want to show that $\mathbf{s}^* \in \arg\max_{\mathbf{s}\in\mathscr{S}_{q+1}(a_0)} P(\mathbf{s}; -v_0^{\mathscr{P}}, -v_0^{\mathscr{A}})$ as well. The ex-ante expected value $v_0^i(B_0,\mu) = B_0 + \mu\int_0^1 \omega_i(a)(a-a_0)da$ is such that $v_0^i(B_0,\mu) = -v_0^i(-B_0,-\mu)$. Hence,

$$P(\mathbf{s}; -v_0^{\mathscr{P}}, -v_0^{\mathscr{A}}) = (-v_0^{\mathscr{A}})\Phi\left(-\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right) + \sigma_{\mathscr{A}}(\mathbf{s})\phi\left(-\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(\mathbf{s})}\right) = -v_0^{\mathscr{A}} + P(\mathbf{s}; v_0^{\mathscr{A}}, v_0^{\mathscr{P}}).$$

Hence, any $\mathbf{s}^*$ that maximizes $P(\mathbf{s}; v_0^{\mathscr{P}}, v_0^{\mathscr{A}})$ also maximizes $P(\mathbf{s}; -v_0^{\mathscr{P}}, -v_0^{\mathscr{A}})$. □

Proof for lemma 4.1. (i) Consider the objective of the agent:

$$P(a; v_0^{\mathscr{P}}, v_0^{\mathscr{A}}) = v_0^{\mathscr{A}}\Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a)}\right) + \sigma_{\mathscr{A}}(a)\phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a)}\right).$$

For $v_0^{\mathscr{P}} = 0$, this objective reduces to $P(a; 0, v_0^{\mathscr{A}}) = \frac{v_0^{\mathscr{A}}}{2} + \frac{\sigma_{\mathscr{A}}(a)}{\sqrt{2\pi}}$. This objective is maximized by $a^* \in \arg\max_a \sigma_{\mathscr{A}}(a)$. The optimal attribute a^* is independent of $v_0^{\mathscr{A}}$.

(ii) When $v_0^{\mathscr{A}} = 0$,

$$P(a; v_0^{\mathscr{P}}, 0) = \sigma_{\mathscr{A}}(a)\phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a)}\right).$$

By the Envelope Theorem, the value $P(a^*; v_0^{\mathscr{P}}, 0)$ is strictly decreasing in v_0^p if $v_0^{\mathscr{P}} > 0$ and strictly increasing otherwise; hence it is strictly decreasing in $|v_0^{\mathscr{P}}|$. The first-order condition with respect to a gives us:

$$\frac{\partial\sigma_{\mathscr{A}}(a)}{\partial a} + \frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}^3(a)}(v_0^{\mathscr{P}}\sigma_{\mathscr{A}}(a) - v_0^{\mathscr{A}}\sigma_{\mathscr{P}}(a))\frac{\partial\sigma_{\mathscr{P}}(a)}{\partial a} = 0.$$

For $v_0^{\mathscr{A}} = 0$, any a that satisfies the FOC is one at which $\frac{\partial\sigma_{\mathscr{A}}(a)}{\partial a}$ and $\frac{\partial\sigma_{\mathscr{P}}(a)}{\partial a}$ have opposite signs. Hence, any local optimum a is a compromise. If $v_0^{\mathscr{P}} = 0$, the FOC simplifies to

$$\frac{\partial\sigma_{\mathscr{A}}(a)}{\partial a} = 0,$$

hence by the assumed single crossing condition, $a^* = a_{\mathscr{A}}$. As $v_0^{\mathscr{P}} \to \infty$, the second term $\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}^3(a)}(v_0^{\mathscr{P}}\sigma_{\mathscr{A}}(a) - v_0^{\mathscr{A}}\sigma_{\mathscr{P}}(a))\frac{\partial\sigma_{\mathscr{P}}(a)}{\partial a}$ dominates the RHS of the FOC. For any a within $(a_{\mathscr{P}}^s - \varepsilon, a_{\mathscr{P}}^s)$, there exists a $v_0^{\mathscr{P}}$ sufficiently large so that $a_s^* = a$.

(iii) The value function of the agent is

$$V(v_0^{\mathscr{P}}, v_0^{\mathscr{A}}) = \begin{cases} \frac{v_0^{\mathscr{A}}}{2} + \frac{\sigma_{\mathscr{A}}(a^*)}{\sqrt{2\pi}} & \text{if } v_0^{\mathscr{P}} = 0, \\ \sigma_{\mathscr{A}}(a^*)\phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a^*)}\right) & \text{if } v_0^{\mathscr{A}} = 0. \end{cases}$$

The difference $V(0,x) - V(x,0)$ is increasing in x for $x \geq 0$ and it is exactly zero at $x = 0$. Hence, for $x > 0$, the agent prefers $v_0^{\mathscr{P}} = 0$ and $v_0^{\mathscr{A}} = x$ to the opposite case. □

Proof of proposition 4.3. The FOC with respect to a is:

$$\frac{\partial \sigma_{\mathscr{A}}(a)}{\partial a} + \frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}^3(a)} \left(v_0^{\mathscr{P}} \sigma_{\mathscr{A}}(a) - v_0^{\mathscr{A}} \sigma_{\mathscr{P}}(a)\right) \frac{\partial \sigma_{\mathscr{P}}(a)}{\partial a} = 0.$$

(i) If $v_0^{\mathscr{A}}$ and $v_0^{\mathscr{P}}$ have opposite signs, the term $v_0^{\mathscr{P}} \left(v_0^{\mathscr{P}} \sigma_{\mathscr{A}}(a) - v_0^{\mathscr{A}} \sigma_{\mathscr{P}}(a)\right) > 0$. Therefore, any a^* that satisfies the FOC is such that $\left.\frac{\partial \sigma_{\mathscr{A}}(a)}{\partial a}\right|_{a^*}$ and $\left.\frac{\partial \sigma_{\mathscr{P}}(a)}{\partial a}\right|_{a^*}$ hold opposite signs. By the single-crossing condition, this implies that a^* is a local compromise.

(ii) If a^* is a compromise, $\left.\frac{\partial \sigma_{\mathscr{A}}(a)}{\partial a}\right|_{a^*}$ and $\left.\frac{\partial \sigma_{\mathscr{P}}(a)}{\partial a}\right|_{a^*}$ hold opposite signs. For a^* to satisfy FOC, it is necessary that $v_0^{\mathscr{P}} \left(v_0^{\mathscr{P}} \sigma_{\mathscr{A}}(a^*) - v_0^{\mathscr{A}} \sigma_{\mathscr{P}}(a^*)\right) > 0$. This yields the condition in the statement. Conversely, suppose that for a given a^*, $v_0^{\mathscr{P}} \left(v_0^{\mathscr{P}} \sigma_{\mathscr{A}}(a^*) - v_0^{\mathscr{A}} \sigma_{\mathscr{P}}(a^*)\right) > 0$. For a^* to satisfy FOC, it is necessary that $\left.\frac{\partial \sigma_{\mathscr{A}}(a)}{\partial a}\right|_{a^*}$ and $\left.\frac{\partial \sigma_{\mathscr{P}}(a)}{\partial a}\right|_{a^*}$ hold opposite signs. Hence, such an attribute a^* must be a compromise. □

Proof of proposition 4.4. Without loss, fix $\omega_{\mathscr{A}}$ and $\omega_{\mathscr{P}}$ such that $a_{\mathscr{A}} = a_{\mathscr{P}}^r$ in $[a_0, 1]$ and such that $a_{\mathscr{A}}^r > a_{\mathscr{P}}^r$. Let $K_i := \mu \int_0^1 \omega_i(a)(a - a_0)da$. By the boundedness of ω_i, the ratio $\frac{\tau_{\mathscr{A}}}{\tau_{\mathscr{P}}}$ is bounded over $a \in [0,1]$.

For B_0 sufficiently close to $-\mu K_{\mathscr{P}}$, the ratio $\frac{v_0^{\mathscr{P}}}{v_0^{\mathscr{A}}}$ is sufficiently close to zero, therefore for all $a \in [0,1]$:

$$\frac{v_0^{\mathscr{P}}}{v_0^{\mathscr{A}}} < \frac{\tau_{\mathscr{A}}(a)}{\tau_{\mathscr{P}}(a)}.$$

By part (ii) of corollary 4.2, this condition guarantees that the optimal attribute either overshoots or undershoots. By lemma 4.1, for $v_0^{\mathscr{P}} = 0$, i.e. for $B_0 = -\mu K_{\mathscr{P}}$, the optimal attribute is exactly $a_{\mathscr{A}}$. For $v_0^{\mathscr{P}}$ sufficiently close to zero, the optimal attribute is in an ϵ-neighborhood of $a^{\mathscr{A}}$. Condition (c) in the statement guarantees that the optimal attribute overshoots rather than undershoots. □

Proof for proposition 4.5.

1. The proof proceeds in two steps. First, we show that there exists a sufficiently large $|B_0|$ above which the constrained-optimal attribute in $[0, a_0]$ is equal to a_0. Second, we show that for any $(v_0^{\mathscr{P}}, v_0^{\mathscr{A}})$, sampling in the compromise region guarantees to the agent a payoff strictly greater than $v_0^{\mathscr{A}}$ (the payoff from sampling a_0).

(i) Suppose, without loss, that $-\mu K_{\mathscr{A}} \geq -\mu K_{\mathscr{P}}$. Suppose, towards a contradiction, that the payoff achieved by sampling in $[0, a_0]$ is weakly greater than the adoption payoff $v_0^{\mathscr{A}} > 0$ for any $B_0 > -\mu K_{\mathscr{A}}$. That is, for any B_0 and any $a \in [0, a_0)$,

$$v_0^{\mathscr{A}} \Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a)}\right) + \sigma_{\mathscr{A}}(a)\phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a)}\right) \geq v_0^{\mathscr{A}}$$

$$\Leftrightarrow \frac{\phi(x)}{1 - \Phi(x)} \geq x \frac{v_0^{\mathscr{A}}}{v_0^{\mathscr{P}}} \frac{\sigma_{\mathscr{P}}(a)}{\sigma_{\mathscr{A}}(a)} \tag{25}$$

where $x := v_0^{\mathscr{P}}/\sigma_{\mathscr{P}}(a)$. As $B_0 \to +\infty$, $x \to +\infty$, while $v_0^{\mathscr{P}}/v_0^{\mathscr{A}}$ converges to 1 from above. For a sufficiently large B_0, for any $a \in [0, a_0)$, the RHS is strictly greater than x. The LHS is the inverse Mill's ratio for the standard normal distribution.

Claim 3. $\frac{\phi(x)}{1-\Phi(x)} \to x$ *as* $x \to +\infty$.

Proof. Using L'Hôpital's rule:

$$\lim_{x\to\infty} \frac{\phi(x)}{1-\Phi(x)} = -\frac{\phi'(x)}{\phi(x)} = -\frac{x\phi(x)}{\phi(x)} = x,$$

where the second-to-last equality uses the property of the standard normal pdf: $\phi'(x) = x\phi(x)$. □

For a sufficiently large x, the LHS of inequality 25 approaches x, while the RHS is strictly greater than x. Hence, this inequality is violated for sufficiently large x. For sufficiently good initial evidence, the payoff from sampling any $a \in [0, a_0)$ is stricly worse than no sampling. By a similar reasoning, we want to show that as $B_0 \to -\infty$, no sampling is strictly preferred to sampling in $[0, a_0)$. By contradiction, suppose that for any $B_0 < -\mu K_{\mathscr{A}}$ and any $a \in [0, a_0)$,

$$v_0^{\mathscr{A}} \Phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a)}\right) + \sigma_{\mathscr{A}}(a)\phi\left(\frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a)}\right) \geq 0,$$

$$\Leftrightarrow \frac{\phi(-x)}{1-\Phi(-x)} \geq -x\frac{v_0^{\mathscr{A}}}{v_0^{\mathscr{P}}}\frac{\sigma_{\mathscr{P}}(a)}{\sigma_{\mathscr{A}}(a)},$$

where x is as before. As $x \to -\infty$, the LHS $\frac{\phi(-x)}{1-\Phi(-x)}$ approaches $-x$, whereas RHS approaches $-x\sigma_{\mathscr{P}}(a)/\sigma_{\mathscr{A}}(a) > -x$. Contradiction.

(ii) In $[a_0, 1]$, the agent's payoff is strictly increasing in both $\sigma_{\mathscr{A}}$ and $\sigma_{\mathscr{P}}$. Therefore, for any $(v_0^{\mathscr{P}}, v_0^{\mathscr{A}})$, the optimum in $[a_0, 1]$ is interior to the two optima, and hence it guarantees a payoff greater than $v_0^{\mathscr{A}}$, which is what the agent obtains by not sampling any attributes.

2. Consider the FOC of the agent's objective with respect to a.

$$\frac{\partial\sigma_{\mathscr{A}}(a)}{\partial a} + \frac{v_0^{\mathscr{P}}}{\sigma^3_{\mathscr{P}}(a)}\left(v_0^{\mathscr{P}}\sigma_{\mathscr{A}}(a) - v_0^{\mathscr{A}}\sigma_{\mathscr{P}}(a)\right)\frac{\partial\sigma_{\mathscr{P}}(a)}{\partial a} = 0.$$

Suppose, without loss, that $a^r_{\mathscr{A}} < a^r_{\mathscr{P}}$. For any $a \in (a_0, 1]$, $T(B_0, a) := v_0^{\mathscr{P}}\left(v_0^{\mathscr{P}}\sigma_{\mathscr{A}}(a) - v_0^{\mathscr{A}}\sigma_{\mathscr{P}}(a)\right) > 0$. For any two $B_0' > B_0 > -\mu\int_0^1 (a-a_0)\omega_{\mathscr{P}}(a)da$ and for any $a_0 < a \leq 1$, $T(B_0', a) > T(B_0, a) > 0$ because $\sigma_{\mathscr{P}}(a) < \sigma_A(a)$ and $v_0^i = B_0 + \mu\int_0^1 (a-a_0)\omega_i(a)da$. Suppose, by way of contradiction, that $a^*(B_0') \leq a^*(B_0)$; let these constrained optima be denoted, respectively, by a' and a. Therefore, $\sigma_{\mathscr{P}}(a') \leq \sigma_{\mathscr{P}}(a)$, and

$$\frac{T(B_0', a')}{\sigma^3_{\mathscr{P}}(a')} > \frac{T(B_0, a)}{\sigma^3_{\mathscr{P}}(a)}.$$

By $a' \leq a$, it also follows that $\frac{\partial\sigma_{\mathscr{A}}(a)}{\partial a}|_a \leq \frac{\partial\sigma_{\mathscr{A}}(a)}{\partial a}|_{a'} < 0$ and $\frac{\partial\sigma_{\mathscr{P}}(a)}{\partial a}|_{a'} \geq \frac{\partial\sigma_{\mathscr{P}}(a)}{\partial a}|_a > 0$. Therefore,

$$\frac{\partial\sigma_{\mathscr{A}}(a')}{\partial a} + \frac{T(B_0', a')}{\sigma^3_{\mathscr{P}}(a')}\frac{\partial\sigma_{\mathscr{P}}(a')}{\partial a} > \frac{\partial\sigma_{\mathscr{A}}(a)}{\partial a} + \frac{T(B_0, a)}{\sigma^3_{\mathscr{P}}(a)}\frac{\partial\sigma_{\mathscr{P}}(a)}{\partial a} = 0.$$

This contradicts the optimality of a' for B_0'. Therefore, $a^*(B_0') > a^*(B_0)$ for any $B_0' > B_0$. By our previous analysis, for any B_0, $a^*(B_0) \in [a^r_{\mathscr{A}}, a^r_{\mathscr{P}}]$. Hence, for any $\epsilon \in (0, a^r_{\mathscr{P}} - a^r_{\mathscr{A}})$, there exists B_0 such that for all $B_0' > B_0$, $a^r_{\mathscr{P}} - a^*(B_0') < \epsilon$.

Similarly, for any $-\mu \int_0^1 (a - a_0)\omega_{\mathscr{P}}(a)da > B_0' > B_0$, $a^*(B_0') < a^*(B_0)$. The proof is similar, so we omit the details. Therefore, we conclude that the constrained optimum approaches $a^r_{\mathscr{P}}$ as $B_0 \to \pm\infty$.

□

Proof of proposition 4.6. Suppose there exists some $v_0 \in \mathbb{R}$ such that $\mathscr{A}$ prefers to disclose no additional attributes. That is, no sampling is preferred to the discovery of any $a \neq a_0$. This means that for any $a \neq a_0$,

$$v_0 \Phi\left(\frac{v_0}{\sigma^{\mathscr{P}}(a)}\right) + \sigma^{\mathscr{A}}(a)\phi\left(\frac{v_0}{\sigma^{\mathscr{P}}(a)}\right) \leq \max\{0, v_0\}.$$

The RHS term captures the fact that absent any additional attribute discovery, $\mathscr{P}$ approves iff $v_0 \geq 0$. Because $\sigma^{\mathscr{A}}(a) \neq 0$ for any $a \neq a_0$, this simplifies to

$$\begin{cases} \frac{v_0}{\sigma^{\mathscr{A}}(a)} \geq \lambda\left(\frac{v_0}{\sigma^{\mathscr{P}}(a)}\right) & \text{for } v_0 \geq 0 \\ -\frac{v_0}{\sigma^{\mathscr{A}}(a)} \geq \lambda\left(-\frac{v_0}{\sigma^{\mathscr{P}}(a)}\right) & \text{for } v_0 < 0, \end{cases}$$

where $\lambda(\cdot)$ denotes the inverse Mill's ratio (or hazard ratio) for the standard normal distribution. It is a known property of this ratio that $\lambda(x) > x$ for any x.[32] It must therefore be that for any $a \neq a_0$,

$$\sigma^{\mathscr{A}}(a) < \sigma^{\mathscr{P}}(a) \Leftrightarrow \tau^{\mathscr{A}}(a) < \tau^{\mathscr{P}}(a).$$

Claim 4. *If $\tau^{\mathscr{A}}(a) < \tau^{\mathscr{P}}(a)$ for all $a \in [0,1] \setminus \{a_0\}$, then*

$$\int_{a_0}^1 (\omega_{\mathscr{A}}(a) - \omega_{\mathscr{P}}(a))da = \int_0^{a_0} (\omega_{\mathscr{A}}(a) - \omega_{\mathscr{P}}(a))da = 0.$$

Proof. Let $\bar{\alpha} \equiv \int_{a_0}^1 (\omega_{\mathscr{A}}(a) - \omega_{\mathscr{P}}(a))da$ and $\underline{\alpha} \equiv \int_0^{a_0} (\omega_{\mathscr{A}}(a) - \omega_{\mathscr{P}}(a))da$. Suppose first that $\bar{\alpha} < 0$. But, because $\int_0^1 \omega_i(a)da = \Omega$ for $i = \mathscr{P}, \mathscr{A}$, it follows that $\bar{\alpha} = -\underline{\alpha}$. Hence, $\underline{\alpha} > 0$. By continuity of the weight functions, there exist $\varepsilon_1, \varepsilon_2$ such that

$$\int_{a_0+\varepsilon_1}^1 (\omega_{\mathscr{A}}(a) - \omega_{\mathscr{P}}(a))da < 0, \quad \int_0^{a_0-\varepsilon_2} (\omega_{\mathscr{A}}(a) - \omega_{\mathscr{P}}(a))da > 0.$$

In particular,

$$\tau^{\mathscr{A}}(a_0 - \varepsilon_2) - \tau^{\mathscr{P}}(a_0 - \varepsilon_2) = \int_0^{a_0-\varepsilon_2} (\omega_{\mathscr{A}}(a) - \omega_{\mathscr{P}}(a))da + \int_{a_0-\varepsilon_2}^{a_0} \frac{a_0 - s}{\varepsilon_2}(\omega_{\mathscr{A}}(a) - \omega_{\mathscr{P}}(a))da.$$

For ε_2 sufficiently small, the first integral of the RHS, which is strictly positive, dominates. Hence, there exists a small enough ε_2 such that $\tau^{\mathscr{A}}(a_0 - \varepsilon_2) - \tau^{\mathscr{P}}(a_0 - \varepsilon_2) > 0$. Therefore it must be that $\bar{\alpha} \geq 0$.

By a similar argument, assuming that $\underline{\alpha} < 0$ would lead us to a contradiction, and hence the conclusion that $\underline{\alpha} \geq 0$. So it must be that $\bar{\alpha} \geq 0$ and $\underline{\alpha} \geq 0$: hence, $\bar{\alpha} = \underline{\alpha} = 0$. □

The conclusion of claim 4 contradicts assumption 4. Hence, there exists some $a \neq a_0$ for which $\tau^{\mathscr{A}}(a) \geq \tau^{\mathscr{P}}(a)$. The agent strictly prefers to discover this attribute compared to no attributes. □

[32] Suppose, toward a contradiction, that $\exists x_1 > 0$ such that $\lambda(x_1) \leq x_1$. But, $\frac{\partial \lambda(x)}{\partial x} = \lambda(x)(\lambda(x) - x)$. Hence, $\frac{\partial \lambda(x)}{\partial x}|_{x=x_1} \leq 0$. Hence, $\lambda(x)$ is weakly decreasing at x_1. Therefore, for any $x > x_1$, $\lambda(x) - x < 0$. But this contradicts the fact that $\lim_{x\to\infty} \frac{\lambda(x)}{x} = 1$.

Proof for lemma 5.1. Using the fact that $(v_1^{\mathscr{A}}, v_1^{\mathscr{P}})$ are distributed according to a bivariate normal distribution, the objective of the agent simplifies as follows:

$$\int_0^\infty \int_0^\infty v_1^{\mathscr{A}} f(v_1^{\mathscr{A}}, v_1^{\mathscr{P}}) dv_1^{\mathscr{P}} dv_1^{\mathscr{A}} = \int_0^\infty v_1^{\mathscr{A}} \frac{1}{\sigma_{\mathscr{A}}} \phi\left(\frac{v_1^{\mathscr{A}} - v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}}\right) \Phi\left(\frac{\rho}{\sqrt{1-\rho^2}} \frac{v_1^{\mathscr{A}} - v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}} + \frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}\sqrt{1-\rho^2}}\right) dv_1^{\mathscr{A}}$$

$$= \int_{-v_0^{\mathscr{A}}/\sigma_{\mathscr{A}}}^{\infty} (x\sigma_{\mathscr{A}} + v_0^{\mathscr{A}})\phi(x)\Phi\left(\frac{\rho}{\sqrt{1-\rho^2}} x + \frac{v_0^{\mathscr{P}}}{\sqrt{1-\rho^2}\sigma_1^{\mathscr{P}}}\right) dx$$

The following two identities (10,010.3) and (10,011.1) from **?** are useful to evaluate this integral:

$$\int \phi(x)\Phi(a+bx)dx = T(x, \frac{a}{x\sqrt{1+b^2}}) + T(\frac{a}{\sqrt{1+b^2}}, \frac{x\sqrt{1+b^2}}{a}) - T(x, \frac{a+bx}{x})$$
$$- T(\frac{a}{\sqrt{1+b^2}}, \frac{ab + x(1+b^2))}{a}) + \Phi(x)\Phi\left(\frac{a}{\sqrt{1+b^2}}\right),$$

$$\int x\phi(x)\Phi(a+bx)dx = \frac{b}{\sqrt{1+b^2}}\phi\left(\frac{a}{\sqrt{1+b^2}}\right)\Phi\left(x\sqrt{1+b^2} + \frac{ab}{\sqrt{1+b^2}}\right) - \Phi(a+bx)\phi(x).$$

where

$$T(h,a) = \int_0^a \frac{\phi(h)\phi(hx)}{1+x^2} dx.$$

Using the first identity, and also the fact that $T(\infty,0) = 0$, $T(\infty,1) = 0$:[33]

$$\int_{-c_{\mathscr{A}}}^{\infty} \phi(x)\Phi\left(\frac{\rho}{\sqrt{1-\rho^2}}x + \frac{c_{\mathscr{P}}}{\sqrt{1-\rho^2}}\right) dx = \Phi(c_{\mathscr{P}}) - \Phi(-c_{\mathscr{A}})\Phi(c_{\mathscr{P}}) - \Big(T(-c_{\mathscr{A}}, -\frac{c_{\mathscr{P}}}{c_{\mathscr{A}}}) +$$
$$T(c_{\mathscr{P}}, -\frac{c_{\mathscr{A}}}{c_{\mathscr{P}}}) - T(-c_{\mathscr{A}}, \frac{\rho c_{\mathscr{A}} - c_{\mathscr{P}}}{c_{\mathscr{A}}\sqrt{1-\rho^2}}) - T(c_{\mathscr{P}}, \frac{c_{\mathscr{A}}\rho - c_{\mathscr{P}}}{c_{\mathscr{A}}\sqrt{1-\rho^2}})\Big)$$
$$= \Phi(c_{\mathscr{P}})\Phi(c_{\mathscr{A}}) + T(c_{\mathscr{A}}, \frac{c_{\mathscr{P}}}{c_{\mathscr{A}}}) + T(c_{\mathscr{P}}, \frac{c_{\mathscr{A}}}{c_{\mathscr{P}}})$$
$$- T(c_{\mathscr{A}}, \frac{c_{\mathscr{P}} - \rho c_{\mathscr{A}}}{c_{\mathscr{A}}\sqrt{1-\rho^2}}) - T(c_{\mathscr{P}}, \frac{c_{\mathscr{A}} - \rho c_{\mathscr{P}}}{c_{\mathscr{P}}\sqrt{1-\rho^2}})$$
$$= BvN\left(c_{\mathscr{A}}, c_{\mathscr{P}}; \rho\right),$$

where $c_i = v_0^i/\sigma_i$. The function $T(h,a)$ is even in h and odd in a. The last step follows from identity 3.2 in **?**, where BvN denotes the cdf of the standard bivariate normal. Using the second identity, we obtain:

$$\int_{-v_0^{\mathscr{A}}/\sigma_{\mathscr{A}}}^{\infty} x\phi(x)\Phi\left(\frac{\rho}{\sqrt{1-\rho^2}}x + \frac{v_0^{\mathscr{P}}}{\sqrt{1-\rho^2}\sigma_1^{\mathscr{P}}}\right) dx = \rho\phi(c_{\mathscr{P}}) - \rho\phi(c_{\mathscr{P}})\Phi\left(\frac{c_{\mathscr{P}}\rho - c_{\mathscr{A}}}{\sqrt{1-\rho^2}}\right)$$
$$+ \phi(-c_{\mathscr{A}})\Phi\left(\frac{c_{\mathscr{P}} - \rho c_{\mathscr{A}}}{\sqrt{1-\rho^2}}\right)$$
$$= \rho\phi(c_{\mathscr{P}})\Phi\left(\frac{c_{\mathscr{A}} - c_{\mathscr{P}}\rho}{\sqrt{1-\rho^2}}\right) + \phi(c_{\mathscr{A}})\Phi\left(\frac{c_{\mathscr{P}} - \rho c_{\mathscr{A}}}{\sqrt{1-\rho^2}}\right)$$

[33] These follow from the fact that for any h, $T(h,0) = 0$ and $T(h,1) = \Phi(h)(1-\Phi(h))/2$.

Putting together all the expressions we derived, the objective of the agent simplifies to

$$v_0^{\mathscr{A}} BvN(c_{\mathscr{A}}, c_{\mathscr{P}}; \rho) + \sigma_{\mathscr{A}} \left(\rho\phi(c_{\mathscr{P}})\Phi\left(\frac{c_{\mathscr{A}} - c_{\mathscr{P}}\rho}{\sqrt{1-\rho^2}}\right) + \phi(c_{\mathscr{A}})\Phi\left(\frac{c_{\mathscr{P}} - \rho c_{\mathscr{A}}}{\sqrt{1-\rho^2}}\right)\right).$$

□

Proof for lemma 5.2. For $q = 1$, $\rho = 1$. But $BvN(h, k, 1) = \Phi(\min\{h, k\})$ from identity 3.4 in **?**. Hence the payoff of the agent simplifies to

$$\begin{cases} v_0^{\mathscr{A}} \Phi(c_{\mathscr{A}}) + \sigma_{\mathscr{A}} \phi(c_{\mathscr{A}}) & \text{if } c_{\mathscr{P}} > c_{\mathscr{A}}, \\ v_0^{\mathscr{A}} \Phi(c_{\mathscr{P}}) + \sigma_{\mathscr{A}} \phi(c_{\mathscr{P}}) & \text{otherwise.} \end{cases}$$

□

Proof for proposition 5.1. If

$$\frac{v_0^{\mathscr{A}}}{\sigma_{\mathscr{A}}(a_{\mathscr{A}})} < \frac{v_0^{\mathscr{P}}}{\sigma_{\mathscr{P}}(a_{\mathscr{A}})}$$

the agent's payoff from sampling $a_{\mathscr{A}}$ in the principal-agent problem is the same as the maximum payoff attained in her single-player problem. The agent prefers this payoff to the payoff attained from sampling any other attribute. □

Proof for proposition 5.2. We need to distinguish two cases, based on which player requests higher B_0 in order to be ex-ante indifferent between adoption and rejection.

(i) $-\mu K_{\mathscr{A}} > -\mu K_{\mathscr{P}}$: For any $B_0 \in (-\mu K_{\mathscr{P}}, -\mu K_{\mathscr{A}})$, the optimal attribute is $a_{\mathscr{A}}$. For any such B_0, $v_0^{\mathscr{A}}/\sigma_{\mathscr{A}}$ is negative for all a, and $v_0^{\mathscr{P}}/\sigma_{\mathscr{P}}$ is positive for all a. By proposition 5.1, the optimal attribute is $a_{\mathscr{A}}$.

(ii) $-\mu K_{\mathscr{A}} \le -\mu K_{\mathscr{P}}$: Using proposition 5.1, if $\tau_{\mathscr{P}}(a_{\mathscr{A}})/\tau_{\mathscr{A}}(a_{\mathscr{A}}) < 1$, $a_{\mathscr{A}}$ is optimal for any $B_0 \in (\bar{B}, \infty)$, where $\bar{B} > -\mu K_{\mathscr{P}}$. If $\tau_{\mathscr{P}}(a_{\mathscr{A}})/\tau_{\mathscr{A}}(a_{\mathscr{A}}) \ge 1$, $a_{\mathscr{A}}$ is optimal for any $B_0 \in (-\infty, \underline{B})$, where $\underline{B} < -\mu K_{\mathscr{A}}$.

□

C Calculation and proofs for in-text examples

C.1 Calculations for example 1

Consider a subsample of size k on $[a_0, 1]$. Let $\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_k$ denote the distance of a_1 from a_0, a_2 from a_1, and so on, until a_k from a_{k-1}. The variance of this sample for $\omega(a) = 1$ is:

$$\varepsilon_1\left(1 - a_0 - \frac{\varepsilon_1}{2}\right)^2 + \varepsilon_2\left(1 - a_0 - \varepsilon_1 - \frac{\varepsilon_2}{2}\right)^2 + \ldots + \varepsilon_k\left(1 - a_0 - \varepsilon_1 - \varepsilon_2 - \ldots - \varepsilon_{k-1} - \frac{\varepsilon_k}{2}\right)^2$$

Let $T_i = 1 - a_0 - \varepsilon_1 - \ldots - \varepsilon_i$ denote the remaining distance from a_i to 1. Taking the first-order conditions with respect to ε_i for $i = 1, \ldots, k$, and then solving backwards starting from ε_k^*, we obtain that $\varepsilon_{k-i}^* = \frac{2}{2i+3} T_{k-i}$. Therefore, starting from ε_1^* such that $\varepsilon_1^* = \frac{2}{2(k-1)+3}(1 - a_0)$, we can iteratively pin down all $(\varepsilon_i^*)_{i=1}^k$. The general term is not dependent on i, i.e., all distances are equal: $\varepsilon_i^* = \frac{2}{2k+1}(1 - a_0)$. Then,

$$a_i = a_0 + \frac{2i}{2k+1}(1 - a_0).$$

Plugging ε_i^* into the initial expression for the variance of the sample, we obtain $\frac{2}{(2k+1)^3}(1-a_0)^3\sum_{i=1}^k(2(k-i)+2)^2$. Observing that $\sum_{i=1}^k(2(k-i)+2)^2 = 4\frac{k(k+1)(2k+1)}{6}$ gives us the desired expression for the variance of k attributes on $[a_0,1]$:

$$\frac{4}{3}(1-a_0)^3\frac{k(k+1)}{(2k+1)^2}.$$

The optimal sample $\mathbf{s}^*$ under capacity q is the one among all constrained-optimal samples (that allocate k attributes in $[0,a_0]$) that maximizes the variance $\sigma^2(\mathbf{s})$. That is, the agent chooses $k \le q$ that solves the following problem:

$$\max_k \quad \frac{4}{3}\left((1-a_0)^3\frac{(q-k)(q+1-k)}{(2(q-k)+1)^2} + a_0^3\frac{k(k+1)}{(2k+1)^2}\right).$$

C.2 Proofs for example 2 in section 4

This subsection formally states and proves the claims made in example 2, in which both players have linear weight functions.

Claim 5 (Ordering of single-agent constrained optima)**.** *$a_r^{\mathscr{A}} < a_r^{\mathscr{P}}$ and $a_\ell^{\mathscr{A}} < a_\ell^{\mathscr{P}}$.*

Proof of claim 5. Consider first the claim $a_0 < a_r^{\mathscr{A}} < a_r^p < 1$. The proof proceeds through the following steps: (i) as $a \to a_0$, $f^{\mathscr{P}}(a) > f^{\mathscr{A}}(a)$; (ii) for $a = 1$, $f^{\mathscr{P}}(1) < f^{\mathscr{A}}(1)$; (iii) $f^{\mathscr{P}}$ and $f^{\mathscr{A}}$ cross only once on $[a_0,1]$, (iv) the unique crossing attribute $\tilde{a} \in (a_0,1)$ is such that $f^{\mathscr{A}}(\tilde{a}) = f^{\mathscr{P}}(\tilde{a}) < 0$. First, note that $f^{\mathscr{A}}(a) < f^{\mathscr{P}}(a)$ given that $a \in [a_0,1]$ and $k_{\mathscr{A}} b_{\mathscr{P}} < b_{\mathscr{A}} k_{\mathscr{P}}$ if and only if the following inequality holds:

$$10a^2 + (3-2a_0)a_0 < a(9+2a_0).$$

(i) At $a = a_0$, the inequality is true. Hence $f^{\mathscr{A}}(a_0) < f^{\mathscr{P}}(a_0)$ as we approach a_0 from the right.

(ii) At $a = 1$, the inequality is false. Hence $f^{\mathscr{A}}(1) > f^{\mathscr{P}}(1)$.

(iii) The two curves $f^{\mathscr{P}}$ and $f^{\mathscr{A}}$ intersect only once at

$$\tilde{a} = \frac{1}{20}\left(9 + 2a_0 + \sqrt{81 - 84(1-a_0)a_0}\right).$$

(iv) It is straightfoward to show that $f^{\mathscr{A}}(\tilde{a}) = f^{\mathscr{P}}(\tilde{a}) < 0$.

Therefore, $f^{\mathscr{A}}(a) < f^{\mathscr{P}}(a)$ for any a for which any of the two functions is positive. Hence, $f^{\mathscr{A}}$ hits zero at a smaller attribute, i.e. $a_r^{\mathscr{A}} < a_r^{\mathscr{P}}$. The proof for the second part $a_\ell^{\mathscr{A}} < a_\ell^{\mathscr{P}}$ is analogous. □

Claim 6 (Monotonicity in compromise region)**.** *a_ℓ^* is monotonically increasing in $|v_0|$. It approaches $a_\ell^{\mathscr{P}}$ as $|v_0| \to \infty$.*

Proof for claim 6. From claim 5, $a_\ell^{\mathscr{P}} > a_\ell^{\mathscr{A}}$. For $v_0 = 0$, $a_\ell^*(0) = a_\ell^{\mathscr{A}}$. Also, from the proof of claim 5, $f^{\mathscr{P}}(a) > f^{\mathscr{A}}(a)$ for any $a \in [a_\ell^{\mathscr{A}}, a_\ell^{\mathscr{P}}]$. Suppose for some v_0, $a^*(v_0)$ is an optimal attribute, and consider a small increase of $|v_0|$ by ε. Consider the FOC condition:

$$\left(\frac{v_0}{(\sigma^{\mathscr{P}}(a;a_0))}\right)^2\left(\frac{\tau^{\mathscr{A}}(a;a_0)}{\tau^{\mathscr{P}}(a;a_0)} - 1\right)f^{\mathscr{P}}(a) + f^{\mathscr{A}}(a) = 0. \tag{FOC}$$

Because $a^*(v_0)$ is interior, it satisfies FOC. As $|v_0| \to |v_0| + \varepsilon$, LHS of FOC evaluated at $a^*(v_0)$ becomes strictly positive, because $f^{\mathscr{P}}(a^*(v_0)) > 0$. We want to argue that $a^*(|v_0|+\varepsilon) > a^*(|v_0|)$.

Claim 7. *For any $a \in (a_0,1]$, the term $\frac{\tau^{\mathscr{A}}(a;a_0)}{\tau^{\mathscr{P}}(a;a_0)} - 1$ is strictly decreasing in a.*

Proof. Observe that:

$$\frac{\tau^{\mathscr{A}}(a;a_0)}{\tau^{\mathscr{P}}(a;a_0)}-1=\frac{(2a^2+(a+a_0)(-3+2a_0))(-b_{\mathscr{A}}k_{\mathscr{P}}+b_{\mathscr{P}}k_{\mathscr{A}})}{(3(-2+a+a_0)b_{\mathscr{P}}+(-3+a^2+aa_0+a_0^2)k_{\mathscr{P}})(2b_{\mathscr{A}}+k_{\mathscr{A}})}.$$

The first derivative of this with respect to a is:

$$\frac{\partial\left(\frac{\tau^{\mathscr{A}}(a;a_0)}{\tau^{\mathscr{P}}(a;a_0)}-1\right)}{\partial a}=\frac{3(-1+a)(-3+a+2a_0)(2b_{\mathscr{P}}+k_{\mathscr{P}})(-b_{\mathscr{A}}k_{\mathscr{P}}+b_{\mathscr{P}}k_{\mathscr{A}})}{(3(-2+a+a_0)b_{\mathscr{A}}+(-3+a^2+aa_0+a_0^2)k_{\mathscr{P}})^2(2b_{\mathscr{A}}+k_{\mathscr{A}})}.$$

By assumption in this example, $b_{\mathscr{P}}k_{\mathscr{A}}<b_{\mathscr{A}}k_{\mathscr{P}}$. Moreover, $a_0<a<1$. Hence, this derivative is strictly negative. □

By increasing the attribute marginally from $a^*(v_0)$ to $a^*(v_0)+\tilde{\varepsilon}$, the first additive term in FOC goes down because both $\frac{\tau^{\mathscr{A}}}{\tau^{\mathscr{P}}}-1$ and $f^{\mathscr{P}}$ go down, while $\sigma^{\mathscr{P}}(a)$ goes up. The second additive term $f^{\mathscr{A}}$ goes down as well (becomes more negative). Hence, the LHS of FOC goes down. □

Claim 8 (Undershooting)**.** *Consider the constrained problem on* $[a_0,1]$*. Then,*

(i) $a_r^*(v_0)\in[a_0,a_r^{\mathscr{A}}]$ *for sufficiently inconclusive* v_0*;*

(ii) There exists a sufficiently large c *such that* $a_r^*(v_0)=a_0$ *for* $|v_0|>c$ *and* $a_r^*(v_0)>a_0$ *otherwise;*

Proof. (i) For $v_0=0$, $a_r^*(0)=a_r^{\mathscr{A}}$. By continuity, for a sufficiently small $|v_0|$, $a_0<a_r^*(v_0)<a_r^{\mathscr{A}}$.

(ii) From lemma 4.5, we know that there exists a sufficiently large $|v_0|=c$ such that for all $|v_0|>c$, the optimal attribute $a_r^*=a_0$. □

D Model variations

D.1 No initial attributes are perfectly known

This section briefly remarks on the single-agent problem in which: 1) no prior attributes are known, 3) the player has a normal prior on $B(0)\sim\mathcal{N}(0,1)$. The parameters of the underlying Brownian process are $(\mu,\sigma)=(0,1)$. The realization of each attribute $a>0$ is distributed normally according to $B(a)\sim\mathcal{N}(m,1+a)$. Given a sampled realization of attribute $a>0$, the posterior distribution of $B(0)$ has density:

$$\Pr(B(0)\mid B(a))=\frac{\Pr(B(a)\mid B(0))\Pr(B(0))}{\Pr(B(a))}=\frac{\phi\left(\frac{B(a)-B(0)}{a}\right)\phi(B(0))}{\phi\left(\frac{B(a)}{\sqrt{1+a}}\right)}.$$

The posterior expected value of the project when $B(a)$ is observed is:

$$\begin{aligned}v(a,B(a)) &:= \int_{-\infty}^{\infty}(B(0)\tau(0;(0,a))+B(a)\tau(a;(0,a)))\Pr(B(0)\mid B(a))dB(0)\\ &= B(a)\tau(a;(0,a))+\tau(0;(0,a))\bar{B},\end{aligned}$$

where $\bar{B}$ is the conditional expectation of $B(0)$ given that $B(a)$ is observed. The project is adopted only if the observed $B(a)$ is sufficiently high so that the posterior expected value is positive.

$$\max_a \quad \Big(1-\Pr(v(a,B(a))\geq 0)\Big)\mathbb{E}\Big(v(a,B(a))\;\Big|\;v(a,B(a))\geq 0\Big).$$

Naturally, for a positive drift, attributes greater than zero have higher mean and higher variance. Yet, the distributions of the corresponding posterior expected value have the same mean due to the law of iterated expectations, that is $\mathbb{E}_{B(a)}\mathbb{E}_{B(0)}[B(0) \mid B(a)] = \mathbb{E}[B(0)] = 0$. Hence,

$$v(a, B(a)) \sim \mathcal{N}\left(0, \quad a\tau^2(a) + \tau^2(0)\,var(\bar{B}) + 2\tau^2(a)\tau^2(0)\,\text{cov}(B(a), \bar{B})\right).$$

The covariance term $\text{cov}(B(a), \bar{B})$ is positive because a higher $B(a)$ implies higher $\mathbb{E}[B(0) \mid B(a)]$. It can be shown easily that $var(\bar{B})$ is increasing in a as well. But as a increases, $\tau(0)$ increases and $\tau(a)$ decreases. The agent explores the attribute that induces the highest variance on v.

D.2 Agent cares only about adoption

Suppose that the agent cares only about adoption. Given a realized path B on $[0,1]$, her payoff function is

$$u(d, B) = \begin{cases} 1 & \text{if } d = 1 \\ 0 & \text{if } d = 0. \end{cases}$$

Given sample a, the probability of adoption is:

$$\Pr(v_0^{\mathcal{P}} \geq 0) = \Phi\left(\frac{v_0^{\mathcal{P}}}{\sigma_{\mathcal{P}}(a)}\right).$$

This is increasing in $\sigma_{\mathcal{P}}$ if $v_0^{\mathcal{P}} < 0$, and decreasing otherwise. If $v_0^{\mathcal{P}} > 0$, the principal is willing to adopt based on what she knows about a_0, hence the agent does not sample further. If $v_0^{\mathcal{P}} < 0$, the principal needs to be convinced to approve. Approval is most likely if the discovered sample is the one that is most informative to the principal. This observation leads to the following proposition.

Proposition D.1. *The agent samples the principal's ideal sample if the project is initially unpromising to the principal* ($v_0^{\mathcal{P}} \leq 0$)*, and samples no additional attributes otherwise (i.e. if* $v_0^{\mathcal{P}} > 0$*).*

Proof. See preceding paragraph. □

D.3 Simultaneous sampling at a fixed per-attribute cost

This section considers sampling at a fixed cost $c > 0$ per attribute for the driftless case ($\mu = 0$). The agent is no longer constrained in the number of attributes she can sample, as in the case of quota-based sampling: she can take as many attributes as she sees fit as long as she affords the cost c. In the case of simultanous costly sampling, the agent decides on a bundle of attributes and pays the cost ($c \times$ size of bundle) upfront, without the possibility of purchasing additional attributes after observing the realizations of this bundle. In the case of sequential costly sampling, the agent discovers attributes one at a time: the decision whether to stop or acquire an additional costly attribute (and if so, which attribute) is potentially informed by the realizations of attributes observed thus far.

Let us first analyze the case of simultaneous sampling. The agent observes (a_0, v_0) and has to decide on 1) the number n of attributes to be discovered, 2) conditional on taking n attributes, which ones to take. But, the analysis of quota-based sampling assures us that for a fixed n, the optimal bundle of n attributes is independent of v_0. The initial attribute a_0 is sufficient to identify the optimal bundle of size n. Hence, the problem of the agent is reduced only to solving for the optimal bundle size n^* for a given initial promise v_0. Hereafter, the optimal bundle among all those of size n is referred to simply as *bundle of size n*.

Let $V(v_0, n)$ denote the benefit attained by the bundle of size n when v_0 is the initial expected value of the project. Also, let σ_n^* denote the standard deviation of the random variable $v_n = \mathbb{E}[v|\mathscr{A}_n]$, where $\mathscr{A}_n$ is the partition created from the bundle of size n and the initial attribute a_0. Then, the payoff from acquiring this bundle is

$$V(v_0, n) - cn = v_0 \Phi\left(\frac{v_0}{\sigma_n^*}\right) + \sigma_n^* \phi\left(\frac{v_0}{\sigma_n^*}\right) - cn$$

In the absence of any further information beyond (a_0, v_0) (i.e. $n = 0$), the payoff is

$$V(v_0, 0) = \begin{cases} v_0 & \text{if } v_0 \geq 0 \\ 0 & \text{if } v_0 < 0. \end{cases}$$

For bundle n to be preferred to no information, the standard deviation of v_n should be sufficiently high compared to the cost, i.e. $\sigma_n^* \geq cn\sqrt{2\pi}$.

Lemma D.1.

(i) Suppose that $V(0,n) - cn > V(0,0) = 0$. Then, there exists $\bar{v}_n > 0$ such that the bundle of size n is preferred to no information for $v_0 \in [-\bar{v}_n, \bar{v}_n]$, and no information is preferred to the bundle of size n otherwise.

(ii) Suppose $v_0 = 0$. If a bundle of size $n > 0$ is preferred to no further information, then any smaller bundle of size $n' < n$ is so as well.

(iii) Suppose n is preferred to n' at $v_0 = 0$. If $n > n'$, there exists a unique $\bar{v}_{n,n'}$ such that n is preferred to n' for $v_0 \in [-\bar{v}_{n,n'}, \bar{v}_{n,n'}]$ and n' is preferred otherwise. If $n < n'$, then n is preferred to n' for all v_0.

Proof. (i) Consider the difference $f(v_0) := V(v_0, n) - V(v_0, 0)$ for $v_0 \in [0, \infty)$. From the premise, $f(0) > cn$. Moreover, it is strictly decreasing in v_0:

$$\frac{\partial f(v_0)}{\partial v_0} = \Phi\left(\frac{v_0}{\sigma_n^*}\right) - 1 < 0.$$

As $v_0 \to \infty$, $f(v_0) \to 0$. Hence, there is a unique $\bar{v}_n > 0$ such that $f(\bar{v}_n) = cn$, i.e.

$$\bar{v}_n \Phi\left(\frac{\bar{v}_n}{\sigma_n^*}\right) + \sigma_n^* \phi\left(\frac{\bar{v}_n}{\sigma_n^*}\right) - cn = \bar{v}_n.$$

Similarly, consider $f(v_0)$ on $(-\infty, 0]$. By a similar argument, $f(v_0)$ is strictly increasing on this domain:

$$\frac{\partial f(v_0)}{\partial v_0} = \Phi\left(\frac{v_0}{\sigma_n^*}\right) > 0.$$

As $v_0 \to -\infty$, $f(v_0) \to 0$. Hence, there exists a unique $\underline{v}_n < 0$ such that $f(\underline{v}_n) = cn$, i.e.

$$\underline{v}_n \Phi\left(\frac{\underline{v}_n}{\sigma_n^*}\right) + \sigma_n^* \phi\left(\frac{\underline{v}_n}{\sigma_n^*}\right) - cn = 0.$$

Finally, observe that applying the facts that $\phi(x) = \phi(-x)$ and $\Phi(-x) = 1 - \Phi(x)$, we obtain that f is symmetric around $v_0 = 0$, i.e. for any $v_0 > 0$, $f(v_0) = f(-v_0)$. Therefore, $\bar{v}_n = -\underline{v}_n$.

(ii) If $V(0,n) - cn > V(0,0)$, then $\frac{\sigma_n^*}{n} \geq c\sqrt{2\pi}$. By the claim in lemma **??**, for any n and $n' < n$,

$$\frac{\sigma_{n'}^*}{n'} \geq \frac{\sigma_n^*}{n}.$$

Therefore, for any $n' < n$, $\frac{\sigma^*_{n'}}{n'} \geq c\sqrt{2\pi}$ as well.

(iii) Suppose $n > n'$. Then $\sigma^*_n > \sigma^*_{n'}$. Define $g(v_0) := V(v_0, n) - V(v_0, n') - c(n - n')$. By the premise, $g(0) > 0$.

$$\frac{\partial g(v_0)}{\partial v_0} = \Phi\left(\frac{v_0}{\sigma^*_n}\right) - \Phi\left(\frac{v_0}{\sigma^*_{n'}}\right).$$

For $v_0 > 0$, g is strictly decreasing. Moreover, as $v_0 \to \infty$, g tends to $-c(n - n') < 0$. Therefore, there exists a unique $\bar{v}_{n,n'}$ such that $g(\bar{v}_{n,n'}) = 0$. Bundle n is preferred to n' for $v_0 \in [0, \bar{v}_{n,n'}]$ and bundle n' is preferred for $v_0 \in (\bar{v}_{n,n'}, \infty)$. Consider now the negative half line, $v_0 < 0$. By a similar argument, g is strictly increasing in v_0 on $(-\infty, 0)$ and $g(v_0) \to -c(n - n') < 0$ as $v_0 \to -\infty$. Therefore, there exists $\underline{v}_{n,n'} < 0$ for which $g(\underline{v}_{n,n'}) = 0$. By the fact that g is an even function, $\bar{v}_{n,n'} = -\underline{v}_{n,n'}$.

Suppose now that $n < n'$, so $\sigma^*_n < \sigma^*_{n'}$. It is immediate to observe that the difference function g is strictly increasing in v_0 for $v_0 > 0$ and strictly decreasing in v_0 for $v_0 < 0$. But $g(0) > 0$, hence, $g(v_0) > g(0) > 0$ for all v_0. □

The previous lemma provides some crucial features of the optimal bundle size for a given v_0. For the most neutral initial promise possible $v_0 = 0$, the agent has a set of costly bundle sizes of the form $\{1, 2, \ldots, k\}$ that she prefers to the option of remaining uninformed. From this set, she chooses the size that solves the following problem:

$$\max_{n \in \{1,\ldots,k\}} V(0, n) - cn.$$

Let n^* be the solution to this problem. First, any bundle of greater size n such that $k \geq n > n^*$ is not taken for any v_0. The greatest bundle is taken precisely for the most neutral initial promise, that is, when the agent is most uncertain about the value of the project. Secondly, the lemma highlights the symmetry of the optimal rule: if n is optimal for a given v_0, then n is optimal for $-v_0$ as well. To see this how this observation follows from the lemma, it might be helpful to make two remarks. First, the bundle of size n is preferred to the absence of any sampling for v_0 within a given distance from $v_0 = 0$, despite the sign of v_0, if preferred at $v_0 = 0$. Secondly, the ranking of any two bundles n and n' is symmetric around $v_0 = 0$: if n is preferred to n' for $v_0 > 0$, then it is so for $-v_0 < 0$ as well.

On the other hand, lemma D.1 does not rule out the possibility that some bundle size $n < n^*$ might be suboptimal for all v_0. The following example illustrates such an instance.

Example 7. *Let $\omega(a) = 1$ for all $a \in [0, 1]$ and $a_0 = 2/5$. Suppose $c = 1/400$. For this cost, the optimal pattern of sampling is:*[34]

$$n^* = \begin{cases} 3 & \textit{for } |v_0| \in [0, 0.239104] \\ 2 & \textit{for } |v_0| \in (0.239104, 0.495794] \\ 0 & \textit{for } |v_0| \in (0.495794, \infty) \end{cases}$$

At $v_0 = 0$, the optimal bundle size is $n^* = 3$. As the initial promise v_0 becomes more conclusive (i.e. its absolute value increases), the optimal bundle size decreases to two and then to no information at all. A bundle of size $n = 1$ is never taken: $n = 1$ is preferred to $n = 2$ for $|v_0| \geq 0.502103$. Yet, such values of v_0 are sufficiently conclusive that the agent prefers to acquire no information at all over acquiring $n = 1$ or $n = 2$ (or any other bundle size, for that matter).

Lemma D.2. *For any v_0, v'_0 such that $|v'_0| > |v_0|$, the optimal bundle size for v_0 is weakly greater than for v'_0.*

[34] For the given value of a_0, the bundle of size $n = 2$ consists of $\{a_1, a_2\} = \{16/25, 22/25\}$ and the bundle of size $n = 3$ consists of $\{a_1, a_2, a_3\} = \{4/7, 26/35, 32/35\}$.

Proof. By part (iii) of lemma D.1, the optimal rule is symmetric around $v_0 = 0$, i.e. if n is optimal for v_0, then n is optimal for $-v_0$ as well. So it is without loss to suppose $v_0' > v_0 > 0$, and let n' and n be the respective optimal bundle sizes for v_0' and v_0 respectively. By way of contradiction, suppose that $n' > n$. Optimality of n' at v_0' implies that

$$V(v_0', n') - V(v_0', n) > c(n' - n).$$

As in the proof of lemma D.1, $V(v_0', n') - V(v_0', n)$ is strictly decreasing in v_0' for $v_0' \in [0, \infty)$. Hence, for $v_0 < v_0'$,

$$V(v_0, n') - V(v_0, n) > c(n' - n).$$

By the optimality of n at v_0, it must be that n' is not affordable at v_0, i.e. $V(v_0, n') < cn'$. But for the inequality to hold, it must be that $V(v_0, n) < cn$ as well. This contradicts the premise that n is affordable and optimal at v_0. □

Lemma D.2 establishes that the agent acquires smaller bundles of attributes for more conclusive v_0. This is an intuitive observation, as a more conclusive initial promise implies that the agent is closer to an approval/rejection decision before any discoveries, so less additional information will be needed to reach a final decision. To put this differently, the added benefit of a greater bundle n' compared to n is decreasing in v_0 for initial promises $v_0 \in [0, \infty)$. Hence, if the agent is willing to pay the cost difference $c(n' - n)$, it must be for higher initial promises.

Lemma D.3 (Ordering of thresholds $\bar{v}_k$)**.**

(i) If the indifference point between n and n′ is such that $\bar{v}_{n,n'} < \min\{\bar{v}_n, \bar{v}_{n'}\}$*, then* $\bar{v}_{n'} < \bar{v}_n$.

(iii) If the indifference point $\bar{v}_{n,n'} > \max\{\bar{v}_n, \bar{v}_{n'}\}$*, then* $\bar{v}_n < \bar{v}_{n'}$.

(ii) If n is preferred to n′ for all $v_0 \in (-\infty, \infty)$*, then* $\bar{v}_n > \bar{v}_{n'}$.

Proof. (i) Suppose by way of contradiction that $\bar{v}_{n,n'} < \min\{\bar{v}_n, \bar{v}_{n'}\}$ and $\bar{v}_{n'} \geq \bar{v}_n$. In other words, suppose

$$-\bar{v}_{n'} < -\bar{v}_n < -\bar{v}_{n,n'}.$$

By the indifference condition,

$$V(\bar{v}_{n,n'}, n') - V(\bar{v}_{n,n'}, n) = c(n' - n).$$

But the function $f(v_0) := V(v_0, n') - V(v_0, n)$ is strictly increasing for $v_0 < 0$ and strictly decreasing for $v_0 > 0$ because due to $\sigma^*_{n'} > \sigma^*_n$,

$$f'(v_0) = \frac{1}{2}\left(\Phi\left(\frac{v_0}{\sigma^*_{n'}}\right) - \Phi\left(\frac{v_0}{\sigma^*_n}\right)\right) < 0.$$

Therefore, $f(v_0) < c(n' - n)$ for all $v_0 < -\bar{v}_{n,n'}$. The assumption that $-\bar{v}_n > -\bar{v}_{n'}$ implies that at $v_0 = \bar{v}_{n'}$,

$$V(v_{n'}, n) - cn < 0.$$

Hence,

$$V(v_{n'}, n') - cn' - (V(v_{n'}, n) - cn) > 0 \Rightarrow f(v_{n'}) > c(n' - n).$$

This contradicts our earlier observation that $f(v_0) < c(n' - n)$ for all $v_0 < -\bar{v}_{n,n'}$.

(ii) It follows from reasoning almost identical to that in the proof of part (ii), assuming initially (towards a contradiction) that $-\bar{v}_{n,n'} < -\bar{v}_n < -\bar{v}_{n'}$.

(iii) By the premise, $V(v_0,n) - V(v_0,n') > 0$ for all v_0. Hence, in particular at $v_0 = -\bar{v}_{n'}$, it must be that $V(-\bar{v}_{n'},n) > 0$. By the fact that $V(v_0,n)$ is strictly increasing in v_0, it follows that $-\bar{v}_n < -\bar{v}_{n'} \Leftrightarrow \bar{v}_n > \bar{v}_{n'}$. □

D.4 Player-specific outside option from rejection

The main analysis assumes that the players have different weight functions, but they both enjoy the same outside value from rejection, normalized to zero. In the single-player problem, the normalization of the outside option to zero is without loss: optimal sampling is the same despite the outside option. Here, we briefly explain the consequences of another form of conflict: the players share the same weight function, but they disagree on the value of rejection. Let (α,π) be the outside option of the agent and the principal, respectively. Suppose that $\pi < v_0$, so in the absence of further information, the principal adopts the project. The objective of the agent is to choose a feasible sample that maximizes the following:

$$\max_{s\in\mathcal{S}_{q+1}(a_0)} \quad v_0 + (\alpha - v_0)\Phi\left(\frac{\pi - v_0}{\sigma(s)}\right) + \sigma(s)\phi\left(\frac{\pi - v_0}{\sigma(s)}\right).$$

The following lemma establishes that conditional on the agent sampling, attribute choice is not distorted by the different outside options. Yet, unlike in the single-player problem, the agent might decide not to sample at all if the principal is too demanding towards the project, i.e. if she has too high of an outside option compared to the agent.

Lemma D.4. *Suppose $v_0 > \pi$. There exists a threshold $c(\omega,a_0) < 0$ such that:*

(i) if $\alpha - \pi \geq c(\omega,a_0)$, the optimal sample coincides with the optimal sample in the single-player problem;

(ii) if $\alpha - \pi < c(\omega,a_0)$, the agent does not sample at all.

Proof. The objective is increasing in σ if $(\pi - v_0)(\pi - \alpha) + \sigma^2 > 0$ and decreasing otherwise. If $(\pi - v_0)(\pi - \alpha) > 0$, i.e. if $\alpha > \pi$, the objective is always increasing in σ, hence the agent picks the sample with the highest σ.

Let $\bar{\sigma} := \arg\max_s \sigma(s)$. Suppose $\alpha < \pi$. If $-\bar{\sigma}^2 < (\pi - v_0)(\pi - \alpha) < 0$, the objective is decreasing in $\sigma(s)$ for $\sigma(s)$ below a certain threshold and increasing otherwise. Therefore, there can only be two local maxima, at $\sigma = 0$ and at $\sigma = \bar{\sigma}$. Because $\alpha < \pi < v_0$, $\alpha - v_0 < 0$. If

$$v_0 \leq v_0 + (\alpha - v_0)\Phi\left(\frac{\pi - v_0}{\bar{\sigma}}\right) + \bar{\sigma}\phi\left(\frac{\pi - v_0}{\bar{\sigma}}\right),$$

the agent prefers to disclose the highest-variance sample; otherwise she prefers to disclose no additional attributes beyond a_0. For $\alpha = \pi$, the strict inequality holds (the problem reduces to a single-agent problem). Therefore, by continuity, it also holds for $\alpha < \pi$ sufficiently close to π.

If $\alpha - \pi < \frac{-\bar{\sigma}^2}{v_0 - \pi}$, the agent prefers not to sample at all. □

No further sampling guarantees adoption. If the agent benefits more than the principal from rejection ($\alpha > \pi$), she is always willing to sample. Even though the principal adopts more than what the agent would after sampling, the lottery on the adoption decision induced by sampling is still preferred to immediate sure adoption. Suppose, instead, that $\alpha << \pi < v_0$. Upon sampling, the principal rejects much more frequently than what the agent would. Because the agent can guarantee herself adoption by not sampling any further, she discovers no attributes and the principal adopts immediately.

Naturally, the opposite holds when the principal favors rejection given v_0, i.e. $\pi > v_0$. If $\pi > \alpha$, the principal adopts less frequently than what the agent would prefer; yet, sampling further promises adoption of a high-quality project, while discovering no additional attributes guarantees a sure rejection. On the other hand, if π is too small relative to α, the agent ex-ante prefers rejection too. Sampling risks a very probable adoption decision, so the agent forgoes sampling completely.