Organizing Data Analytics*

RICARDO ALONSO† ODILON CÂMARA‡

London School of Economics and CEPR *University of Southern California*

May 7, 2021

Abstract

We develop a theory of credible skepticism in organizations to explain the main trade-offs in organizing data generation, analysis and reporting. We study a designer-analyst-principal game where the designer selects the information privately observed by the analyst who can misreport it at a cost, while the principal can (possibly imperfectly) audit the analyst's report. We study the principal's problem of how to allocate tasks, how much to audit the report, whether to limit her own discretion, or even shape the analyst's costs of misreporting. We show that motivating informative experimentation while discouraging misreporting are often conflicting organizational goals. As a result, the principal foregoes a perfect audit and prefers to separate the tasks of experimental design and analysis to incentivize experimentation. Finally, we provide conditions under which the optimal organization involves a dual internal-external audit system.

JEL classification: D8, D83, M10.

Keywords: Strategic experimentation, Bayesian persuasion, tampering, organizational design, information technology, audit.

---

*This manuscript incorporates some results presented by the authors under the title "Tampering with Information." For their suggestions, we thank Heski Bar-Isaac, Rahul Deb, Andrea Galeotti, Jin Li, Tony Marino, Doehong Min, Luca Picariello, Marco Pagano, Eduardo Perez-Richet, João Ramos, Lucas Siga, Vasiliki Skreta, Wing Suen, Roland Strausz, Denis Shishkin and Giorgio Zanarone. We thank attendees to several seminars and participants of the 2017 LA Theory Conference, 2018 Informs, 2020 Kent-Bristol-City Workshop, and 2020 Seminar in Communication and Persuasion.

†LSE, Houghton Street, London WC2A 2AE, United Kingdom. R.Alonso@lse.ac.uk

‡USC FBE Dept, 701 Exposition Blvd, Ste. 231, HOH-231, MC-1422, Los Angeles, CA 90089-1422. ocamara@marshall.usc.edu.

1 Introduction

Employees need to recognize that not all numbers are created equal–some are more reliable than others.

Shah, Horne and Capellá, "Good Data Won't Guarantee Good Decisions", HBR (April 2012)

The Digital and ICT revolution has made organizations awash with data by drastically reducing the costs of data gathering, storage, access, and analysis. It has also changed how managers make decisions, relying less on opinions and intuition and more on insights derived from this data.[1] In spite of these improvements, the information that reaches decision makers is still hampered by incentive conflicts: conflicts of interest over decisions result in disagreement over which data to collect and how to analyze it, and creates frictions when communicating its findings to decision makers. In this paper, we develop a theory of the organization of data analytics in the face of these frictions.

The Covid-19 pandemic offers an example of how conflicting preferences in (public) organizations affect data gathering, analysis and communication. Politicians, business owners and ordinary citizens had to make decisions on whether to re-open, keep a moderate lockdown, or change to a more restrictive lockdown. A key factor was information about the spread of the virus. Nevertheless, in the United States, there were significant differences in how each state collected data and which data was made available to the public.[2] There were also differences in how countries recorded Covid-19 deaths and related statistics, even among members of the European Union (Robbins and Reuben, 2020). In addition, numerous states were accused of manipulating data to deliberately make things look better than they were

[1]Brynjolfsson, Hitt, and Kim (2011) and Brynjolfsson and McElheran (2016) report rapid and widespread adoption of Data-Driven Decision Making (DDM) practices in organizations, where the rate of adoption is heavily influenced by a series of complementary organizational practices.

[2]For example, Johns Hopkins University lists the main approaches to compute the test positivity rate– see https://coronavirus.jhu.edu/testing/differences-in-positivity-rates . The methodology can consider the number of positive tests or consider the number of people who test positive, and the methodology can consider RT-PCR tests only, or also include antigen tests. Only 10 out of the 50 states report the data needed to compute their Approach 2.

(Michelle R. Smith and Amy, 2020), while there were reports of data manipulation in other countries.[3]

Communication frictions in (private) organizations also manifest themselves in managers' mistrust of data–a recent survey by KPMG reports that only one third of respondents trust the insights generated from their business operations[4]–but also in the underutilization of some of the data stored by firms–so called "dark data" that is unstructured, improperly recorded, or simply never actively used by members of the organization.[5]

We are interested in understanding optimal organizational practices in the face of these conflicts. For instance, fostering trust in analytics calls for policies that ensure data accuracy and integrity—e.g., through regular examination of data, access management, and audit trails—or that prevent data tampering (or minimize its effect). The adoption of new technologies such as blockchain can eliminate data tampering within organizations, giving decision makers access to information that is known to be correct (Tapscott and Tapscott, 2017). Likewise, data underutilization can be remedied by incentivizing agents to provide a more insightful analysis; e.g., by rewarding them for experimenting through the promotion of the products, innovations, and business ideas that result from their analysis.[6]

We show that reducing data tampering poses conflicting organizational goals under delegated experimentation: organizational practices that ensure data integrity and truthful reporting can also reduce how much information agents extract from the data in the first place. To wit, agents may gather "just enough" evidence if decision makers find them reliable. We argue that in the face of unresolved conflict, promoting a moderate sense of mistrust can create a culture of "healthy" skepticism in the organization: managers can

[3]For example, there are claims of data manipulation in Brazil (Phillips, 2020) and Russia (Michelle R. Smith and Amy, 2020).

[4]See KPMG, 2016, 2018. There are multiple factors behind this mistrust of data: data breaches and inaccuracies–questioning the integrity of data–but also the lack of experience with certain advanced analytics that lead managers to regard them as a "black box" and doubt the value of their results (KPMG, 2018).

[5]Firms routinely generate and store data from its normal operations which is improperly recorded–e.g., machine and sensor data from devices (IoT), website access and browsing data, social media–and comprises a substantial fraction of all data stored by a firm. See Hand (2020) for an analysis of the sources and implications of the underutilization of this data in organizations.

[6]Most data analytics is used for process or product improvement or related to other types of innovation–see Wu, Hitt, and Lou (2020).

refrain from adopting agents' self-serving recommendations, for instance to launch a new product or project, forcing the latter to provide stronger evidence backing them. However, to credibly do so, firms have a preference for limiting decision makers' ability to audit or evaluate such recommendations.

We see data analytics as composed of different tasks/job roles: (i) a data architect specifying which data to collect and how to process it (*experimental design*), and (ii) a data analyst who processes data according to the specified design and reports the findings to decision makers (*analysis/reporting*). Formally, we model data analytics as a *designer-analyst-principal* game:[7] (i) a principal decides whether to deviate from a status quo decision d_S by scaling up (d_H) or scaling down (d_L) operations; (ii) a designer (data architect) and a (data) analyst have different preferences than the principal and prefer scaling-up to the status quo, with scaling-down their least preferred choice;[8] (iii) to persuade the principal, the uninformed designer strategically designs an experiment that reveals information about a payoff-relevant state—as in Kamenica and Gentzkow (2011), KG henceforth; (iv) the outcome of this experiment is privately observed by the analyst who, before submitting his report, finds how costly it would be to misrepresent (tamper) its outcome. This simple setup captures the main experimentation-tampering trade-off: the principal needs to strike a balance between inducing the designer to select a more informative experiment while at the same time restraining the analyst from tampering. To manage this trade-off, we explore several organizational levers: (i) *task allocation*—namely, whether to integrate or separate the roles of the data architect and the data analyst– (ii) *monitoring*—how much to audit the analyst's report (auditing intensity is given by the probability $\lambda \in [0, 1]$ that the principal also observes the actual experimental outcome, thus both detecting tampering and muting its effect on decision making); (iii) *discretion*–whether to limit the options available to the principal, and (iv) *a system of punitive measures* that punishes tampering.

Agents limit the strength of the evidence supporting scaling-up in order to increase its probability–see KG–but also misrepresent unfavorable evidence if auditing is imperfect. In

[7] Alternatively, borrowing from the strategic communication literature, this would be a *designer-sender-receiver* game.

[8] Agents may favor scaling-up because of a preference for empire building, or because their human capital is tied to this decision, or as a result of improved outside opportunities.

equilibrium, the designer's experiment leads to a recommendation in $\{d_L, d_H\}$ ("up-or-down" experiment) or in $\{d_S, d_H\}$ ("status-quo" experiment); the principal follows the analyst's advice when recommending d_S or d_L but may refrain from doing so when recommending d_H unless it is confirmed by a conclusive audit. Reducing auditing intensity invites misrepresentation, but also allows the principal to credibly mistrust "favorable" evidence after an inconclusive audit. We show that the spectre of a "shadow of a doubt" unambiguously pushes the designer (regardless of task allocation) to experiment more–i.e., to provide stronger evidence when recommending d_H. Moreover, as tampering costs are only incurred by the reporting agent, integrating design and analysis leads the designer to economize on tampering costs by shifting to less informative experiments.

Firms must compare the gains from improved experimentation to the losses from increased tampering when organizing data analytics. Our first result looks at a principal that "*organizes to innovate*"–i.e, a case in which the only alternative to the status-quo is to scale-up. Consistent with our theme of "credible skepticism," she commits to an imperfect audit ($\lambda^* < 1$) as long as there is a positive probability of low tampering costs, and prefers to separate the design task from the analysis/reporting task. This insight resonates with organizations that centralize design in a corporate headquarters while coming short of implementing water-tight auditing measures.[9]

Next, we look at a principal that "*organizes for scale*"–i.e., she entertains both scaling-up or down as viable alternatives to the status-quo. It is still true that separating tasks and committing to an imperfect audit leads to stronger evidence backing a d_H recommendation, but the principal must guard against "adverse switches"–i.e., a switch from an "up-or-down" experiment to a less informative "status-quo" experiment– and it may now be optimal for her to integrate tasks and to perfectly audit the analyst's report (set $\lambda^* = 1$). Adverse switches occur because agents promoting d_H are more willing to compromise on the status quo than the principal. She can avoid such compromise by ruling out the status-quo as an option; in fact, if the principal can commit to ruling-out decisions (i.e., to reduce her discretion), then she prefers an imperfect audit ($\lambda^* < 1$). Thus, in this context, discretion

[9]For instance, Mckinsey and Co. reports on several firms centralizing data analytics around a center of excellence (CoE) tasked with homogenizing data analytics and supporting the different business units. See McKinsey and Co., 2018.

and auditing-intensity act as complements.

Our theory of credible skepticism builds on the premise that allowing some tampering can prove beneficial to firms as it provides decision makers with commitment power to reject self-serving recommendations issued with weak supporting evidence. Thus, firms where experimentation is also delegated to agents would like to incentivize some tampering. We show this by solving for the optimal organization when the firm can specify the distribution of the analyst's tampering costs.[10] Optimally, the firm both makes low tampering costs sufficiently likely and limits its auditing intensity. Under this scheme, the designer always selects a fully informative experiment. That is, organizations in our model would take actions to maximize experimentation while being subject to moderate levels of data misrepresentation. We show that this optimal organization can be implemented through a decoupled internal-external audit system.

We present the model in Section 2. Section 3 characterizes the equilibrium in the design and communication subgame for a fixed organizational structure, and Section 4 analyzes the effect of task allocation and auditing on experimental design. Section 5 covers our main insights on the optimal organization of data analytics. Section 6 discusses several extensions of the model and we conclude with a discussion of the related literature in Section 7. All proofs are in the Appendices.

2 Model

A principal relies on the information gathered, analyzed and communicated by expert, albeit biased, agents to maintain the current "status quo" or to switch to one of two alternatives. To model the different tasks of data analytics, we introduce a "designer-analyst-principal" game in which the data designer (he) specifies what information the data analyst (he) will privately observe and report to the principal (she) prior to making the decision.

Preferences and Prior Beliefs: Players are expected utility maximizers. The state space is binary, with typical realization $\theta \in \Theta = \{0, 1\}$, and players hold a common prior $\mu = \Pr[\theta =$

[10]For instance, through a by-law that defines the punishment for tampering, or through data security measures that make tampering more or less costly and more or less easy to detect.

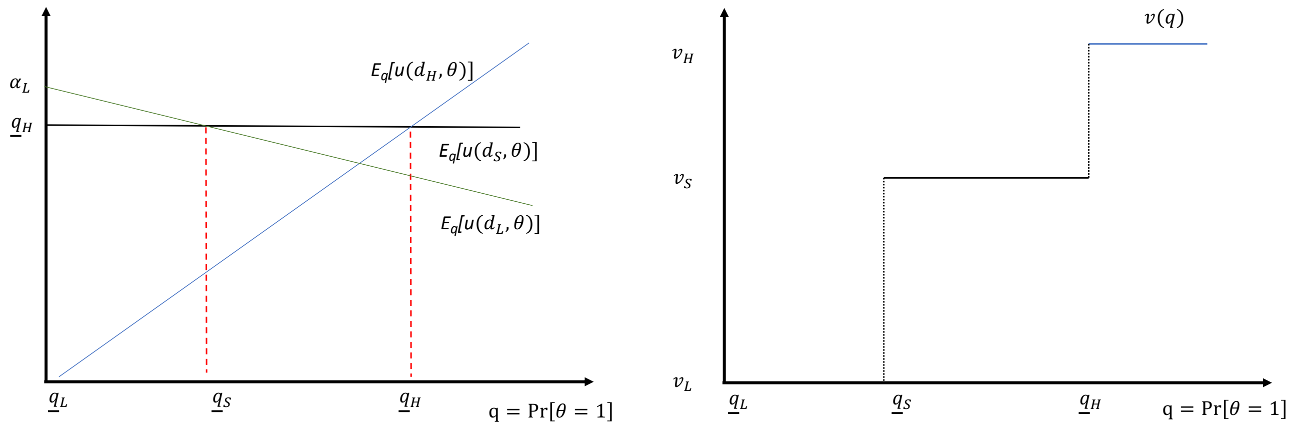

(a) Principal's indirect utility

(b) Agents' indirect utility

Figure 1: Principal and Agents' payoffs

1]. The principal selects d from $\{d_L, d_S, d_H\}$, and has preferences characterized by $u(d, \theta)$, [11]

$$u(d,\theta) = \begin{cases} \underline{q}_H & \text{for } d = d_S, \\ \theta & \text{for } d = d_H, \\ \alpha_L & \text{for } d = d_L \text{ and } \theta = 0, \\ \alpha_L - \mathbb{I}_{\{\alpha_L \geq \underline{q}_H\}} \frac{\alpha_L - \underline{q}_H}{\underline{q}_S} & \text{for } d = d_L \text{ and } \theta = 1, \end{cases}$$

with $0 = \underline{q}_L \leq \underline{q}_S < \underline{q}_H < 1$. In words, the principal decides either to keep the status quo d_S, to scale-up operations by choosing d_H, or to scale-down by choosing d_L–Figure 1 represents the principal's payoff as a function of her posterior q. We will consider two cases. First, if $\alpha_L < \underline{q}_H$, then the principal "organizes to innovate" as she effectively chooses between d_S or d_H (i.e., whether to approve the "innovation" d_H), selecting d_H only if $q \in \left[\underline{q}_H, 1\right]$. Second, if $\alpha_L \geq \underline{q}_H$, then the principal "organizes for scale:" if she deviates from d_S, then she could either scale-up (d_H), or scale-down (d_L), selecting d_L only if $q \in \left[\underline{q}_L, \underline{q}_S\right]$. In either case, $\underline{q}_i$ represents the minimum posterior belief for which the principal still selects d_i.

We capture the conflict of interest between the agents and the principal by positing that the designer and the analyst receive a state-independent payoffs $v(d_i, \theta) = v_i$ with $0 = v_L < v_S < v_H$, so that they benefit from persuading the principal to choose "higher" decisions. To focus on the more interesting case, we assume that $\mu \in (\underline{q}_S, \underline{q}_H)$ so that the

[11] $\mathbb{I}_A$ represents the indicator function of the set A; i.e., $\mathbb{I}_A(x) = 1$ if $x \in A$ and $\mathbb{I}_A(x) = 0$ if $x \notin A$.

principal retains the status quo in the absence of any report, and we let $\Delta_i \equiv v_H - v_i$, $i \in \{L, S\}$ be an agent's gain from inducing his preferred decision d_H when the alternative is d_i.

Strategic Experimentation, Communication and Tampering: All players process information according to Bayes' rule. We consider three stages to this model of data analytics.

First, in the experimental-design stage the data designer (data architect) specifies which data to gather and how it will be processed: he selects an experiment π, consisting of a finite outcome space $S(\pi)$ and a family of likelihood functions over $S(\pi)$, $\{\pi(\cdot|\theta)\}_{\theta\in\Theta}$, with $\pi(\cdot|\theta) \in \Delta(S(\pi))$. Given the common prior, we can without loss set $S(\pi) \subset \Delta(\Theta)$, so that $\pi = \{q, \Pr[q]\}_{q\in S(\pi)}$ is expressed as a distribution over posterior beliefs q induced by observing the experimental outcome, with $S(\pi)$ indexing these outcomes. We say that the designer "experiments more" when he selects a Blackwell-more informative experiment. We make two important assumptions regarding experimental design. First, as in KG, the designer can choose *any experiment* that is correlated with the state. Second, experiments are costless to the designer. This can be the case, for instance, if a fully informative experiment is originally available to the designer and he can garble its outcome at no cost.

Second, the design stage is followed by an analysis/communication stage. The analyst privately observes the outcome $s \in S(\pi)$–we refer to s as the analyst's "type"– and sends a message $m \in S(\pi)$ to the principal, which is potentially subject to misrepresentation: the analyst can tamper with the true outcome s by reporting instead $s' \in S(\pi), s' \neq s$. We will work with a reduced-form model of tampering: the analyst incurs a cost c if he tampers, with c unknown at the design stage and distributed according to $F(c)$, and independent of the experiment π. This can be micro-founded as follows. We see tampering as a question of "opportunity:" there are potentially many different tampering methods $a \in A$ each carrying a different cost $c(a)$. These costs can be physical costs–e.g., effort in "doctoring the books" or "creating a credible alternative story"–tied to punishments if caught misrepresenting–with the severity of the punishment varying with the tampering method–or even psychic costs of misrepresentation.[12] However, only at the analysis/communication stage does the

[12] Gneezy (2005) and Abeler, Nosenzo and Raymond (2020) show experimentally that individuals have some innate preference for honesty.

analyst learn which subset of tampering methods $A' \subset A$ are actually feasible,[13] selecting the one with the lowest cost if he tampers. From an ex ante perspective we would then have $F(c) = \Pr[\min_{a \in A'} c(a) \leq c]$. We will assume that F is absolutely continuous, with $\bar{F}(c) \equiv 1 - F(c)$.[14]

We make two assumptions regarding these tampering costs. First, they are always borne by the analyst when he tampers. In Section 6.1, we show that the main insights of our analysis hold if the analyst incurs the cost c only if tampering is uncovered through auditing. Second, the analyst bears the same cost independently of the actual message sent. In other words, the decision of how to misrepresent the state only depends on the equilibrium inference of the principal, rather than the costs/punishments specifically associated to different messages.

In the third stage, the decision making stage, the principal observes both the designer's experiment and the analyst's message. Key in our model is the principal's ability to evaluate the truthfulness of this message, and undo the effect of any misrepresentation, by auditing the experiment. We assume that with probability λ the audit is conclusive and the principal learns the actual experimental outcome, while with probability $1-\lambda$ the audit is inconclusive and she gains no new information. Importantly, what can be learned from an audit is constrained by the informativeness of π. Thus, auditing differs from seeking a "second opinion" in which the principal may have access to a separate information source.[15] If the audit is conclusive, the principal is informed (of s) and selects (a possibly mixed) $d_I(m, s)$ which depends on the message m and the outcome s. If the audit is inconclusive, she is uninformed and selects $d_U(m)$. To lighten the exposition, we say that "the message/recommendation is (un)audited" when the audit is (in)conclusive.

Organizational Design: Agents perform two tasks—experimental design and analysis—

[13]The fact that some actions may not be available ex post is similar to the modeling of productive activities in Chassang (2010).

[14]Our modeling approach captures the idea of uncertain returns from tampering owing to uncertain tampering costs. An alternative way to micro-found our model is to consider that there is only one tampering action, with deterministic cost c', but the decision-payoff to the sender is given by $\alpha v(d)$ where α is unknown at the time of the design but commonly known to be distributed according to F'. Then, this model is equivalent to ours by defining $c \equiv c'/\alpha$.

[15]See, for instance, Kolotilin, 2018, Kolotilin, Mylovanov, Zapechelnyuk, and Li, 2017, and Guo and Shmaya, 2019 for information-design models where the receiver is privately informed.

and the principal has several organizational levers to incentivize them. First, she sets the auditing intensity $\lambda \in [0,1]$. For instance, she can assign resources at the outset that are used later to audit the analyst's report, thus, dictating the likelihood of a conclusive audit. An important assumption is that the principal can commit to an imperfect audit, i.e., to $\lambda < 1$. Otherwise, once the designer selects an experiment, she can completely eliminate the effect of tampering by perfectly auditing the analyst's message. In anticipation of a perfect audit, however, the designer would select the perfect commitment experiment as in KG. To see this, and for future reference, consider experiments $\{\underline{q}_L, \underline{q}_H\}$ and $\{\underline{q}_S, \underline{q}_H\}$ with

$$p_i^C \equiv \Pr[s = \underline{q}_i] = \frac{\underline{q}_H - \mu}{\underline{q}_H - \underline{q}_i}, \tag{1}$$

the probability of outcome $s = \underline{q}_i$, $i = L, S$. If $\lambda = 1$, then the designer selects the experiment $\{\underline{q}_i, \underline{q}_H\}$ that minimizes $p_i^C \Delta_i$. If the organization hopes to induce more experimentation, it must be able to guarantee that the success rate of an audit is limited to λ. Our interpretation is that further resources cannot be deployed once the auditing intensity is announced, so that λ cannot be increased neither in reaction to the chosen experiment, nor to the reported outcome.

Second, the principal can choose to either integrate design and analysis/reporting, by letting the same agent perform both tasks, or to separate them, by allocating each to a different agent.[16] Let k denote the principal's task allocation, with $k \in (\mathcal{I}, \mathcal{S})$. Instead of changing the number of agents for each task allocation, we keep our designer-analyst-principal game throughout all task allocations and assume that the designer also bears the tampering costs incurred by the analyst under integration ($k = \mathcal{I}$), while he does not bear them under separation ($k = \mathcal{S}$). In terms of organizational structure, task separation would correspond to a firm in which experimental design is centralized in a corporate headquarters and the designer mandates each operating unit which analysis to perform, while the actual data collection and reporting is decentralized to those units.[17]

[16] A maintained assumption of our analysis is that task allocation does not affect the agents preferences over decisions. That is, task allocation cannot be used to reduce the conflict of interest between principal and agents.

[17] For instance, the design of customer surveys or the specification of which data to be collected by ERP systems could be performed by an enterprise-wide data architect, while the analysis of the results is performed at the divisional level. Integration would have both tasks been decentralized to lower level units, so that

Timing: The principal selects whether to separate or integrate tasks and the auditing intensity λ. Then, the designer publicly selects $\pi = \{q, \Pr[q]\}_{q \in S(\pi)}$—this is the design subgame. The communication subgame follows: Nature draws θ and the analyst privately observes outcome $s \in S(\pi)$, generated according to π, and the cost c, and selects a message $m \in S(\pi)$. The principal observes the actual outcome s with probability λ and, given the outcome of the audit and the analyst's message, she updates her beliefs according to Bayes' rule, selects a decision, payoffs are realized and the game ends. We look for Perfect Bayesian Equilibria that constitute a Perfect Bayesian Equilibrium in every subgame.

3 Tampering and Equilibrium Experimentation

We start the organizational-design analysis by studying how agents respond to a given organizational structure. That is, we study the equilibria in the *designer-analyst-principal* game corresponding to a fixed task allocation and auditing intensity. We work backwards by first characterizing the equilibria in the communication subgame for any π, which will determine both expected tampering and the distribution over the principal's decisions. We then turn to the designer's optimal choice by introducing the set of robust experiments—a reduced set of experiments with binary outcomes which contains a solution to the designer's problem.

3.1 Equilibrium Tampering

The analyst decides whether to tamper by comparing the gain from misrepresenting his type to the realized tampering cost. Let $\bar{\Delta}$ be the maximum gain from tampering in any equilibrium (both on- and off- the equilibrium path).[18] Assumption 1 ensures a positive

local agents have discretion in deciding which data to collect and which analysis to perform. As an example in the public sector, David Cameron created the Behavioral Insights Team (BIT) under the supervision of the Cabinet Office (see Alonso and Câmara, 2016 for details). In an example of task-integration, the BIT would both design and conduct small-scale experiments for the UK Government.

[18] If the audit is inconclusive, the largest gain from tampering comes from inducing the principal to select d_H when truthful communication would have led to his least preferred decision. Tampering incentives are also shaped by the principal's decision after a conclusive audit: if the analyst's type is a "threshold" type $\underline{q}_i, i = \{S, H\}$, then the principal may punish/reward him after a conclusive audit by randomizing differently between decisions as a function of his report. Then, the maximum gain from tampering is $\bar{\Delta} = (1 - \lambda)\,\Delta_S$

probability of truthful reporting for every experiment and experimental outcome, by positing the existence of tampering costs that make tampering unprofitable for every analyst's type.

Assumption 1 (All messages on-path) *The tampering cost distribution satisfies*

$$\bar{F}(\bar{\Delta}) > 0. \tag{2}$$

In the absence of threshold types–i.e., whenever $\underline{q}_i \notin S(\pi), i = \{S, H\}$–the analyst that tampers sends a message that induces the "highest" decision after an inconclusive audit. Thus, if different types choose to tamper by sending different messages, it must be that they all induce the same unaudited decision (or mixtures over decisions). With this observation, we now characterize tampering behavior in the communication subgame.

Proposition 1. *Suppose that Assumption 1 holds. Then, in any equilibrium of the communication subgame following the choice* $\pi = \{q, \Pr[q]\}_{q \in S(\pi)}$ *we have:*
(i) For each $q \in S(\pi)$*, there exists* $\bar{c}(q)$*, with* $\bar{F}(\bar{c}(q)) > 0$*, such that* $m^*(q, c) = q$ *if* $c > \bar{c}(q)$ *and* $m^*(q, c) \neq q$ *if* $c < \bar{c}(q)$*;*
(ii) Let $M_T(\pi) \subset S(\pi)$ *be the set of "tampered outcomes:"*

$$M_T(\pi) = \{q \in S(\pi) : \exists (q_z, c),\ m^*(q_z, c) = q, q_z \neq q\}.$$

If $d_I(m, \underline{q}_i)$ *is independent of* m *whenever* $\underline{q}_i \in S(\pi), i = \{S, H\}$*, then for* $q, q' \in M_T(\pi)$ *(a)* $d_U(q) = d_U(q')$*, and (b)* $\bar{c}(q) = 0$*.*

Assumption 1 guarantees that all messages in $S(\pi)$ are sent with positive probability on the equilibrium path. This limits the scope of the principal to discipline the analyst's tampering by holding "optimistic or pessimistic" beliefs after an off-the-equilibrium-path message. Then, Proposition 1-i shows that the analyst's tampering behavior is monotonic: he reports truthfully if the realized cost exceeds an outcome-dependent threshold, $\bar{c}(q)$, and will surely tamper if the cost falls below this threshold. Proposition 1-ii(a) makes formal

---

if the principal organizes to innovate, but $\bar{\Delta} = \Delta_L - \lambda \Delta_S$ if she organizes for scale. To see this last case, consider the experiment $\left\{\underline{q}_S, q\right\}$ with $q > \underline{q}_H$, and the following sequentially-rational decision making: after a conclusive audit, the principal selects d_H if $s = q$ but if $s = \underline{q}_S$ she selects d_L if the analyst truthfully reported $m = \underline{q}_S$ and d_S if he tampered $m = q$. If the audit is inconclusive, and the only tampered outcome is q, the principal's consistent belief after an unaudited $m = \underline{q}_S$ must be precisely $\underline{q}_S$, in which case she selects d_L. Then, the gain from tampering after $s = \underline{q}$ is $\lambda v_S + (1 - \lambda) v_H - v_L = \Delta_L - \lambda \Delta_S$.

the above-mentioned property that "tampered outcomes"–messages that are transmitted by some other type with positive probability—may induce different posterior beliefs but must all lead to the same unaudited mixture over decisions; this is true as long as the principal does not condition his audited decision on the analyst's message. Additionally, there shouldn't be any gain from tampering for a "tampered-outcome" type; that is, types that others would like to mimic always report truthfully. This is Proposition 1-ii(b).

3.2 Optimal and Robust Experimentation

The fact that all messages are on-path does not rule out multiple equilibria of the communication subgame as the principal may randomize after a conclusive or inconclusive audit. For instance, if the designer selects $\{\underline{q}_S, q\}$ with $q > \underline{q}_H, \underline{q}_S > 0$, then the principal could choose different (mixed) decisions as a function of the analyst's report after a conclusive audit yields $s = \underline{q}_S$, or as a function of the audit after the analyst reports $m = \underline{q}_S$. Thus, the designer's payoff from experiment $\{\underline{q}_S, q\}$ varies with the equilibrium in the communication subgame.

To characterize the designer's equilibrium choice, we introduce the set of robust experiments $\Pi_i,\ i \in \{L, S\}$. These are binary experiments of the form $\{\underline{q}_i, q\}$ such that the principal when indifferent–e.g., after message $m = \underline{q}_i$ or after a conclusive audit determines $s = \underline{q}_i$–selects the most favorable action to the analyst.

Definition (Robust Experiments) *Define the class $i \in \{L, S\}$ of robust experiments Π_i indexed by $\tau \in [0, 1]$, as*

$$\Pi_i = \left\{ \pi = \{\underline{q}_i, q\} : q \geq \underline{q}_H \text{ and } \bar{F}\left(\bar{c}_i(\tau)\right) = p_i^C \left(\frac{q - \underline{q}_i}{q - \mu} \right), \tau \in [0, 1] \right\} \tag{3}$$

with $\bar{c}_i(\tau)$ given by

$$\bar{c}_S(\tau) = \tau\,(1 - \lambda)\,\Delta_S \text{ and } \bar{c}_L(\tau) = \bar{c}_S(\tau) + (1 - \lambda)\left(\Delta_L - \Delta_S\right),\ \tau > 0, \tag{4}$$

and $\bar{c}_i(\tau) = 0$ if $\tau = 0$, alongside the equilibrium behavior

$$d_U\left(\underline{q}_i\right) = d_I\left(m, \underline{q}_i\right) = d_i; d_I\left(m, q\right) = d_H; d_U\left(q\right) = \tau d_H + (1 - \tau) d_S.$$

Finally, let $\Pi \equiv \Pi_L \cup \Pi_S$ denote the set of robust experiments.

We will refer to Π_S as "status-quo" (robust) experiments and to Π_L as "up-or-down" (robust) experiments. In equilibrium, an inconclusive audit leads the principal to posterior

$\underline{q}_H$ following $m = q$ and she selects d_H with probability τ and d_S with probability $1-\tau$. This also determines the incentives to tamper after an unfavorable outcome—which given scale-up probability τ translate to thresholds (4)—and the probability $\bar{F}\left(\bar{c}_i(\tau)\right)$ that following $s = \underline{q}_i$ the analyst reports truthfully−see (3). This also implies that the set Π_i depends on the auditing intensity λ; for instance, if $\lambda = 1$ then the commitment experiment is the only robust experiment–i.e, $\Pi_i = \{\underline{q}_i, \underline{q}_H\}$. Moreover, experiments in Π_i can be equivalently indexed by either (i) the posterior q, (ii) the probability of outcome $s = \underline{q}_i$, with $\Pr\left[s = \underline{q}_i\right] = \frac{q-\mu}{q-\underline{q}_i}$, (iii) the scale-up probability τ, or (iv) the induced tampering threshold $\bar{c}_i$. Finally, all experiments in Π_i are ordered according to their informativeness: trivially, an experiment with a higher q (equivalently higher $\bar{c}_i$, higher τ, or higher $\Pr[s = \underline{q}_i]$) corresponds to a Blackwell-more informative experiment. In what follows, we refer to the class of "up-or-down" experiments Π_L as the "more informative" class.[19]

We now present our main equilibrium characterization. For $i \in \{L, S\}$, let $v_i(\tau, \mu; \lambda, k)$ be the designer's equilibrium payoff in a communication subgame after he selects $\pi_i(\tau) \in \Pi_i$, with $k \in \{\mathcal{S}, \mathcal{I}\}$ the principal's task allocation and λ her auditing intensity, and let

$$V_i(\mu; \lambda, k) \equiv \max_{\tau \in [0,1]} v_i(\tau, \mu; \lambda, k), \tag{5}$$

$$\bar{V}(\mu; \lambda, k) \equiv \max\left\{V_S(\mu; \lambda, k), V_L(\mu; \lambda, k)\right\}, \tag{6}$$

with $V_i(\mu; \lambda, k)$ the designer's maximum expected payoff when restricted to Π_i, and $\bar{V}(\mu; \lambda, k)$ his maximum expected payoff when choosing a robust experiment.

Proposition 2. *Let $\lambda > 0$ and $\mu \in (\underline{q}_S, \underline{q}_H)$. Then,*
(i) there is always an equilibrium of the design subgame in which the designer selects a robust experiment,
(ii) if the designer obtains payoff V^* in some equilibrium of the design subgame, then (a) $V^* = \bar{V}(\mu; \lambda, k)$ when the principal organizes to innovate, and (b) $V^* \geq \bar{V}(\mu; \lambda, k)$ when the principal organizes for scale. If $V^* > \bar{V}(\mu; \lambda, k)$ and π^* is an equilibrium experiment, then $\underline{q}_S \in S(\pi^*)$.

Proposition 2 justifies our restriction to robust equilibria when analyzing the principal's organizational design problem. This is based on two observations. First, there is always

[19]This terminology captures the fact that for each experiment in Π_S there is always an experiment in Π_L that is Blackwell-more informative.

an equilibrium in which the designer responds by selecting an experiment in Π following the principal's organizational choice. We prove this claim in the appendix by constructing from an arbitrary $\pi' = \{q, \Pr[q]\}_{q \in S(\pi')}$–and considering communication equilibria in which the principal selects the designer's most favorable action after a conclusive audit–a robust experiment that gives the designer a (weakly) higher payoff. Our construction involves two steps: we first show that there is an equilibrium of an experiment supported on $\left\{\underline{q}_L, \underline{q}_S, q\right\}$, $q \geq \underline{q}_H$, that (weakly) improves the designer's payoff obtained from π'. We then show that the designer's payoff is quasiconvex when restricted to the (convex) set of experiments supported on $\left\{\underline{q}_L, \underline{q}_S, q\right\}$. This ensures that the maximum payoff is achieved at one of the extreme points–i.e., at either $\{\underline{q}_L, q\} \in \Pi_L$ or $\{\underline{q}_S, q\} \in \Pi_S$.

Second, Proposition 2-ii shows that the designer can always guarantee himself $\bar{V}(\mu; \lambda, k)$ in any equilibrium. Indeed, experiment $\{\underline{q}_S + \varepsilon, q\}$, with ε small, results in a unique designer's payoff as the principal never randomizes after a conclusive audit. This uniqueness of equilibrium payoff guarantees that $V^* \geq \bar{V}(\mu; \lambda, k)$ in any equilibrium. In fact, when organizing to innovate, all equilibria of the design subgame give him the same expected payoff; thus, to find the designer's optimal payoff we can restrict attention to robust experiments. This is not the case when organizing for scale: if $\underline{q}_S \in S(\pi)$, the principal could use her indifference after a conclusive audit shows $s = \underline{q}_S$ to minimize tampering by announcing that she would treat favorably truth-telling and unfavorably tampering. In fact, whenever $\lambda \geq \Delta_S/\Delta_L$ the designer can obtain the commitment payoff from experiment $\{\underline{q}_S, \underline{q}_H\}$–this is the case if the principal threatens to implement decision d_L if the analyst is caught tampering. This proves that $V^* > \bar{V}(\mu; \lambda, k)$ in this case. We differ to Section 6.2 a discussion of organizational design with non-robust experiments.

3.3 Designer's Equilibrium Payoffs

To solve for the designer's optimal experiment using Proposition 2, we now characterize $v_i(\tau(\bar{c}), \mu; \lambda, k)$–the designer's payoff as a function of the induced tampering threshold. To this end, define $\eta(\bar{c})$ as the product of the expected tampering cost conditional on tampering times the odds of tampering,

$$\eta(\bar{c}) \equiv E[c | c \leq \bar{c}] \frac{F(\bar{c})}{\overline{F}(\bar{c})} = \frac{\int_0^{\bar{c}} c dF(c)}{\overline{F}(\bar{c})}. \tag{7}$$

Lemma 1. *For $\mu \in \left[\underline{q}_i, \underline{q}_H\right]$, consider $\{\underline{q}_i, q\} \in \Pi_i$ that induces threshold $\bar{c}$. Then,*

$$v_i(\tau(\bar{c}), \mu; \lambda, k) = v_H - (1-\lambda)\Delta_i + \bar{c} - m_i(\bar{c}; \lambda, k)(\underline{q}_H - \mu), \tag{8}$$

with m_i the slope of the designer's payoff with respect to the prior μ:

$$m_i(\bar{c}; \lambda, k) \equiv \frac{\Delta_i}{\underline{q}_H - \underline{q}_i}\left(\frac{\lambda}{\bar{F}(\bar{c})} + \frac{\bar{c}}{\Delta_i} + \mathbb{I}_{\{k=\mathcal{I}\}}\frac{\eta(\bar{c})}{\Delta_i}\right). \tag{9}$$

To understand (8-9), for each robust experiment $\{\underline{q}_i, q\} \in \Pi_i$ we introduce an equivalent "full commitment" game with a "modified" designer's utility so that the designer's payoff from $\{\underline{q}_i, \underline{q}_H\}$ gives him the same expected payoff–see Figure 2.[20] This allows us to appeal to the geometric intuition of the full commitment case when studying the designer's preferences in the imperfect commitment case. Fix $\pi(\bar{c}) \in \Pi_i$ which induces threshold $\bar{c}$, and define the task allocation-dependent utility $\tilde{v}_{i,k}$ for d_i,

$$\tilde{v}_{H,k} \equiv v_H - (1-\lambda)(1-\tau(\bar{c}))\Delta_S, \tag{10}$$

$$\tilde{v}_{i,k} \equiv v_i - \lambda\Delta_i\frac{F(\bar{c})}{\bar{F}(\bar{c})} - \mathbb{I}_{\{k=\mathcal{I}\}}\eta(\bar{c}),\ i \neq H. \tag{11}$$

Figure 2 describes the relation between $v(q)$ in the original game and the indirect utility $\tilde{v}_k(q)$ in the equivalent "full commitment" game.

To show payoff-equivalence, start with the designer's payoff under full commitment from $\{\underline{q}_i, \underline{q}_H\}$, $\left(1 - p_i^C\right) v_H + p_i^C v_i$, with p_i^C given by (1). Equilibrium tampering alters this payoff in two ways. First, upon observing $\underline{q}_H$, the principal now keeps the status quo with probability $(1-\lambda)(1-\tau(\bar{c}))$ –this explains (10). Second, tampering and auditing change the distribution of outcomes so that the probability of observing $\underline{q}_i$ increases from p_i^C to $p_i^C/\bar{F}(\bar{c})$—see (3). This higher probability of a low outcome reduces the designer's payoff by

$$\lambda\left[\left(1-p_i^C\right)v_H + p_i^C v_i\right] - \lambda\left[\left(1 - \frac{p_i^C}{\bar{F}(\bar{c})}\right)v_H + \frac{p_i^C}{\bar{F}(\bar{c})}v_i\right] = \lambda\Delta_i\frac{F(\bar{c})}{\bar{F}(\bar{c})},$$

which explains (11) for $k = \mathcal{S}$. Finally, tampering costs are incurred if $s = \underline{q}_i$, which occurs with probability $p_i^C/\bar{F}(\bar{c})$, and explains (11) for $k = \mathcal{I}$.

[20] We cannot apply the concavification argument when $\lambda < 1$ as the probability of a message corresponding to a "tampered outcome," as well as the principal's posterior belief when unaudited, are determined in equilibrium by the analyst's global tampering behavior, which in turn depends on the entire distribution of experimental outcomes.

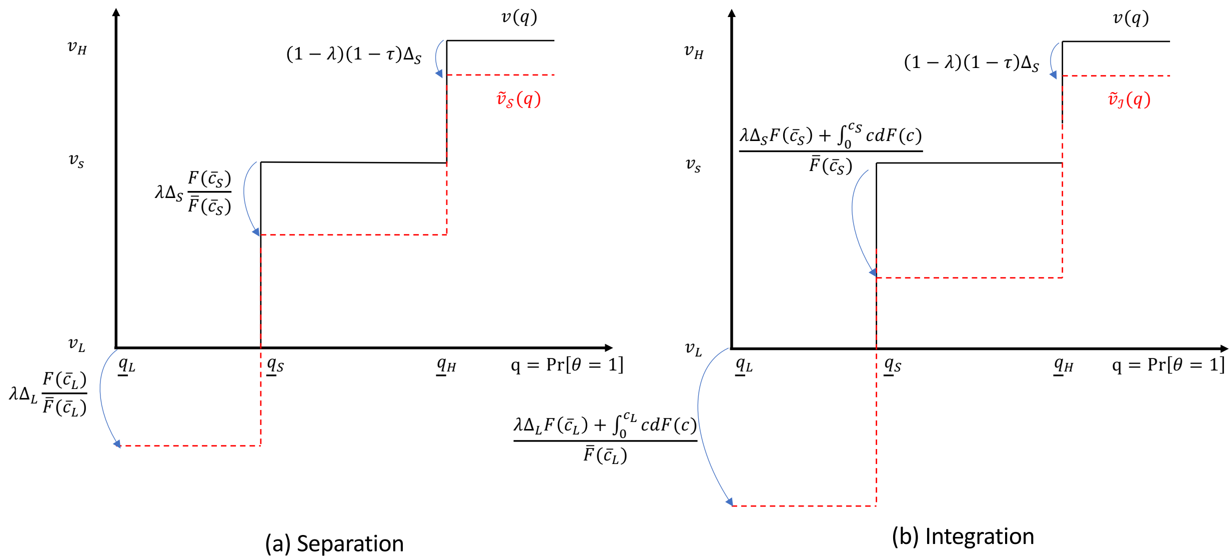

Figure 2: Payoffs for the equivalent "full commitment" game.

Expressions (10-11) capture the trade-off that the designer faces: a higher scale-up probability τ increases the payoff after an unaudited message q (outcome $\underline{q}_H$ in the equivalent "full commitment" game), but a higher scale-up probability can only result from a higher tampering threshold. This forces the designer to offer a more informative experiment to sustain the higher τ, thus increasing the likelihood of observing an unfavorable outcome.

Figure 3 depicts the designer's payoff under separation $V_S(\mu; \lambda, \mathcal{S})$ and $V_L(\mu; \lambda, \mathcal{S})$. $\bar{V}(\mu; \lambda, \mathcal{S})$ is then computed as the upper envelope of these payoffs. These graphs highlight two features that distinguish our model with tampering from a model with analyst's commitment. First, in Figure 3-a, $\{\underline{q}_S, \underline{q}_H\}$ is the optimal experiment if $\lambda = 1$, which remains optimal if $\lambda < 1$—this is a general result formally shown in Proposition 4-ii(a). However, the equilibrium expected payoff of the designer is now strictly convex in the prior for $\mu \in \left[\underline{q}_S, \underline{q}_H\right]$. This is a reflection that he actually changes the experiment (switching to one with a higher τ and hence higher $\bar{c}$) as μ increases. In fact, as shown in Figure 3-b, a higher prior can lead to switches in the optimal class of experiment. Indeed, in spite of an "up-or-down" experiment being optimal in Figure 3-b when $\lambda = 1$, if $\lambda < 1$ then the designer may actually switch to a "status-quo" experiment–see Section 5.2.

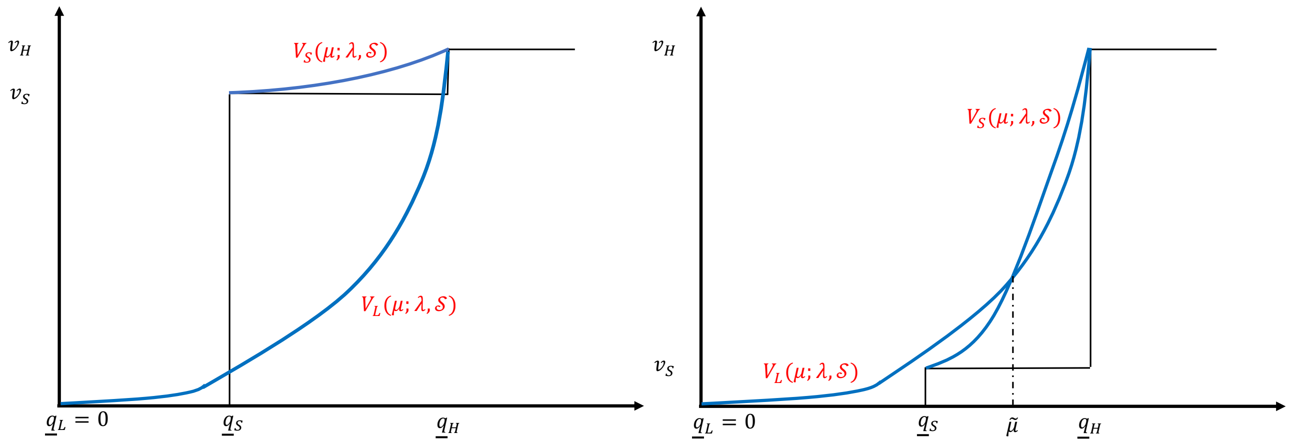

(a) Status-quo experiment optimal if $\lambda = 1$.

(b) Up-or-down experiment optimal if $\lambda = 1$.

Figure 3: Designer's equilibrium payoff for different experiment classes.

4 Motivating Informative Experimentation

Firms have two often-conflicting goals when managing data analytics: to encourage experimentation by their members and to foster faithful reporting of their findings. In this section, we focus on the first goal and show how different organizational levers—in particular, whether to integrate or separate design and analysis, and how much to audit the analyst's report—affect the informativeness of experimentation.

For the remainder, let $\{\underline{q}_i, q_i^*(\lambda, k)\} \in \Pi_i$ be the designer's optimal experiment under a k-allocation when restricted to Π_i, with $\bar{c}_i^*(\lambda, k)$ the induced tampering threshold, and let $i^*(\lambda, k)$ be the class of the designer's optimal experiment. We show below that separating tasks or decreasing auditing intensity always increases within-class experimentation—so that $q_i^*(\lambda, k)$ increases for fixed i—but this may lead to adverse class switches—e.g., the designer switching to an experiment in the "less informative" class Π_S. With these insights, we then solve the principal's organizational design problem in Section 5.

4.1 Motivating Experimentation through Task Allocation

Assigning the tasks of experimental design and analysis to the same agent forces him to economize on tampering costs when choosing an experiment; for instance, he may sacrifice scaling-up probability in order to reduce the incentives to tamper—this will always be the case when he is restricted to an experiment in Π_i— or may switch the class of the experiment. To understand how internalizing tampering costs affects the choice of experiment, we first study how equilibrium tampering costs vary across experiments in Π_S and Π_L.

Consider experiments $\pi_L(\tau) \in \Pi_L$ and $\pi_S(\tau) \in \Pi_S$ that induce the same scale-up probability after an inconclusive audit. From (11), expected tampering costs are

$$C_i(\tau) \equiv \frac{p_i^C}{\bar{F}(\bar{c}_i(\tau))} \int_0^{\bar{c}_i(\tau)} c dF(c) = p_i^C \eta(\bar{c}_i(\tau)), \tag{12}$$

with $\eta(c)$ defined by (7). Conditional on an unfavorable outcome—$\underline{q}_L$ in the case of $\pi_L(\tau)$ and $\underline{q}_S$ in the case of $\pi_S(\tau)$—expected tampering costs are higher for $\pi_L(\tau)$ as the analyst faces a larger gain from tampering; so we must have $\eta(\bar{c}_L(\tau)) > \eta(\bar{c}_S(\tau))$. Nevertheless, the probability of an unfavorable outcome is also lower for $\pi_L(\tau)$ −which implies that $p_L^C < p_S^C$. Therefore, expected tampering costs may not increase, and actually decrease, when the designer moves from experiment $\pi_S(\tau)$ to experiment $\pi_L(\tau)$. However, if

$$\inf_{\tau\in[0,1]} \frac{\eta(\bar{c}_L(\tau))}{\eta(\bar{c}_S(\tau))} \geq \frac{p_S^C}{p_L^C}, \tag{13}$$

then $C_L(\tau) \geq C_S(\tau)$ for all $\tau \in [0,1]$.[21] With these insights, and recalling from (5) that $V_i(\mu; \lambda, k)$ is the designer's maximum expected payoff from an experiment in Π_i, we now characterize situations in which integration reduces experimentation.

Proposition 3. *(i) We have* $q_i^*(\lambda, \mathcal{S}) \geq q_i^*(\lambda, \mathcal{I})$.

(ii) Let m_i *be the slope of the designer's payoff defined in (9). If (ii-a)* $m_S(\bar{c}_S(\tau); \lambda, \mathcal{I}) - m_L(\bar{c}_L(\tau); \lambda, \mathcal{I})$ *is single-crossing in* $\tau \in [0,1]$ *(from negative to positive), and (ii-b) condition (13) holds; then* $V_S(\mu; \lambda, \mathcal{I}) > V_L(\mu; \lambda, \mathcal{I})$ *whenever* $V_S(\mu; \lambda, \mathcal{S}) > V_L(\mu; \lambda, \mathcal{S})$.

[21] Lemma 5 in the Appendix provides sufficient conditions for (13) to hold. For instance, (13) follows if $\eta(c)$ increases sufficiently rapidly (e.g., if $\eta(c)$ is convex and $p_L^C > (\Delta_S/\Delta_L) p_S^C$; or if $\eta(c)$ is log-convex in $[\underline{c}, (1-\lambda)\Delta_L]$) or if the relative gain $\Delta_L - \Delta_S$ is sufficiently large (e.g., if $\eta(c)$ is log-concave and $p_L^C \eta((1-\lambda)\Delta_L) \geq p_S^C \eta((1-\lambda)\Delta_S)$.

How would a concern with reducing tampering costs affect experimentation? First, if restricted to experiments in Π_i, reducing tampering costs can only be achieved through a reduction in scale-up probability,[22] implying that integration (weakly) reduces within-class experimentation–this is Proposition 3-i. Then, if task allocation doesn't change the class of the designer's optimal experiment, integration unambiguously reduces the informativeness of experimentation. Put differently, integration can lead to more informative experimentation only if it leads to a change in the optimal experiment class.

Second, (13) guarantees that the designer can reduce expected tampering costs while preserving scale-up probability by switching to a status-quo experiment. One would then conjecture that if the designer selects an status-quo experiment under task separation, he will certainly do so if tasks are instead integrated. However, if tasks are integrated the designer will seek to reduce scale-up probability, and for lower scale-up probabilities an "up-or-down" experiment may actually be preferable. Proposition 3-ii then provides a sufficient condition to rule out this case and preserve the designer's preferences for the "less informative" class Π_S when tasks are integrated.[23]

4.2 Motivating Experimentation through lax auditing

The standard rationale for auditing data analytics processes is both to ensure data integrity and to dissuade tampering. This remains true in our model: for a fixed experiment, increasing λ allows the principal to both shield herself after a conclusive audit from any misrepresentation of the data, and to reduce the incentives to tamper by making an inconclusive audit less likely. In fact, for a fixed experiment, increasing λ can only increase the information that reaches the principal.

However, once experimental design is delegated, varying λ also changes the designer's incentives to experiment. Indeed, reducing λ both: (i) changes the set of robust experiments,

[22] Decreasing scale-up probability lowers both the equilibrium tampering threshold and the likelihood that a tampering outcome occurs, leading to lower expected tampering costs.

[23] We actually prove the contrapositive in the proof of Proposition 3-ii: if $V_L(\mu;\lambda,\mathcal{I}) \geq V_S(\mu;\lambda,\mathcal{I})$ then $V_L(\mu;\lambda,\mathcal{S}) \geq V_S(\mu;\lambda,\mathcal{S})$. The single crossing condition on the slope guarantees that $V_L(\mu;\lambda,\mathcal{I}) - V_S(\mu;\lambda,\mathcal{I})$ is single-crossing, and that moving to separating tasks increases both the scale-up probability and the difference $C_L(\tau) - C_S(\tau)$, both changes favoring "up-or-down" experiments.

allowing for more informative experiments; and (ii) (weakly) reduces the principal's scale-up probability for each experiment. To understand the overall effect on experimentation, the following proposition decomposes the designer's problem into two subproblems. First, we analyze how auditing changes within-class experimentation. Second, we study the optimal class $i^*(\lambda, k)$. This decomposition uncovers the two countervailing effects that auditing has on equilibrium experimentation: moving away from a perfect audit leads the designer to seek stronger evidence supporting decision d_H but (possibly) weaker evidence favoring d_L.

Proposition 4. *(i) Fix $\mu \in [\underline{q}_i, \underline{q}_H]$, $i \in \{L, S\}$ and $k \in \{\mathcal{S}, \mathcal{I}\}$. Then,*

(i-a) $v_i(\tau(\bar{c}), \mu; \lambda, k)$ is supermodular in $(\bar{c}, -\lambda)$.

(i-b) Outcome $q_i^*(\lambda, k)$ is non-increasing in λ.

(ii) Suppose that $p_L^C \geq (\Delta_S/\Delta_L)\, p_S^C$–so that the designer selects a status-quo experiment under a perfect audit. Then, for $\lambda > 0$:

(ii-a) If tasks are separated, then $V_S(\mu; \lambda, \mathcal{S}) \geq V_L(\mu; \lambda, \mathcal{S})$.

(ii-b) If tasks are integrated and (13) holds, then $V_S(\mu; \lambda, \mathcal{I}) \geq V_L(\mu; \lambda, \mathcal{I})$.

Proposition 4-i shows that, when restricted to Π_i, increasing λ : (a) reduces the designer's incremental payoff from an experiment with a higher tampering threshold; and (b) leads the designer to select less informative experiments. To see this, consider the marginal change in the designer's payoff (8) when, as a result of intensifying auditing, scale-up probability increases to preserve the same tampering threshold $\bar{c}$. The designer's marginal expected payoff conditional on an inconclusive audit decreases by $v(\underline{q}_i)$, but increases after a conclusive audit by $\Delta_i \Pr[s = q] + v(\underline{q}_i)$. Overall, the designer's marginal payoff is proportional to the probability that π generates a favorable outcome, which decreases with $\bar{c}$. Therefore, increasing λ reduces the designer's marginal payoff from more informative experiments. A more intense auditing also changes the set of robust experiments, however. Proposition 4-i(b) then shows that the combined effect of increasing λ—i.e., lower incentives to experiment but more informative experiments available in Π_i−unambiguously discourages experimentation. It follows that whenever the class of the designer's experiments is the same for two different auditing intensities, the lower auditing intensity induces a Blackwell-more informative experiment.

Proposition 4-ii shows that the designer's preference for the less informative class are preserved when moving from a perfect audit to a lax audit.[24] An important implication is that task separation is ineffective as a tool to induce more information about the "downscale" decision d_L if the designer is not willing to provide this information under a perfect audit.

4.2.1 Designer's responsiveness to Auditing

We end this Section by asking: when can lax auditing increase experimentation? To this end, we introduce the notion of designer's responsiveness to auditing.

Definition (Responsiveness) *The designer under a $k-$allocation is responsive to $\lambda-$auditing in class $i \in \{L, S\}$ if he strictly prefers experiment $\{\underline{q}_i, q\}$ to $\{\underline{q}_i, \underline{q}_H\}$ for some $q > \underline{q}_H$. With $\Lambda_i(k)$ the set of auditing intensities satisfying this condition, we say that the designer is responsive to auditing iff* $\Lambda_i(k) \neq \emptyset$ for some $i \in \{L, S\}$.

The following Lemma will be used in Section 5. The proof of the lemma also provides a sharp characterization of the sets $\Lambda_i(k), i = \{L, S\}$.

Lemma 2. *If $f(0) > 0$, then, regardless of the task allocation, the designer is responsive to auditing in Π_S.*

Two conditions ensure that the designer is responsive to auditing: (i) there must be experiments in Π_i for which the principal is willing to scale-up with positive probability, and (ii) the designer must gain when switching to one of these experiments. As the likelihood of high tampering costs can help the analyst commit not to tamper, the assumption $f(0) > 0$ ensures that the the principal will never scale-up after an inconclusive audit if the designer were to select experiment $\{\underline{q}_i, \underline{q}_H\}$. Moreover, in the class Π_S, the marginal payoff from an experiment with positive scaling-up probability can be made arbitrarily large by lowering auditing intensities. It follows that the designer is always responsive to auditing when restricted to "status-quo" experiments if $f(0) > 0$.[25]

[24] This is always the case if tasks are separated. If tasks are integrated, however, we must now consider how equilibrium expected tampering costs vary across classes. Nevertheless, condition (13) guarantees that expected tampering costs decrease when moving to a status-quo experiment with the same scale-up probability. Combined with Proposition 4-ii(a), we then have that preferences for a less informative class are also preserved under integration if expected tampering costs are also lower.

[25] On the other hand, for low λ no "up-or-down" experiment may have a positive probability of scaling-up as the analyst will always tamper if tampering costs fall below $\bar{c}_L(0)$. Then, the proof of the lemma provides

5 Organizational Design

We now turn to the issue of organizational design. To understand the trade-offs that the principal faces, consider her expected utility from experiment $\pi = \{\underline{q}_i, q\} \in \Pi_i$:

$$\begin{aligned} U(\lambda, k; \pi) &= \underline{q}_H + \Pr[s = q]\lambda(q - \underline{q}_H) + \\ &\quad + \Pr[s = \underline{q}_i][1 - (1-\lambda)\Pr[m = q | s = \underline{q}_i]](E_{\underline{q}_i}[u(d_i, \theta)] - \underline{q}_H) \end{aligned} \quad (14)$$

The principal benefits from π in two ways. First, she selects d_H after auditing a scale-up recommendation so that the more convincing the evidence in favor of d_H—the larger $q - \underline{q}_H$ is—the greater her gain. Second, when $\pi \in \Pi_L$, she selects d_L after $s = \underline{q}_L$ except when the analyst tampers and her audit is inconclusive. Therefore, the gain from scaling-down can be increased by discouraging tampering. Thus, she needs to strike a balance between inducing the designer to experiment more while restraining at the same time the analyst from tampering.

How do task allocation and auditing intensity help her resolve this trade-off? Using (14), we can write the principal's equilibrium expected utility before the design subgame:[26]

$$U(\lambda, k) = \underline{q}_H + \left(\underline{q}_H - \mu\right)\left(\lambda \frac{F(\bar{c}^*)}{\bar{F}(\bar{c}^*)} + \mathbb{I}_{\{i^*=L\}}\left(1 + \lambda \frac{F(\bar{c}^*)}{\bar{F}(\bar{c}^*)}\right)\left(\frac{\alpha_L - \underline{q}_H}{\underline{q}_H - \underline{q}_L}\right)\right). \quad (15)$$

This expression showcases our main insight: fostering experimentation while discouraging tampering are conflicting goals. The principal can always eliminate frictions in communication by perfectly auditing the experiment—for a fixed experiment and costless auditing, she will certainly do so—but this will reduce the information she receives regarding decision d_H. In fact, an imperfect audit allows her to credibly withhold scaling-up if the evidence in favor of d_H are not convincing, forcing the designer to provide more compelling evidence. Thus, she would like to incentivize tampering by having a "shadow of a doubt" on the claims of the analyst, but such skepticism can only be credible if $\lambda < 1$.

Expression (15) clarifies that, given equilibrium behavior, the principal is always willing to trade-off more distortions in communication for more informative experimentation. The

the set of auditing intensities that ensures that there are "up-or-down" experiments with $\tau > 0$.

[26] Recall that for a $k-$allocation and $\lambda-$auditing, $i^* = i^*(\lambda, k)$ is the class of the designer's optimal experiment, and $\bar{c}^* = \bar{c}^*_{i^*}(\lambda, k)$ the analyst's tampering threshold.

reason is twofold. First, she strictly benefits from selecting d_H only if the audit is conclusive, thus making any distortion in communication irrelevant. Second, an experiment with a higher tampering threshold $\bar{c}^*$ lowers her expected utility conditional on $s = \underline{q}_L$, but also makes this outcome more likely; the combined effect on her expected utility is proportional to $\left(1 + \lambda F\left(\bar{c}^*\right)/\bar{F}\left(\bar{c}^*\right)\right)$ which increases in $\bar{c}^*$.

We now study the optimal organizational structure by analyzing separately the case that the principal organizes to innovate and when she organizes for scale. We end this section by considering two additional organizational levers: the principal can specify the distribution of the analyst's tampering costs as well as restrict her choice set.

5.1 Organizing to Innovate

Suppose that the firm "organizes to innovate," so that the principal decides whether to "approve" d_H or retain the status-quo d_S. The absence of a scale-down option d_L means that there are no concerns regarding adverse class switches and robust experiments are always optimal—see Proposition 2. Then, task separation always increases within-class experimentation—see Proposition 3—and this explains the principal's preference for separating design and analysis. From (15), she gains from experimentation only if lax auditing compels the designer to select an experiment that will be tampered with positive probability. This indirect benefit of lax auditing leads the principal to lower the auditing intensity below the one she would set if she controlled experimental design. We thus reach one of our main results: if auditing is costless and the designer is responsive to auditing–see Lemma 2–then in every equilibrium the principal commits to an imperfect audit, i.e. $\lambda^* < 1$.

Proposition 5. *(i) Consider a subgame with $\lambda \in (0,1)$. Then, the principal prefers to separate tasks.*
(ii) Suppose that the principal can select λ at no cost, and let λ^* denote her equilibrium choice. Then, $\lambda^* < 1$ in every equilibrium if and only if the designer is responsive to auditing. In particular, if $f(0) > 0$, then $\lambda^* < 1$.

5.1.1 Tampering Incentives and Optimal Auditing

How much should the principal audit the analyst's report given that she separates tasks? To derive her optimal audit λ^*, we first characterize the designer's equilibrium experiment for any $\lambda \in (0,1)$. To this end, define $L(c) \equiv f(c)/\left(\bar{F}(c)\right)^2$ and $\bar{c}_{FR}$ implicitly by $\bar{F}\left(\bar{c}_{FR}\right) = p_S^C/\left(1-\mu\right)$, so that $\bar{c}_{FR}$ is the threshold induced by a fully informative experiment.[27]

Lemma 3. *Fix $\lambda \in (0,1)$ and suppose that $L(c) - (\phi_S/\lambda)$ is single-crossing in $[0, \Delta_S]$, from negative to positive, with $\phi_S \equiv \left(1 - p_S^C\right)/\left(p_S^C \Delta_S\right)$. Then, the designer under separation selects the commitment experiment, corresponding to $\bar{c}^* = 0$, if $L(0) \geq \phi_S/\lambda$. Otherwise, he selects an experiment that induces tampering threshold*

$$\bar{c}^* = \min\left[L^{-1}(\phi_S/\lambda), (1-\lambda)\,\Delta_S, \bar{c}_{FR}\right]. \tag{16}$$

Consistent with Proposition 4-i.b, the designer's optimal experiment induces less tampering, but is less informative, as auditing intensifies–the equilibrium tampering threshold (16) decreases with λ. The single-crossing condition on $L(c)$ guarantees that the designer's expected utility is quasiconcave in the tampering threshold and is always satisfied, for instance, if the hazard rate $f(c)/\bar{F}(c)$ is increasing. The equilibrium threshold $\bar{c}^*$ is the minimum of three possible choices. The term $\bar{c}_{FR} = \bar{F}^{-1}\left(p_S^C/\left(1-\mu\right)\right)$ corresponds to a fully informative experiment, while $(1-\lambda)\,\Delta_S$ corresponds to the case that the principal rubberstamps the analyst's recommendation. The first term in (16) reflects the designer's choice when it leads to a lower approval probability. In fact, if $L(c)$ is large—in particular, $L(0) \geq \phi_S/\lambda$—then the principal only approves when she audits and the designer's experiment induces $\bar{c}^* = 0$. Therefore, imperfect, albeit intense, auditing—specifically, when $\lambda \geq \phi_S/f(0)$—can still completely crowd-out valuable experimentation. This imposes an upper bound on the range of auditing intensities that the principal might entertain.

From (15), the principal's problem when organizing to innovate reduces to

$$\lambda^* \in \arg\max_{\lambda \in [0,1]} \lambda \frac{F\left(\bar{c}^*\right)}{\bar{F}\left(\bar{c}^*\right)}, \text{ with } \bar{c}^* \text{ given by (16).} \tag{17}$$

The optimal auditing will, in general, be sensitive to the cost distribution and preferences of agents. To illustrate (17), we study a case where tampering costs are uniformly distributed.

[27] Recall that when organizing to innovate, $\underline{q}_S = 0$.

Example: Uniform Distribution. Let $\Delta_S = 1$ with $v_S = 0$, and suppose that c is uniformly distributed in $[0,1]$, so that $\bar{F}[(1-\lambda)\Delta_S] = \lambda$. From (8), the designer's utility under task separation when the experiment induces $\overline{c} \in [0, 1-\lambda]$ is

$$v_S(\tau(\overline{c}), \mu; \lambda, S) = \lambda + \overline{c} - m_S(\overline{c}; \lambda, S)\left(\underline{q}_H - \mu\right) = \lambda + \left(1 - p_S^C\right)\overline{c} - p_S^C \frac{\lambda}{1-\overline{c}},$$

which is concave in $\overline{c}$. Denote by $\overline{c}_{crit} \equiv 1 - \sqrt{\lambda \frac{p_S^C}{1-p_S^C}}$ its unconstrained maximum. Then, mirroring (16), the designer's optimal experiment leads to a tampering threshold

$$\overline{c}^* = \min\left\{\max\left\{0, \overline{c}_{crit}\right\}, 1-\lambda, \frac{\mu}{1-\mu}\frac{1-\underline{q}_H}{\underline{q}_H}\right\}.$$

We can now fully characterize the equilibrium experiment as a function of λ. Recall that $\frac{1-p_S^C}{p_S^C}$ are the approval odds of the innovation if the principal were to set $\lambda = 1$. If $\lambda \geq \frac{1-p_S^C}{p_S^C}$, then $\overline{c}_{crit} \leq 0$ and the designer selects $\left\{0, \underline{q}_H\right\}$, i.e., selects the commitment experiment. If $\lambda \leq \frac{p_S^C}{1-p_S^C}$, then $\overline{c}_{crit} \geq 1-\lambda$ and the designer selects the most informative robust experiment. This would lead to either a fully informative experiment, or to an experiment for which the principal's rubberstamps the analyst's recommendation. Finally, if $\frac{p_S^C}{1-p_S^C} \leq \lambda \leq \frac{1-p_S^C}{p_S^C}$, then $\overline{c}^* = \min\left[\overline{c}_{crit}, \frac{\mu}{1-\mu}\frac{1-\underline{q}_H}{\underline{q}_H}\right]$ and the designer limits the informativeness of the experiment, leading to intermediate approval probabilities after an inconclusive audit.

Figure 4 describes two cases, with $\frac{1-p_S^C}{p_S^C}$ taking values 2 and 1/2.[28] If $\frac{1-p_S^C}{p_S^C} = 2$, then the innovation idea is a *good prospect*: it is likely to be perceived after experimentation as a profitable alternative to the current status quo. Then, the designer reacts to more intense auditing by switching to experiments that are less informative (consistent with Proposition 4) but that lead to a higher probability of approval. Figure 4-a shows the principal's utility, which is maximized for $\lambda = 0.57$. So, for good-prospect ideas, the principal engages in somewhat intense auditing and the designer restricts experimentation as, for such intense auditing, the principal is willing to rubberstamp the analyst's recommendations.

If $\frac{1-p_S^C}{p_S^C} = 1/2$, then the innovation idea is a *poor prospect*: it is unlikely that experimentation will uncover evidence showing it to be more profitable than the status quo. Again, the designer reacts to more intense auditing by experimenting less but approval probability is now non-monotonic: it increases for low values of λ—as the designer always selects

[28] In both cases, we take $\mu = 1/4$. We have $\underline{q}_H = 3/8$ if $\frac{1-p_S^C}{p_S^C} = 2$, while $\underline{q}_H = 3/4$ if $\frac{1-p_S^C}{p_S^C} = 1/2$,

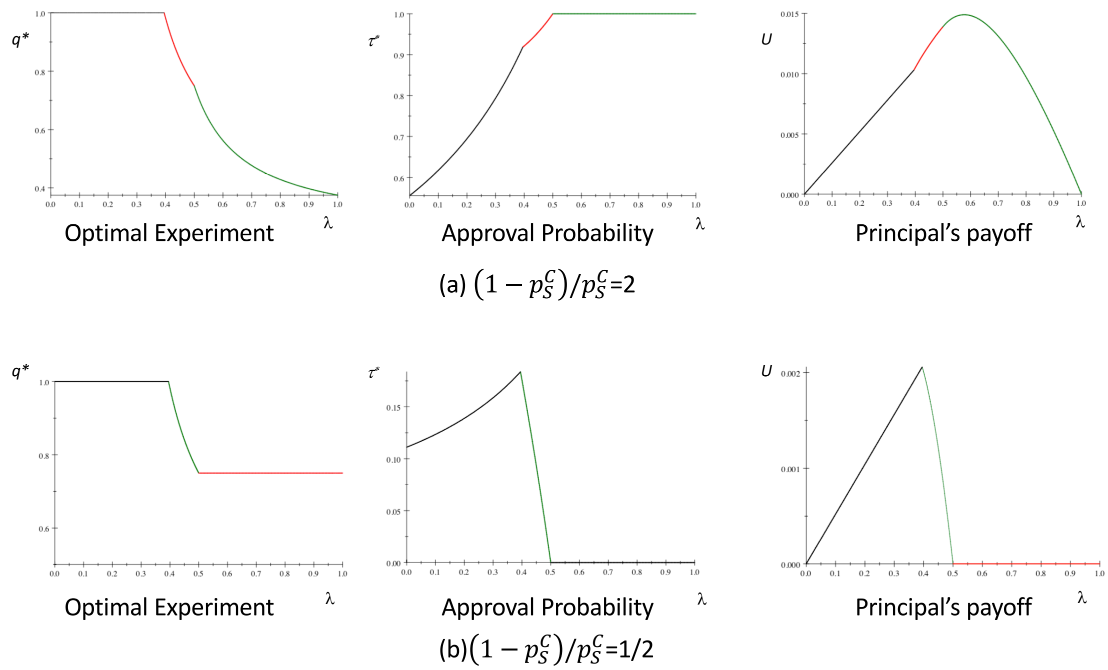

Figure 4: Equilibrium experimentation as a function of auditing for (a) a good prospect, and (b) a bad prospects.

a fully informative experiment and increased auditing simply raises approval probability—but it monotonically decreases when the designer actually switches to a less informative experiment. In fact, for $\lambda > 1/2$, the designer selects the commitment experiment so that experimentation creates no value for the principal. Figure 4-b describes the principal's utility which is maximized for $\lambda = 0.39$. So, for poor-prospect ideas, the principal seldom audits the experiment and the designer in response does not reduce experimentation—i.e., the designer's experiment fully reveals the state. Nevertheless, such lax auditing implies that approval largely relies on the principal vetting the analyst's recommendation.

5.2 Organizing for Scale

Suppose now that the firm "organizes for scale." From (15), it remains true that she would separate tasks and commit to an imperfect audit in an effort to incentivize experimentation. When organizing for scale, however, she must guard against *adverse* class switches—namely, a designer's switch from an "up-or-down" experiment to a "status-quo" experiment—in response to an imperfect audit. This terminology is motivated by the following observation: under task separation, a switch to the designer's optimal "status-quo" experiment makes the principal (weakly) worse off.

Lemma 4. *Let* $\phi_i \equiv \left(1 - p_i^C\right) / \left(p_i^C \Delta_i\right)$*,* $i = L, S$*, and suppose that* $L(c) - (\phi_i/\lambda)$ *is single-crossing in* $[0, \Delta_i]$*, and a perfect audit leads the designer to select an "up-or-down" experiment. Then, if tasks are separated, for any* $\lambda < 1$ *the principal is weakly better off if the designer is constrained to "up-or-down" experiments.*

The concern with adverse class switches underlies the main differences between organizing to innovate and organizing for scale. First, to ensure that the designer selects an "up-or-down" experiment, she may now prefer to integrate tasks. Second, she may prefer to perfectly audit, even if the designer is responsive to auditing, if lax auditing triggers an adverse class switch.

Proposition 6. *(i) Suppose that for* $\lambda < 1$*, the conditions of Lemma 4 and Proposition 3-ii hold. Then, the principal separates tasks.*
(ii) Let

$$W(\Delta, \bar{c}) \equiv \frac{\int_0^{\bar{c}} c dF(c)}{\lambda\Delta + \bar{c}\overline{F}[\bar{c}]}, \tag{18}$$

and suppose that $W(\Delta_S, (1-\lambda)\Delta_S) > W(\Delta_L, (1-\lambda)\Delta_L)$*. Then, there exist* $\underline{q}_S$ *and* μ *with* $\underline{q}_S < \mu < \underline{q}_H$ *such that the principal integrates tasks.*
(iii) Suppose that either (iii-a) $p_L^C > (\Delta_S/\Delta_L)\, p_S^C$ *and* $f(0) > 0$*, or (iii-b)* $p_L^C < (\Delta_S/\Delta_L)\, p_S^C$*, the designer is responsive to auditing and* $i^*(k, \lambda) = L$*. Then* $\lambda^* < 1$ *in every equilibrium of the organizational design game.*

As long as an adverse class switch does not occur, task separation remains optimal if the designer's preference for the less informative class are preserved when integrating tasks—this is Proposition 6-i. Integration can be optimal if the principal's auditing leads the designer under separation to select a status-quo experiment while he would instead select an "up-or-down" experiment under integration. Proposition 6-ii provides a sufficient condition for such case. Finally, Proposition 6-iii provides sufficient conditions for optimal auditing to be imperfect. First, if $p_L^C > (\Delta_S/\Delta_L)\, p_L^C$ then the designer already selects an status-quo experiment under a perfect audit, eliminating any concern that an imperfect audit might trigger a switch; then, a similar argument as in Proposition 5-ii guarantees that $\lambda^* < 1$. If $p_L^C < (\Delta_S/\Delta_L)\, p_L^C$, however, the designer selects experiment $\{\underline{q}_L, \underline{q}_H\}$ when the audit is perfect and lowering auditing intensity might induce an adverse switch. In this case, $\lambda^* < 1$ obtains as long as there is an auditing intensity that avoids a switch to a status-quo experiment.

5.2.1 Discretion, Auditing, and Experimentation.

The only reason for setting $\lambda = 1$ is that any imperfect audit that motivates valuable experimentation also leads to an adverse class switch. The same is true of integration: the principal integrates tasks only if separation would otherwise lead to the selection of a status-quo experiment. Nevertheless, the principal can always avoid adverse switches if she can commit ex-ante to ruling out the status-quo, thus committing to selecting from extreme options. Then, imperfect auditing and task separation can boost experimentation and become optimal.

Corollary 1. *Suppose that $f(0) > 0$. If the principal can ex-ante commit to ruling out decisions, then in any equilibrium we have $\lambda^* < 1$ and the principal prefers to separate tasks.*

This result resonates with insights from the delegation literature in which the principal rules out intermediate decisions to improve the informational content of delegated decision making (see, e.g., Szalay (2005) and Alonso and Matouschek (2008)). For example, Szalay (2005) shows that, to boost incentives to acquire information, a principal may rule out the (agent's) optimal decision if uninformed. In our case, agents interested in promoting a specific option (i.e., scaling up) are more willing to compromise on the status quo than the principal. Then, limiting discretion boosts experimentation by eliminating such compromise. Note also that restricting the principal's choice is accompanied by less intense monitoring: i.e., she sets $\lambda^* < 1$ if she is able to rule-out decisions. Thus, in this context, discretion and monitoring act as complements.

5.3 Organizational Design—Optimal Distribution of Tampering Costs.

So far, we have considered organizations that adjust their structure in response to the tampering costs of their agents. In some instances, however, these costs can result from organizational practices: the principal can establish both the cost to the agent from each tampering action–say, through by-laws exacting punishments upon detection of tampering–and its likelihood–for example, through improved data security systems, audit trails, and database access controls.

Suppose that the distribution of tampering costs can be fully specified by the principal. Consistent with our theme of "credible skepticism" to motivate experimentation, she will incentivize some tampering in equilibrium by making low tampering costs sufficiently likely.

Proposition 7. *Suppose that either the principal organizes for innovation, or she organizes for scale and* $p_L^C < (\Delta_S/\Delta_L)\, p_S^C$. *If she can specify the task allocation, auditing intensity, and the distribution of tampering costs, then:*
(i) She sets a prior-independent auditing intensity

$$\lambda^*_{opt} = \frac{1}{2 - \underline{q}_H}. \tag{19}$$

(ii) The designer selects a fully informative experiment $\pi = \{0, 1\}$.
(iii) There is a multiplicity of optimal cost distributions but, among them, the following minimizes expected tampering costs,

$$F^*_{opt}(c) = \begin{cases} \frac{\mu\left(1-\underline{q}_H\right)}{\underline{q}_H(1-\mu)} \text{ for } c \in [0, \frac{1-\underline{q}_H}{2-\underline{q}_H}), \\ 1 \text{ for } c \geq \frac{1-\underline{q}_H}{2-\underline{q}_H}. \end{cases} \tag{20}$$

For this cost distribution, the principal is indifferent between separating or integrating tasks.
(iv) The principal can implement (20) through a dual internal-external audit: Tampering is always costless, but an internal audit privately verifies the agent's report with probability $\frac{\underline{q}_H-\mu}{\underline{q}_H(1-\mu)}$ *and rectifies a tampered report.*

An important principle in organizing data analytics is that, under delegated experimentation, the organization must also allow, to some extent, tampering by agents. To do so optimally, the organization both raises the likelihood of low tampering costs and engages in

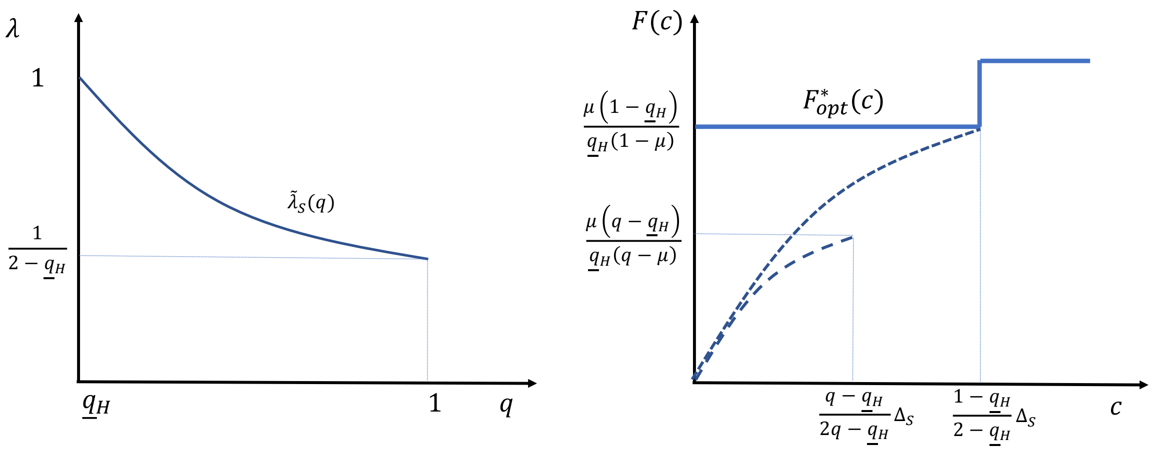

(a) Maximum λ that induces {0,q}.

(b) Minimum cost distribution to induce {0,q} when $\lambda = \tilde{\lambda}_S(q)$.

Figure 5: Optimal auditing and tampering cost distribution.

lax auditing: the optimal auditing intensity (19) is always lower than 1, but higher than 1/2, and increases with the principal's approval threshold.

We prove this proposition by first solving an auxiliary problem: to find the maximum auditing intensity that induces the selection of experiment $\pi = \{0, q\}$ for some cost distribution. The solution is $\widetilde{\lambda}_S(q) = q/(2q - \underline{q}_H)$–depicted in Figure 5–which is obtained by ensuring that switching to the commitment experiment $\{0, \underline{q}_H\}$ is never profitable for the designer. As shown in the proof of the proposition, there are many different distributions that would lead him to select $\{0, q\}$ when auditing intensity is $\widetilde{\lambda}_S(q)$. In all of them, tampering for low realizations must be sufficiently likely so that the principal can commit to high approval rates only if experiments are sufficiently informative. Formally, experiment $\{0, q\}$ is the designer's incentive compatible choice if the distribution $F(c)$ exceeds some lower bound–see Figure 5-b where the minimum values of $F(c)$ for two experiments are represented by dashed lines.

Optimizing over q gives (19) and we obtain $q^* = 1$–this is Proposition 7-ii. To wit, under an optimal organization, the designer has no incentive to garble an experiment that reveals the underlying state and the principal rubberstamps any scale-up recommendation after an inconclusive audit. If the organization wants to minimize the costs imposed upon agents—say because of concerns with increasing hiring costs—then the optimal distribution makes tampering either costless or completely deters tampering, with $\Pr[c=0] = \frac{\mu\left(1-\underline{q}_H\right)}{\underline{q}_H(1-\mu)}$– see Figure 5. This distribution also makes task allocation irrelevant, as expected tampering costs are always zero.

The optimal organization that satisfies (20) can be afforded an intuitive implementation: the analyst faces no cost of tampering but his report is subjected to a *decoupled internal-external audit.* First, the report is internally audited, albeit the probability of a elucidating the true outcome is restricted to $\frac{\underline{q}_H-\mu}{\underline{q}_H(1-\mu)}$. If the internal audit is conclusive, however, it ensures that the report is consistent with the experimental outcome. Second, this report is subjected to an imperfect external audit, which is conclusive with probability $1/(2-\underline{q}_H)$. In summary, the organization conducts more intense internal audits for bad prospects—i.e., when scaling-up is less likely under a perfect audit—but commits to a prior-independent external audit.

Importantly, the outcome of the internal audit must be unknown to the principal. This decoupling of audits is essential to incentivize experimentation: if the outcome of the internal audit were known to the principal, the designer in anticipation would then select the commitment experiment $\left\{0, \underline{q}_H\right\}$. The accounting literature is also concerned with the possible effects of internal control audits and, in particular, whether the public disclosure of internal control audits should be mandatory. For example, Lennox and Wu (forthcoming) study the effects of regulation mandating the disclosure of internal control audits in China. They present evidence that mandatory disclosure of internal control audits can significantly reduce the quality of information.

If $p_L^C \geq (\Delta_S/\Delta_L)\, p_S^C$, so that the designer selects a status-quo experiment when $\lambda = 1$, then to optimally induce an "up-or-down" experiment the principal integrates tasks and the optimal cost distribution must make tampering strictly costly. The reason is given by Proposition 4-ii(a): regardless of the cost distribution and auditing intensity, if tasks are separated, the designer always selects a status-quo experiment if he prefers the status-quo experiment when $\lambda = 1$. The principal must then incentivize a class switch by integrating tasks and equilibrium tampering must be strictly costly. Nevertheless, and consistent with Corollary 1, if she can commit to ruling out decisions–so that she can induce an "up-or-down experiments" by ruling out decision d_S– then the optimal organization would always lead to full experimentation and rely on an internal-external dual audit.

Corollary 2. *Suppose that the principal can ex-ante commit to ruling out decisions. Then, the optimal organization satisfies (19) and (20) in Proposition 7.*

6 Extensions

We now consider two extensions to our main analysis: tampering costs are incurred only if the analyst is caught tampering, and we allow for equilibria involving non-robust experimentation.

6.1 Tampering costs incurred only if audit is conclusive.

A conclusive audit reveals both the true experimental outcome and whether tampering took place. Suppose then that tampering costs are only borne if the analyst is found to have tampered. We can adjust our analysis by positing that setting the expected tampering cost to λc when the cost realization is c. A full analysis of this case can be found in the online Appendix B.

For a fixed $\lambda < 1$, Proposition 1 still applies, albeit with different threshold values $\bar{c}(q)$– this is also true of robust experiments in Proposition 2. To see how this affects the designer's payoff, consider $\{\underline{q}_i, q\} \in \Pi_i$ that induces tampering threshold $\bar{c}$. Then, expressions (3-4) translate to

$$\bar{F}\left(\bar{c}\right) p = p_i^C \text{ with } p = \frac{q-\mu}{q-\underline{q}_i},$$
$$\bar{c} = \frac{1-\lambda}{\lambda}\left(\Delta_i - \left(1-\tau\right)\Delta_S\right).$$

The first condition is identical to (4): the tampering threshold determines the probability of truthful reporting given an unfavorable outcome–i.e., $\bar{F}\left(\bar{c}\right)$– so that the probability of observing message $m = \underline{q}_i$ must be equal to p_i^C, irrespective of whether expected tampering costs are c or λc. However, the second condition shows that the corresponding scale-up probability must be lower to account for the lower tampering cost.

Nevertheless, the comparative statics of experimentation with respect to the two organizational levers remain the same: separating tasks or decreasing auditing intensity always increases within-class experimentation but may lead to an adverse class switch.[29] What are the organizational implications if the analyst's tampering cost reduces to λc? Restricting attention to Sections 5.1 and 5.3, it is still true that, when organizing to innovate, the principal prefers to separate tasks and to commit to an imperfect audit whenever the designer is responsive to auditing–however, the conditions for designer responsiveness are now more stringent. Moreover, if the principal can freely shape the distribution of tampering costs and can commit to ruling out decisions, then the same organizational design as in Proposition 7 remains optimal-see online Appendix B.

6.2 Organizing for scale with non-robust experimentation.

Our organizational-design analysis has focused on equilibria in which the principal resolves any indifference after a conclusive audit by selecting the agents' preferred decision. While this analysis is without loss when the principal organizes to innovate, these equilibria provide only a lower bound on the designer's equilibrium payoff when organizing for scale–see Proposition 2. Thus, one wonders if our insights may change if the principal instead conditions her audited decision on whether the analyst tampered. We show that, while equilibrium auditing intensity may vary, our main qualitative results hold when we allow for "non-robust" equilibria.

Proposition 2-ii shows that if the designer's equilibrium payoff exceeds that of a robust experiment, then $\underline{q}_S \in S(\pi^*)$, i.e., the designer's experiment is (possibly a mixture involving) a experiment $\{\underline{q}_S, q\}$. The following lemma shows that the designer's maximum payoff from a experiment $\{\underline{q}_S, q\}$ comes when the principal punishes, to some extent, tampering.

Lemma 5. *Suppose that the principal organizes for scale and the designer selects $\{\underline{q}_S, q\}, q > \underline{q}_H$. Let $\overline{c}_S$ be the tampering threshold for $\{\underline{q}_S, q\} \in \Pi_S$; i.e., $\overline{c}_S$ satisfies $\bar{F}\left(\overline{c}_S\right) = p_S^C\left(\frac{q-\underline{q}_S}{q-\mu}\right)$– see (3). If $\tau_I(q)$ denotes the probability that the principal selects d_S after the analyst is found tampering, then at the equilibrium which maximizes the designer's payoff we have*

[29]The only notable difference is that the designer's payoff is no longer supermodular in (minus) the tampering threshold and the auditing intensity. Moreover, the fact that tampering is now less costly implies that the conditions for designer's responsiveness are more stringent.

$\tau_I(q) = \tau_I^*(q)$ *with*

$$\tau_I^*(q) \equiv \min\left\{\tau_I(q) : \bar{c}_S = (1-\lambda)\tau_U(q)\Delta_S - \lambda(1-\tau_I(q))(\Delta_L - \Delta_S), \tau_U(q) \in [0,1]\right\}.$$

In particular, if the principal rubberstamps a scale-up recommendation for $\{\underline{q}_S, q\} \in \Pi_S$*, then* $\tau_I^*(q) = 1$.

If $\tau_I^*(q) < 1$ then the principal punishes tampering by lowering the probability of selecting d_S if the analyst tampered. The designer obtains then a higher payoff from $\{\underline{q}_S, q\}$ not because it reduces the tampering threshold–indeed, the tampering threshold $\bar{c}_S$ is the same as that of the robust status-quo experiment–but rather because it allows the principal to scale-up more often after an unaudited scale-up recommendation. In turn, the principal could use the designer's preference for equilibria with $\tau_I(q) = \tau_I^*(q)$ as follows: she punishes tampering (i.e. sets $\tau_I(q) = \tau_I^*(q)$) if the designer's experiment is sufficiently informative–i.e., if the designer selects $\{\underline{q}_S, q\}$–while she is lenient if the designer selects a less informative status-quo experiment (i.e., $\tau_I(q') = 1$, if $q' < q$). While moving to non-robust experiments may lead the principal to choose a different auditing intensity, our main qualitative results of Section 5 nevertheless hold for non-robust experimentation.

Proposition 8. *Suppose that we restrict attention to PBE of the design subgame that maximize the principal's expected payoff. Then: (i) relative to Proposition 6, the principal is now more likely to set* $\lambda^* < 1$ *and to separate tasks, and (ii) the auditing intensity and cost distribution in Proposition 7 remain optimal.*

The only reason for the principal to set $\lambda^* = 1$ when organizing for scale is to avoid an adverse class switch. However, the principal can induce a (weakly) more informative non-robust status-quo experiments for the same λ, so that the previous class switch may now be profitable for the principal. This explains Proposition 8-i. Moreover, the principal can increase the designer's payoff from experiment $\{\underline{q}_S, q\}$ only if she does not rubberstamp a scale-up recommendation for a robust experiment–i.e., only if $\tau_U(q) < 1$ for $\tau_I(q) = 1$. Nevertheless, the optimal organization in Proposition 7 is based on finding cost distributions that implement $\{\underline{q}_S, q\}$ with $\tau_U(q) = 1$, thus the principal cannot increase the designer's payoff for such experiments by punishing tampering. That is, if the principal could shape the distribution of tampering costs and select her preferred PBE for each subgame, then the auditing intensity (19) and cost distribution (20) remain optimal.

7 Discussion and Concluding Remarks

In this paper, we develop a model of data analytics and argue that organizations that delegate experimentation to their agents must also create a culture of "credible skepticism" by limiting decision-makers' ability to assess the truthfulness of the information they receive. We now discuss these findings in the context of several strands of the literature, after which we conclude.

7.1 Related Literature

Literature on decision-making processes in organizations:

Our analysis contributes to the study of decision making processes in organizations and, in particular, to how organizations optimally react to the incentive conflicts that members face (see Gibbons, Matouschek, and Roberts (2013) and Bolton and Dewatripont (2013) for excellent surveys of this literature). For instance, in models of strategic delegation, the organization would like to assign authority to a party whose preferences may differ from those of the organization as these affects the production and communication of information (for instance, in Dessein, 2002, delegation to a biased intermediary can improve cheap-talk communication with experts).[30] One recent example is Nayeem (2017), who quantifies the value of appointing a decision maker that is harder to convince to approve a project–e.g., as his preference for a "good project" are weaker than those of the organization. That is, there is value in appointing a "skeptic" for project approval. In our model, however, the principal cannot credibly delegate the decision to someone else nor commit to biasing decisions in favor of agents. Skepticism arises not because of differing preferences, but as an attitude to (rationally) doubt the claims made by others.

Our paper also contributes to the literature that studies how "light monitoring" of agents' recommendations may avoid crowding-out their efforts to experiment (see, e.g., Aghion and Tirole (1997)). In our case, imperfect auditing allows the principal to refrain from adopting the agent's self-serving recommendation, thus, spurring experimentation.

The literature on task allocation has emphasized that task separation can allow for the

[30] More generally, decision makers may be able to commit to ex-post biasing decisions in favor of experts, e.g., in a relational setting as in Alonso and Matouschek, 2007.

provision of higher power incentives in each task (Holmstrom and Milgrom, 1991, Dewatripont, Jewitt, and Tirole, 2000) or improve information acquisition (Dewatripont and Tirole, 1999). Moreover, in sequential tasks, task separation may increase the information generated in the first task to incentivize the second (Lewis and Sappington, 1997, Landier, Sraer, and Thesmar (2009)), or can be optimal under effort externalities between tasks (Schmitz, 2013). We also find that task separation allows for stronger incentives to experiment, even though we do not allow for explicit incentives, as separation provides a "coarse" instrument to lower the costs of experimentation.

Literature on Information acquisition and Communication

We contribute to the literature that studies models of delegated expertise (Demski and Sappington (1987))–in particular, models in which a decision maker relies on the information actively gathered and communicated by experts. For instance, Pei (2015), Argenziano, Severinov, and Squintani (2016), and Deimen and Szalay (2019) consider models where an agent decides what information to gather if communication with the principal takes the form of cheap talk, while Che and Kartik (2009) considers certifiable disclosure.[31] Argenziano et al. (2016) and Deimen and Szalay (2019) use the threat of off-path "bad" communication (e.g., a reversion to a "babbling" equilibrium) if the expert acquires less information to motivate information acquisition. In Pei (2015) communication is "frictionless:" the agent reveals all the information gathered if acquiring a less informative signal is always feasible (and less costly) (see also Gentzkow and Kamenica (2016)). In Che and Kartik (2009), incentives to acquire information come from players having different priors: an expert has a stronger incentive to be informed relative to the common prior case as he expects that better information will lead the principal to, on average, embrace his point of view.[32]

A main insight in these papers is that frictions in communication can be used to discipline agents if they underinvest in information acquisition.[33] While this insight resonates with our

[31] Our communication stage is also related to models of communication with lying costs–e.g., Kartik, Ottaviani, and Squintani (2007) and Kartik (2009). Relative to these models, our communication model is simpler, as we consider a message independent tampering cost, but we incorporate an information acquisition stage prior to communication.

[32] Alonso and Câmara, 2016 also show that differences of opinion generically give rise to incentives to persuade a principal.

[33] Frictions in communication can also enhance the amount of information transmitted, see e.g., Blume et

main finding, our mechanism is markedly different. In contrast to Pei (2015), Argenziano et al. (2016), and Che and Kartik (2009), the agent faces no explicit cost in acquiring more information in our model–this matches our main application where data becomes available to the organization automatically through its normal operation. In contrast to Deimen and Szalay (2019), we consider an explicit cost of misrepresentation when the analyst communicates the results, as well as the principal's ability to audit the analyst's message and to allocate tasks to different agents.

Literature on relaxing the commitment assumption in models of Bayesian persuasion.

Our paper contributes to the literature that relaxes the sender-commitment assumption in Kamenica and Gentzkow's model of Bayesian Persuasion.[34] Papers in this recent literature differ on the modeling of imperfect commitment. For instance, Guo and Shmaya (Forthcoming) consider a model of costly miscalabration: the sender decides the statistical properties of an experiment and can deviate from the "asserted" meaning for each outcome at a cost related to the difference between the asserted meaning and its true meaning. That is, they allow for a sender's private experimental design rather than our public experimental design subject to private output-tampering. Min (2020) considers the output-tampering case but tampering only occurs with some exogenous probability and explores the effect of changes in this probability in Crawford and Sobel (1982) uniform-quadratic case. In these papers, there is no tampering or misrepresentation in equilibrium.[35] Instead, in our paper tampering is a generic equilibrium phenomenon resulting from the principal's choice of auditing intensity. Perez-Richet and Skreta (2021) study test design under costly state falsification: a designer selects a test and an agent can change its input at a cost. That is, in contrast to our setup with output-tampering, the agent engages in input-tampering. Fréchette, Lizzeri, and Perego (2019) analyze experiments in which the level of commitment can vary across treatments,

al (2007)

[34]See also the literature on strategic sample selection, e.g., Tillio, Ottaviani, and Sørensen (2017), Tillio, Ottaviani, and Sørensen (2021), Adda, Decker, and Ottaviani (2020), Felgenhauer and Loerke (2017) and Libgober (Forthcoming).

[35]Tampering-proof equilibria are the focus of Min (2020), while Guo and Shmaya (Forthcoming) show that there is always a Sender-optimal equilibrium with a calibrated strategy–i.e., such that receiver correctly anticipates its meaning. See also Sobel (2020) for an analysis that distinguishes between "lying" and "deception".

albeit the ability to tamper is exogenously given, while it is an equilibrium outcome in our paper.

Closest to our modeling of limited commitment are Lipnowski, Ravid, and Shishkin (2018) and Nguyen and Tan (2018). Lipnowski et al. (2018) consider an information design setup with output-tampering where the tampering probability depends on the actual message/state and provide an elegant geometric characterization of the sender's value of persuasion. While in our setup the tampering probability is also message/state dependent, it arises endogenously from the agent's equilibrium incentives to tamper. Furthermore, while they characterize the optimal level of credibility from the sender's perspective, our focus lies on the receiver/principal's perspective. Nguyen and Tan (2018) also study public experimentation subject to private output-tampering. They consider a setup with a fixed experimental outcome space and message space and a communication technology where each message carries a cost that depends both on the message and the experimental outcome. They focus on conditions on this technology for the Sender's preferred equilibrium to be supported without tampering (Condition 1 in Nguyen and Tan (2018)). Out setup does not satisfy Condition 1 (as the tampering cost is the same regardless of the message sent) and, thus, we cannot apply their results.

One overarching theoretical difference with this literature is that we endogenize the sender's commitment power by allowing the receiver to select among different organizational practices; for instance, how much to audit of the sender's message. Thus, while the literature shows that exogenously relaxing the sender's commitment can be beneficial for the receiver, we show the extent to which imperfect commitment is an equilibrium outcome of the receiver's organizational practices.

7.2 Concluding Remarks

The ICT revolution–by lowering the costs of data acquisition, storage and processing–has made managers more reliant on the insights derived from analyzing these data rather than the intuitions and opinions of other members of the organization. It would then seem that many of the trade-offs that drive the optimal organization to process information are no longer relevant. We argue that unresolved conflict still makes organizational structure meaningful as members handling data still decide which data to use and how to analyze it. We show that this poses a fundamental trade-off: dissuading misrepresentation also reduces data utilization, limiting the insights that agents derive from the data. Optimal organizations are then based on a culture of "credible skepticism:" managers have limited ability to audit the data and analytics behind the recommendations issued by agents, which invites tampering and misrepresentation in equilibrium.

The adoption of new technologies such as blockchain can eliminate tampering by effectively imposing an infinitely high tampering cost (Tapscott and Tapscott, 2017). Nevertheless, under delegated experimentation, this is never optimal for the firm as the optimal distribution of tampering costs must lead to some tampering in equilibrium. We showed that this optimal organization can be implemented through a decoupled internal-external audit: tampering is costless, but an internal (imperfect) audit can limit its effect by rectifying the tampered outcome with the true outcome. Then, an external audit is triggered with some probability without knowing whether the internal audit rectified the report. This system of consecutive audits strikes a perfect balance between experimentation and tampering and minimizes the tampering costs of agents. Importantly, under an optimal internal-external audit, the data architect engages in full experimentation.

To focus on the trade-off between experimentation and misrepresentation, we offer a streamlined model. In particular, managers do not have access to alternative sources of information (i.e., they do not "seek a second opinion") nor do they induce competition between agents to persuade them. We see these extensions as promising and leave them for future work.

References

Adda, J., Decker, C., & Ottaviani, M. (2020). P-hacking in clinical trials and how incentives shape the distribution of results across phases. *Proceedings of the National Academy of Sciences*, *117*(24), 13386–13392.

Aghion, P., & Tirole, J. (1997). Formal and real authority in organizations. *Journal of Political Economy*, *105*(1), 1–29.

Alonso, R., & Câmara, O. (2016). Bayesian Persuasion with Heterogeneous Priors. *Journal of Economic Theory*, *165*, 672–706.

Alonso, R., & Matouschek, N. (2007). Relational delegation. *The RAND Journal of Economics*, *38*(4), 1070–1089.

Alonso, R., & Matouschek, N. (2008). Optimal delegation. *The Review of Economic Studies*, *75*(1), 259–293.

Argenziano, R., Severinov, S., & Squintani, F. (2016). Strategic information acquisition and transmission. *American Economic Journal: Microeconomics*, *8*(3), 119–155.

Bolton, P., & Dewatripont, M. (2013). Authority in organizations a survey. In R. Gibbons & J. Roberts (Eds.), *The handbook of organizational economics* (pp. 342–372).

Brynjolfsson, E., Hitt, L. M., & Kim, H. H. (2011). Strength in numbers: How does data-driven decisionmaking affect firm performance? *Available at SSRN https://ssrn.com/abstract=18194*

Brynjolfsson, E., & McElheran, K. (2016). The rapid adoption of data-driven decision-making. *American Economic Review*, *106*(5), 133–39.

Che, Y.-K., & Kartik, N. (2009). Opinions as incentives. *Journal of Political Economy*, *117*(5), 815–860.

Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica*, *50*(6), 1431.

Deimen, I., & Szalay, D. (2019). Delegated expertise, authority, and communication. *American Economic Review*, *109*(4), 1349–1374.

Demski, J. S., & Sappington, D. E. M. (1987). Delegated expertise. *Journal of Accounting Research*, *25*(1), 68.

Dessein, W. (2002). Authority and communication in organizations. *Review of Economic Studies*, *69*(4), 811–838.

Dewatripont, M., Jewitt, I., & Tirole, J. (2000). Multitask agency problems: Focus and task clustering. *European Economic Review*, *44*(4-6), 869–877.

Dewatripont, M., & Tirole, J. (1999). Advocates. *Journal of Political Economy*, *107*(1), 1–39.

Felgenhauer, M., & Loerke, P. (2017). Bayesian persuasion with private experimentation. *International Economic Review*, *58*(3), 829–856.

Fréchette, G., Lizzeri, A., & Perego, J. (2019). *Rules and commitment in communication: An experimental analysis*. Working Paper.

Gentzkow, M., & Kamenica, E. (2016). Disclosure of endogenous information. *Economic Theory Bulletin*, *5*(1), 47–56.

Gibbons, R., Matouschek, N., & Roberts, J. (2013). Decisions in organizations. In R. Gibbons & J. Roberts (Eds.), *The handbook of organizational economics* (pp. 373–431).

Guo, Y., & Shmaya, E. (Forthcoming). Costly miscalibration. *Theoretical Economics*.

Guo, Y., & Shmaya, E. (2019). The Interval Structure of Optimal Disclosure. *Econometrica*, *87*(2), 653–675.

Hand, D. J. (2020). *Dark data*. Princeton Univers. Press.

Holmstrom, B., & Milgrom, P. (1991). Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design. *The Journal of Law, Economics, and Organization*, *7*(special_issue), 24–52.

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

Kartik, N. (2009). Strategic communication with lying costs. *Review of Economic Studies*, *76*(4), 1359–1395.

Kartik, N., Ottaviani, M., & Squintani, F. (2007). Credulity, lies, and costly talk. *Journal of Economic Theory*, *134*(1), 93–116.

Kolotilin, A. (2018). Optimal Information Disclosure: A Linear Programming Approach. *Theoretical Economics*, *13*(2), 607–635.

Kolotilin, A., Mylovanov, T., Zapechelnyuk, A., & Li, M. (2017). Persuasion of a Privately Informed Receiver. *Econometrica*, *85*(6), 1949–1964.

KPMG. (2016). *Building trust in analytics*. KPMG International Data Analytics.

KPMG. (2018). *Guardians of trust*. KPMG International Data Analytics.

Landier, A., Sraer, D., & Thesmar, D. (2009). Optimal dissent in organizations. *Review of Economic Studies*, *76*(2), 761–794.

Lewis, T., & Sappington, D. M. (1997). Information management in incentive problems. *Journal of Political Economy*, *105*(4), 796–821.

Libgober, J. (Forthcoming). False positives and transparency. *American Economic Journal: Microeconomics*.

Lipnowski, E., Ravid, D., & Shishkin, D. (2018). Persuasion via weak institutions. *SSRN Electronic Journal 3168103*.

McKinsey, & Co. (2018). *Building an effective analytics organization*. McKinsey on Payments.

Michelle R. Smith, C. L., & Amy, J. (2020). Numerous states accused of manipulating or bungling covid-19 data. *Chicago Tribune*.

Min, D. (2020). Bayesian persuasion under partial commitment. *Unpublished*.

Nayeem, O. A. (2017). Bend them but don't break them: Passionate workers, skeptical managers, and decision making in organizations. *American Economic Journal: Microeconomics*, *9*(3), 100–125.

Nguyen, A., & Tan, T. Y. (2018). Bayesian persuasion with costly messages. *SSRN Electronic Journal 3298275*.

Pei, H. D. (2015). Communication with endogenous information acquisition. *Journal of Economic Theory*, *160*, 132–149.

Perez-Richet, E., & Skreta, V. (2021). Test design under falsification. *CEPR Discussion Papers 15627*.

Phillips, D. (2020). Brazil stops releasing covid-19 death toll and wipes data from official site. *The Guardian*.

Robbins, C., & Reuben, A. (2020). Coronavirus: Why are international comparisons difficult? *BBC News*.

Schmitz, P. W. (2013). Job design with conflicting tasks reconsidered. *European Economic Review*, *57*, 108–117.

Sobel, J. (2020). Lying and deception in games. *Journal of Political Economy*, *128*(3), 907–947.

Szalay, D. (2005). The economics of clear advice and extreme options. *The Review of Economic Studies*, *72*(4), 1173–1198.

Tapscott, D., & Tapscott, A. (2017). How blockchain will change organizations. *MIT Sloan Management Review*, *58*(2), 10.

Tillio, A. D., Ottaviani, M., & Sørensen, P. N. (2017). Persuasion bias in science: Can economics help? *The Economic Journal*, *127*(605), F266–F304.

Tillio, A. D., Ottaviani, M., & Sørensen, P. N. (2021). Strategic sample selection. *Econometrica*, *89*(2), 911–953.

Wu, L., Hitt, L., & Lou, B. (2020). Data analytics, innovation, and firm productivity. *Management Science*, *66*(5), 2017–2039.

A Appendix

Proof of Proposition 1: Given $\pi = \{q, \Pr[q]\}_{i\in S(\pi)}$, with type space $S(\pi)$ consider a PBE of the communication subgame where the analyst's reporting strategy is $m^*(q,c)$, which leads to decisions $d_I^*(m,s)$ and $d_U^*(m)$. Proposition 1-i follows immediately as the gain from tampering is the same for all analysts that observe the same experimental outcome: if type (q,c) finds it profitable to send $q_z \neq q$ instead of q, all types with $c' < c$ will strictly prefer to tamper.

For part (ii), consider the set of tampered outcomes M_T defined in the proposition. Suppose that $q, q' \in M_T$ but the distributions $d_U^*(q)$ and $d_U^*(q')$ lead to different expected payoffs for the analyst.[36] If $\underline{q}_i \notin S(\pi)$, $i \in \{S,H\}$, then the principal never mixes after a conclusive audit and the analyst's payoff in this event is independent of the message sent. This is also the case if the audited decision $d_I^*(m, \underline{q}_i)$ is independent of m whenever $\underline{q}_i \in S(\pi), i \in \{S,H\}$. This implies that the analyst only benefits from tampering in the event that the audit is inconclusive, but if $d_U^*(q)$ and $d_U^*(q')$ yield a different payoff, then $m^*(q,c)$ cannot be part of an equilibrium. Therefore, we must have that $d_U^*(q) = d_U^*(q')$ for $q, q' \in M_T$. Finally, suppose that $q \in M_T$, $q \neq \underline{q}_i, i \in \{S,H\}$. Then, the sender never gains

[36]As the principal only mixes after an inconclusive audit when her posterior is either $\underline{q}_S$ (thus, mixing between d_L and d_S) or $\underline{q}_H$ (thus, mixing between d_S and d_H) the analyst must obtain a different expected payoff after an inconclusive audit when reporting q and q' if these distributions are different.

from tampering, as the audited decision is independent of m and the unaudited decision would be the same if he had instead truthfully reported his type. ■

Proof of Proposition 2: (i) We show that for each λ and $k \in \{\mathcal{S}, \mathcal{I}\}$, there is always an equilibrium of the design subgame in which the designer selects a robust experiment. To this end, let Π^{Σ} be the set of experiments with outcomes $\left\{\underline{q}_L, \underline{q}_S, q\right\}$, indexed by (p_L, p_S, p) with $p_L, p_S, p \geq 0$ and $p_L + p_S = 1$,and defined as follows:

$$\Pi^{\Sigma} \equiv \left\{\left\{\underline{q}_L, \underline{q}_S, q\right\} : q \geq \underline{q}_H \textit{ and} \Pr[q] = 1 - p, \Pr\left[\underline{q}_L\right] = p * p_L, \Pr\left[\underline{q}_S\right] = p * p_S\right\},$$

and equilibrium decision making

$$d_U\left(\underline{q}_i\right) = d_I\left(m, \underline{q}_i\right) = d_i; d_I(m, q) = d_H; d_U(q) = \tau d_H + (1 - \tau) d_S,$$

and note that $\Pi \subset \Pi^{\Sigma}$.

Now consider an arbitrary finite experiment $\tilde{\pi} = \{q, \Pr[q]\}_{q \in S(\tilde{\pi})}$ and suppose that players follow a PBE of the communication subgame in which the principal, if indifferent after a conclusive audit–which only applies if either $\underline{q}_S \in S(\tilde{\pi})$ or $\underline{q}_H \in S(\tilde{\pi})$–always selects the agents' preferred decision–i.e., $d_I^*(m, \underline{q}_i) = d_i$ $i = \{S, H\}$. We show that there exists $\tilde{\pi}^R \in \Pi$ that (weakly) improves the designer's payoff relative to $\tilde{\pi}$. Therefore, if $\pi^* \in \Pi$ maximizes the designer's payoff when restricted to Π, then selecting π^* is a PBE of the design subgame, as the designer's expected utility cannot be improved by any alternative $\tilde{\pi}$.

We proceed in two steps. In step 1 we derive an experiment $\tilde{\pi}^{\Sigma} \in \Pi^{\Sigma}$ that improves the designer's payoff relative to $\tilde{\pi}$. In step 2, we show that the designer's maximum expected payoff in Π^{Σ} is achieved by an experiment in Π.

Step 1: Define $S_T(\tilde{\pi})$ as the set of tampering types:

$$S_T(\tilde{\pi}) = \{q \in S(\tilde{\pi}) : \Pr[m^*(q, c) \neq q] > 0\},$$

and recall that, from Proposition 1, $M_T(\tilde{\pi})$ is the set of tampered outcomes. Thus, type $q \in S_T(\tilde{\pi})$ will tamper with positive probability while some tampering type will report $q' \in M_T(\tilde{\pi})$ with positive probability. Proposition 1 shows that, if $d_I^*(m, \underline{q}_i) = d_i$ whenever $\underline{q}_i \in S(\pi), i = \{S, H\}$, then $S_T(\tilde{\pi}) \cap M_T(\tilde{\pi}) = \emptyset$. We first show that tampering types correspond to low realizations while tampered outcomes are associated with high realizations of the experiment, i.e.,

$$q_{S_T} \equiv \max\{q : q \in S_T(\tilde{\pi})\} < \min\{q : q \in M_T(\tilde{\pi})\} \equiv q_{M_T}. \tag{21}$$

To see this, let d'_U be the decision following an unaudited tampered outcome—see Proposition 1-ii(a)—and suppose, by contradiction, that there are $q' < q''$ with $q' \in M_T(\tilde{\pi})$ and $q'' \in S_T(\tilde{\pi})$. Assumption 1 implies that message $m = q''$ is sent with positive probability and, as $q'' \notin M_T(\tilde{\pi})$, we must have that the posterior belief of the principal if the audit is inconclusive must be q''. Since $q' < q'' \leq q_{S_T}$, and Proposition 1-ii(b) shows that type $s = q' \in M_T(\tilde{\pi})$ sends $m = q'$, the principal's posterior belief after an unaudited $m = q'$ must be strictly lower than q_{S_T}. But then, we must have $q_{S_T} \notin S_T(\tilde{\pi})$ as type q_{S_T} prefers to induce decision $d_U(q_{S_T})$ rather than tamper to induce d'_U.

Next, partition $S(\tilde{\pi})$ by defining $X_L(\tilde{\pi}) = S(\tilde{\pi}) \cap (\underline{q}_L, \underline{q}_S)$, $X_S(\tilde{\pi}) = S(\tilde{\pi}) \cap \left(\underline{q}_S, \underline{q}_H\right)$ and $X_H(\tilde{\pi}) = S(\tilde{\pi}) \cap [\underline{q}_H, 1]$. We now show that (21) implies that all messages in $X_i(\tilde{\pi})$ lead to the same unaudited (mixture over) decision(s)–which means that all types in $X_i(\tilde{\pi})$ face the same gain from tampering and must therefore have the same tampering threshold. Proposition 1-ii(a) implies that this is true if $X_i(\tilde{\pi}) \subset M_T(\tilde{\pi})$. We will show now by contradiction that there cannot be tampered outcomes as well as non-tampered outcomes in $X_i(\tilde{\pi})$, $i \in \{L, S\}$. To see this, suppose that q_{M_T} defined in (21) satisfies $q_{M_T} \in X_i(\tilde{\pi})$ and there is some $q' \in X_i(\tilde{\pi})$ but $q' \notin M_T(\tilde{\pi})$. Then we must have $q' < q_{M_T}$, but $d^*_U(q') = d_i$ as the posterior after an unaudited message q' is precisely q'. However, the posterior after unaudited $q_{M_T} \in M_T(\tilde{\pi})$ must be strictly lower than q_{M_T}. But then we must have that $d^*_U(q_{M_T}) = d_i$, otherwise tampering types would send message q' instead of q_{M_T}. Thus, for all $q, q' \in X_i(\tilde{\pi})$, $d^*_U(q) = d^*_U(q')$.

We now construct $\tilde{\pi}_c$ that has an equilibrium that gives the designer the same expected utility as the equilibrium of $\tilde{\pi}$. We do so by replacing all realizations in $X_i(\tilde{\pi})$, $i = L, S, H$, with a realization $s = \tilde{q}^{X_i}$ that is its conditional expectation, i.e.,

$$\tilde{q}^{X_i} = \frac{\sum_{q \in X_i(\tilde{\pi})} \Pr[q]\, q}{\sum_{q \in X_i(\tilde{\pi})} \Pr[q]}, \;\; \Pr\left[\tilde{q}^{X_i}\right] = \sum_{q \in X_i(\tilde{\pi})} \Pr[q],$$

and adjusting the equilibrium (mixture over) messages to

$$m_c(\tilde{q}^{X_i}, c) = \frac{\sum_{q \in X_i(\tilde{\pi})} \Pr[q] \sum_{j=\{L,S,H\}} \sum_{q' \in X_j(\tilde{\pi})} \Pr[m(q,c) = q']\, \tilde{q}^{X_j}}{\sum_{q \in X_i(\tilde{\pi})} \Pr[q]}.$$

That is, the probability that type $s = \tilde{q}^{X_i}$ sends message $m = \tilde{q}^{X_j}$ when the cost realization is c, is the conditional probability that a type in $X_i(\tilde{\pi})$ would send a message corresponding to a type in $X_j(\tilde{\pi})$. We complement the definition by having threshold types $\underline{q}_i$ send message

$m = \tilde{q}^{X_j}$ whenever they were sending a message $m \in X_j(\tilde{\pi})$. As all messages in $X_i(\tilde{\pi})$ led to the same unaudited decision, the same decision must now be optimal for the principal with experiment $\tilde{\pi}_c$, as the tampering threshold corresponding to $\tilde{q}^{X_i}$ is the same as the threshold for $q \in X_i$. Thus, the designer's expected payoff from $\tilde{\pi}$ and $\tilde{\pi}_c$ coincide.

To obtain an improvement within the set Π^Σ, suppose that $X_L(\tilde{\pi})$ or $X_S(\tilde{\pi})$ is non-empty–otherwise $\tilde{\pi}_c \in \Pi^\Sigma$. By either lowering $\tilde{q}^{X_L} > \underline{q}_L$ or $\tilde{q}^{X_S} > \underline{q}_S$, we can raise the probability of realization $\tilde{q}^{X_H}$ in a way that tampering incentives remain constant but this transformed experiment raises the designer's payoff by raising the probability of the favorable outcome $s = \tilde{q}^{X_H}$. Therefore, if $X_L(\tilde{\pi})$ or $X_S(\tilde{\pi})$ is non-empty, there is an experiment in Π^Σ that gives the designer a higher payoff.

Step 2: Let $\tilde{\pi}^\Sigma \in \Pi^\Sigma$, described by (p_L, p_S, p) with associated scale-up probability τ after an unaudited $m = q$. As the principal will select the agents' preferred decision after a conclusive audit reveals $s = \underline{q}_i$, the expected gain from tampering is $(1-\lambda)(\tau v_H + (1-\tau) v_S - v_i)$ and this establishes the tampering threshold

$$\bar{c}_i \equiv \bar{c}(\underline{q}_i) = (1-\lambda)(\tau v_H + (1-\tau) v_S - v_i) = (1-\lambda)[\tau \Delta_S + (\Delta_i - \Delta_S)]. \tag{22}$$

Suppose first that $\tau > 0$. This requires that the principal's posterior after an unaudited $m = q$ must not fall below $\underline{q}_H$, so that Bayesian updating requires that

$$\frac{(1-p) q + p \sum_{i=L,S} p_i F(\bar{c}_i) q_i}{(1-p) + p \sum_{i=L,S} p_i F(\bar{c}_i)} \geq \underline{q}_H,$$

which, giving the Bayesian consistency constraint $(1-p) q = \mu - p \sum_{i=L,S} p_i q_i$, leads to

$$\begin{aligned} \mu - p \sum_{i=L,S} p_i q_i + p \sum_{i=L,S} p_i F(\bar{c}_i) q_i &\geq \underline{q}_H \left((1-p) + p \sum_{i=L,S} p_i F(\bar{c}_i) \right) \\ p \left(\underline{q}_H \left(1 - \sum_{i=L,S} p_i F(\bar{c}_i) \right) - \sum_{i=L,S} p_i (1 - F(\bar{c}_i)) q_i \right) &\geq \underline{q}_H - \mu, \\ p \left(\sum_{i=L,S} p_i \bar{F}(\bar{c}_i) \left(\underline{q}_H - q_i \right) \right) &\geq \underline{q}_H - \mu, \end{aligned}$$

which, using (1) can be expressed as

$$p \sum_{i=L,S} p_i \frac{\bar{F}(\bar{c}_i)}{p_i^C} \geq 1. \tag{23}$$

Note that if this constraint is slack, then the unaudited posterior is strictly above $\underline{q}_H$ and the principal's sequential rationality implies that $\tau = 1$. But then, experiment $\tilde{\pi}^\Sigma$ cannot be optimal for the designer. Indeed, consider $\pi' \in \Pi^\Sigma$, described by (p'_L, p'_S, p'), that differs from $\tilde{\pi}^\Sigma$ only in that $p' < p$, while $p'_i = p_i$, but such that the constraint (23) is still slack (so that $\tau' = 1$). As tampering thresholds and decisions have not changed, conditional on each realization the designer's expected utility has not changed, but the favorable outcome $s = q$ is now more likely, thus reaching a contradiction. Thus, in a designer's optimal experiment in Π^Σ, (23) must hold with equality.

We now show that the maximum expected utility of the designer is attained on the boundary of Π^Σ, i.e., by setting $p_i = 0$ or $p_i = 1$. But these are precisely the experiments in Π.

Let

$$W_i(\bar{c}_i; \lambda, k) = \lambda\Delta_i + \bar{c}_i\bar{F}(\bar{c}_i) - \mathbb{I}_{\{k=\mathcal{I}\}}\int_0^{\bar{c}_i}\bar{F}(c)\,dc. \tag{24}$$

For experiment $\tilde{\pi}^\Sigma$, the designer's utility can be written as

$$\begin{aligned} V\left(\tilde{\pi}^\Sigma; \lambda, k\right) &= \sum_{i=L,S} p_i\{\lambda\left[(1-p)\,v_H + pv_i\right] \\ &\quad + (1-\lambda)\left[p\bar{F}(\bar{c}_i)\,v_i + \left(1 - p\bar{F}(\bar{c}_i)\right)(\tau v_H + (1-\tau)\,v_S)\right]\} \\ &\quad - \mathbb{I}_{\{k=\mathcal{I}\}}p\sum_{i=L,S} p_i\int_0^{\bar{c}_i}\bar{F}(c)\,dc \\ &= \sum_{i=L,S} p_i\left\{v_i + \lambda(1-p)\Delta_i + \left(1 - p\bar{F}(\bar{c}_i)\right)\bar{c}_i - \mathbb{I}_{\{k=\mathcal{I}\}}p\int_0^{\bar{c}_i}\bar{F}(c)\,dc\right\} \\ &= v_H - (1-\lambda)(1-\tau)\,\Delta_S - p\sum_{i=L,S} p_i\left\{\lambda\Delta_i + \bar{c}_i\bar{F}(\bar{c}_i) - \mathbb{I}_{\{k=\mathcal{I}\}}\int_0^{\bar{c}_i}\bar{F}(c)\,dc\right\} \\ &= v_H - (1-\lambda)(1-\tau)\,\Delta_S - p\sum_{i=L,S} p_i W_i(\bar{c}_i; \lambda, k). \end{aligned} \tag{25}$$

in which we have used (22) for $\bar{c}_i$. Fix an scale-up probability $\tau > 1$–which also determines the tampering thresholds $\bar{c}_i$, see (22)–and consider the optimal $\hat{\pi} \in \Pi^\Sigma$ that maximizes $V(\pi; \lambda, k)$ among experiments in Π^Σ with scale-up probability τ. Then, replacing p with the binding Bayesian updating constraint (23), experiment $\hat{\pi}$ solves

$$\max_{(p_L, p_S)} \frac{\sum_{i=L,S} p_i W_i(\bar{c}_i; \lambda, k)}{\sum_{i=L,S} p_i \frac{\bar{F}(\bar{c}_i)}{p_i^C}}, \text{s.t. } p_L + p_S = 1,$$

which is a quasiconcave program as the objective function is quasilinear—since it is the ratio of linear functionals—and the constraint set is convex (see Boyd and Vandenberghe 2004). Then, there is always an extreme point of the simplex that solves this program. In other words, there is always a robust experiment that maximizes the designer's payoff when selecting experiments in Π^{Σ}. This concludes the proof of part i of the Proposition.

(ii) If the principal organizes to innovate, then the principal is never indifferent after a conclusive audit yields $s = \underline{q}_S$ and the proof of part i then implies that there is always a robust experiment that gives the designer a (weakly) higher payoff. This establishes $V^* = \bar{V}(\mu; \lambda, k)$.

Now suppose that the principal organizes for scale. Note that the payoff from any experiment in Π can be approximated by experiments of the form $\left\{\underline{q}_i + \epsilon, q\right\}$ which have a unique communication equilibrium–i.e., we can express $\bar{V}(\mu; \lambda, k)$ as $\bar{V}(\mu; \lambda, k) = \sup V(\pi; \lambda, k)$, s.t., $\pi = \left\{\underline{q}_i + \epsilon\right.$ $\{L, S\}$. Thus, we must have $V^* \geq \bar{V}(\mu; \lambda, k)$. The proof of part 1 showed that for any experiment $\tilde{\pi}$ such that the principal is never indifferent after a conclusive audit there is a robust experiment that gives the designer a (weakly) higher payoff. Since $V^* > \bar{V}(\mu; \lambda, k)$ implies that the principal must be indifferent after a conclusive audit, then we must have $\underline{q}_S \in S(\pi^*)$ for any equilibrium experiment π^* that yields a payoff $V^* > \bar{V}(\mu; \lambda, k)$.

■

Proof of Lemma 1: Setting $p_S = 0$ and $p_L = 0$ in (25) and using (24) we obtain $v_i(\tau(\bar{c}), \mu; \lambda, k)$ for $i \in \{L, S\}$. Noting from (3) that $\Pr\left[s = \underline{q}_i\right] = \frac{q-\mu}{q-\underline{q}_i}$ gives (8-9). ■

The following lemma complements Proposition 3 as it provides sufficient conditions for expected equilibrium tampering costs to increase when switching from $\pi_S(\tau) \in \Pi_S$ and $\pi_L(\tau) \in \Pi_L$.

Lemma 6. *Suppose that either (a) $\eta(c)$ is convex and $p_L^C > (\Delta_S/\Delta_L)\, p_S^C$; (b) $\eta(c)$ is log-convex; or (c) $\eta(c)$ is log-concave and $p_L^C \eta((1-\lambda)\Delta_L) \geq p_c^S \eta((1-\lambda)\Delta_S)$. Then $C_L(\tau) \geq C_S(\tau)$ for all $\tau \in [0, 1]$.*

Proof of Lemma 6: From (12), $C_i(\bar{c}_i) = p_i^C \eta(\bar{c}_i)$ are the equilibrium expected tampering cost of an experiment in Π_i that induces threshold $\bar{c}_i$. We show that condition (13) holds which implies that $C_L(\tau) \geq C_S(\tau)$ for all $\tau \in [0, 1]$.

(a) Suppose $\eta(c)$ is convex. Then, since $\eta(c)$ is increasing we must have

$$\eta'(c) \leq \frac{\eta(c') - \eta(c)}{c' - c} \leq \eta'(c'), \text{ for } c < c'.$$

To economize on notation, we drop the explicit dependence of $\overline{c}_S(\tau)$ and $\overline{c}_L(\tau)$ on τ and we let $x_i = \underline{q}_H - q_i$. We then have

$$
\begin{aligned}
C_L(\overline{c}_L) - C_S(\overline{c}_S) &= \frac{\underline{q}_H - \mu}{x_L}\eta(\overline{c}_L) - \frac{\underline{q}_H - \mu}{x_S}\eta(\overline{c}_S) \\
&= \frac{\underline{q}_H - \mu}{x_L x_S}\left[x_S\left(\eta(\overline{c}_L) - \eta(\overline{c}_S)\right) - (x_L - x_S)\,\eta(\overline{c}_S)\right] \\
&\geq \frac{\underline{q}_H - \mu}{x_L x_S}\left[x_S \eta'(\overline{c}_S)\left(\overline{c}_L - \overline{c}_S\right) - (x_L - x_S)\,\eta(\overline{c}_S)\right] \\
&\geq \frac{\underline{q}_H - \mu}{x_L x_S}\left[x_S \frac{\eta(\overline{c}_S)}{\overline{c}_S}\left(\overline{c}_L - \overline{c}_S\right) - (x_L - x_S)\,\eta(\overline{c}_S)\right].
\end{aligned}
$$

From (3) we have $\overline{c}_L(\tau) - \overline{c}_S(\tau) = (1-\lambda)(\Delta_L - \Delta_S)$ so

$$
\begin{aligned}
C_L(\overline{c}_L) - C_S(\overline{c}_S) &\geq \frac{\left(\underline{q}_H - \mu\right)\eta(\overline{c}_S)}{x_L x_S}\left[x_S \frac{(1-\lambda)(\Delta_L - \Delta_S)}{\overline{c}_S} - (x_L - x_S)\right] \\
&= \frac{\left(\underline{q}_H - \mu\right)\eta(\overline{c}_S)}{x_L}\left[\frac{\Delta_L - \Delta_S}{\tau \Delta_S} - \frac{x_L - x_S}{x_S}\right].
\end{aligned}
$$

By assumption, $p_L^C > (\Delta_S/\Delta_L)\, p_S^C$, which can can be written as $\Delta_L/x_L > \Delta_S/x_S$ so that

$$
\frac{x_L}{x_S} < \frac{\Delta_L}{\Delta_S} \Rightarrow \frac{x_L - x_S}{x_S} < \frac{\Delta_L - \Delta_S}{\Delta_S} < \frac{\Delta_L - \Delta_S}{\tau \Delta_S}
$$

which implies that $C_L(\overline{c}_L) - C_S(\overline{c}_S) \geq 0$.

(b) Define $\underline{c}$ implicitly by $p_L^C \eta\left((1-\lambda)(\Delta_L - \Delta_S)\right) = p_S^C \eta(\underline{c})$. Then, for $\tau \in [0,1]$ with $\overline{c}_S(\tau) \leq \underline{c}$ and since $\overline{c}_L(\tau) = \overline{c}_S(\tau) + (1-\lambda)(\Delta_L - \Delta_S)$, then

$$
p_L^C \eta\left(\overline{c}_L(\tau)\right) \geq p_L^C \eta\left((1-\lambda)(\Delta_L - \Delta_S)\right) = p_S^C \eta(\underline{c}) \geq p_S^C \eta\left(\overline{c}_S(\tau)\right).
$$

If $\eta(c)$ is log-convex in $[\underline{c}, (1-\lambda)\Delta_L]$

$$
\left(\frac{p_L^C}{p_S^C} \geq\right) \frac{\eta(\underline{c})}{\eta\left(\underline{c} + (1-\lambda)(\Delta_L - \Delta_S)\right)} \geq \frac{\eta\left(\overline{c}_S(\tau)\right)}{\eta\left(\overline{c}_L(\tau)\right)}, \text{ for } \overline{c}_S(\tau) > \underline{c}.
$$

Combining both results we have that $p_S^C \eta\left(\overline{c}_S(\tau)\right) \leq p_L^C \eta\left(\overline{c}_L(\tau)\right)$ for all $\tau \in [0,1]$ so (13) holds.

(b) If $\eta(c)$ is log-concave, then for $\tau < 1$

$$
\frac{\eta\left(\overline{c}_L(\tau)\right)}{\eta\left(\overline{c}_S(\tau)\right)} \geq \frac{\eta\left(\overline{c}_L(1)\right)}{\eta\left(\overline{c}_S(1)\right)} = \frac{\eta\left((1-\lambda)\Delta_L\right)}{\eta\left((1-\lambda)\Delta_S\right)},
$$

so that (13) is satisfied as long as

$$
\frac{\eta\left((1-\lambda)\Delta_L\right)}{\eta\left((1-\lambda)\Delta_S\right)} \geq \frac{p_S^C}{p_L^C}.
$$

■

Proof of Proposition 3: To ease notation, we drop the explicit dependence of $v_i(\tau, \mu; \lambda, k)$, $V_i(\mu; \lambda, k)$, and $m_i(\bar{c}; \lambda, k)$ on λ, and also write $m_i(\tau; k)$ for $m_i(\bar{c}(\tau); k)$.

(i) Using (8) and (9), the difference in the designer's marginal payoff from a higher tampering threshold $\bar{c}$ when moving from integration to separation is

$$\frac{\partial \left(v_i(\tau(\bar{c}), \mu; \mathcal{I}) - v_i(\tau(\bar{c}), \mu; \mathcal{S})\right)}{\partial \bar{c}} = \frac{\partial \left(m_i(\bar{c}; \mathcal{I}) - m_i(\bar{c}; \mathcal{S})\right)}{\partial \bar{c}} \left(\underline{q}_H - \mu\right) = -p_i^C \eta'(\bar{c}) \leq 0.$$

Therefore, the optimal tampering threshold under integration is lower than under separation, $\bar{c}_i^*(\lambda, \mathcal{I}) \leq \bar{c}_i^*(\lambda, \mathcal{S})$, which implies $q_i^*(\lambda, \mathcal{S}) \geq q_i^*(\lambda, \mathcal{I})$.

(ii) Define $\Delta_m(\tau) \equiv m_S(\tau; \mathcal{I}) - m_L(\tau; \mathcal{I})$ as the difference of the slopes in (9), and $\tau_i^*(\mu; k)$ be the optimal scale-up probability under a $k-$allocation when the designer selects an experiment in Π_i. By assumption, $\Delta_m(\tau)$ is single-crossing for $\tau < 1$; let τ_{cr} be the minimum value at which $\Delta_m(\tau) = 0$–we set $\tau_{cr} = 1$ if $\Delta_m(\tau)$ never crosses zero.

We prove the claim in a series of steps.

Step 1: We show that if $\Delta_m(\tau)$ is single-crossing (from negative to positive), then $V_L(\mu; \mathcal{I}) - V_S(\mu; \mathcal{I})$ is single crossing in $\mu \in [\underline{q}_S, \underline{q}_H)$ (from negative to positive); and if μ_{cr} is the minimum value at which $V_L(\mu; \mathcal{I}) = V_S(\mu; \mathcal{I})$, then $\tau_L^*(\mu_{cr}; \mathcal{I}) \geq \tau_{cr}$.

To prove that $V_L(\mu; \mathcal{I}) - V_L(\mu; \mathcal{I})$ is single-crossing, we will make use of the following two results: (a) for $\tau' > \tau_{cr} > \tau''$

$$v_L(\tau', \mu_0; \mathcal{I}) > v_S(\tau'', \mu_0; \mathcal{I}) \Rightarrow v_L(\tau', \mu; \mathcal{I}) > v_S(\tau'', \mu; \mathcal{I}) \text{ for all } \mu > \mu_0, \tag{26}$$

and (b) if $V_L(\mu; \mathcal{I}) \geq V_S(\mu; \mathcal{I})$ then $\tau_L^*(\mu'; \mathcal{I}) \geq \tau_{cr}$ while if $V_L(\mu; \mathcal{I}) < V_S(\mu; \mathcal{I})$ then $\tau_S^*(\mu'', \mathcal{I}) < \tau_{cr}$.

These two results follow from three facts (i) the slope $m_i(\tau; \mathcal{I})$ is increasing in τ; (ii) the designer's payoff at $\mu = \underline{q}_H$ is the same across experiments with the same scale-up probability, i.e., $v_L(\tau, \underline{q}_H; \mathcal{I}) = v_S(\tau, \underline{q}_H; \mathcal{I})$—both facts follow immediately from (9)– and (iii) if $v_L(\tau', \mu; \mathcal{I}) \geq v_S(\tau', \mu; \mathcal{I})$ then $m_S(\tau'; \mathcal{I}) \geq m_L(\tau'; \mathcal{I})$ so that $\tau' \geq \tau_{cr}$.

By way of contradiction, suppose that there exist $\mu', \mu'' \in [\underline{q}_S, \underline{q}_H)$ with $\mu' < \mu''$ and $V_L(\mu'; \mathcal{I}) > V_S(\mu'; \mathcal{I})$ and $V_L(\mu''; \mathcal{I}) < V_S(\mu''; \mathcal{I})$. The first condition implies that $\tau_L^*(\mu'; \mathcal{I}) \geq \tau_{cr}$ and

$$v_L(\tau_L^*(\mu'; \mathcal{I}), \mu'; \mathcal{I}) > \max_{\tau \in [0,1]} v_S(\tau, \mu'; \mathcal{I}) \geq v_S(\tau_S^*(\mu'', \mathcal{I}), \mu'; \mathcal{I}),$$

while the second condition requires $\tau_S^*(\mu'', \mathcal{I}) < \tau_{cr}$. But then, using (26) we obtain

$$V_L(\mu''; \mathcal{I}) \geq v_L(\tau_L^*(\mu'; \mathcal{I}), \mu''; \mathcal{I}) \geq v_S(\tau_S^*(\mu'', \mathcal{I}), \mu''; \mathcal{I}) = V_S(\mu''; \mathcal{I}),$$

thus reaching a contradiction.

Step 2: Let $V_i^{[\tau',1]}(\mu; k) \equiv \max_{\tau \in [\tau',1]} v_i(\tau, \mu; k)$ be the designer's maximum payoff when restricted to experiments with scale-up probability in $[\tau', 1]$. If $\Delta_m(\tau)$ is single-crossing and $C_L(\tau) \geq C_S(\tau)$, for $\tau \in [0,1]$, then we show that $V_L^{[\tau_{cr},1]}(\mu; \mathcal{S}) \geq V_S^{[\tau_{cr},1]}(\mu; \mathcal{S})$, $\mu \geq \mu_{cr}$.

Using the definitions in (9) and (12), we can express the slopes m_i in terms of expected costs as

$$m_i(\tau; \mathcal{I}) = m_i(\tau; \mathcal{S}) + \frac{1}{\underline{q}_H - \mu} C_i(\tau).$$

Then, for $\tau \geq \tau_{cr}$ we have $m_S(\tau; \mathcal{I}) \geq m_L(\tau; \mathcal{I})$ which implies

$$m_S(\tau; \mathcal{S}) \geq m_L(\tau; \mathcal{S}) + \frac{1}{\underline{q}_H - \mu} (C_L(\tau) - C_S(\tau)) \geq m_L(\tau; \mathcal{S}).$$

Using (8) then we must have $v_L(\tau, \mu; \mathcal{S}) \geq v_S(\tau, \mu; \mathcal{S})$ for $\tau \geq \tau_{cr}$ proving the claim in Step 2.

Step 3: We conclude the proof by showing that $V_L(\mu; \mathcal{S}) \geq V_S(\mu; \mathcal{S})$, for $\mu \geq \mu_{cr}$. As step 1 showed that $V_L(\mu; \mathcal{I}) \geq V_S(\mu; \mathcal{I})$ iff $\mu \geq \mu_{cr}$, then taking the contrapositive, would imply that if $V_S(\mu; \mathcal{S}) > V_L(\mu; \mathcal{S})$ then we must have $\mu < \mu_{cr}$, implying that $V_S(\mu; \mathcal{I}) > V_L(\mu; \mathcal{I})$.

To prove Step 3, recall from Step 1 that $\tau_L^*(\mu; \mathcal{I}) \geq \tau_{cr}$ for $\mu \geq \mu_{cr}$. Suppose first that for some $\mu' \geq \mu_{cr}$ we have $\tau_S^*(\mu'; \mathcal{S}) \geq \tau_{cr}$. Then, Step 2 implies.

$$V_S(\mu'; \mathcal{S}) = V_S^{[\tau_{cr},1]}(\mu'; \mathcal{S}) \leq V_L^{[\tau_{cr},1]}(\mu'; \mathcal{S}) = V_L(\mu'; \mathcal{S}).$$

Suppose now that $\tau_S^*(\mu'; \mathcal{S}) < \tau_{cr}$ for some $\mu' \geq \mu_{cr}$. Then, we must have $\tau_L^*(\mu; \mathcal{I}) \geq \tau_{cr} > \tau_S^*(\mu'; \mathcal{S})$ and

$$V_L(\mu'; \mathcal{S}) \geq V_L(\mu'; \mathcal{I}) + C_L(\tau_L^*(\mu; \mathcal{I})) \geq V_S(\mu'; \mathcal{I}) + C_L(\tau_S^*(\mu'; \mathcal{S})) \geq V_S(\mu'; \mathcal{S}),$$

where the first and last inequality exploit the relation $v_i(\tau, \mu; \mathcal{I}) = v_i(\tau, \mu; \mathcal{S}) - C_i(\tau)$ and properties of the maximum, and the second inequality uses the single-crossing property derived in Step 1 and the assumption that $C_L(\tau) \geq C_S(\tau)$ for all $\tau \in [0,1]$. ■

The following lemma will be used in the proof of proposition 4. It guarantees a preference-preserving property among experiments that approve with probability τ.

Lemma 7. *Let $\pi_S(\tau) \in \Pi_S$ and $\pi_L(\tau) \in \Pi_L$ be experiments that approve with probability $\tau \in [0,1]$ and suppose that $p_L^C > (\Delta_S/\Delta_L)\, p_S^C$. Then, for every $\tau \in [0,1]$ and $\lambda > 0$,*
(i) the designer under separation prefers $\pi_S(\tau)$ to experiment $\pi_L(\tau)$.
(ii) if (13) holds, then the designer under integration prefers $\pi_S(\tau)$ to experiment $\pi_L(\tau)$.

Proof of Lemma 7: (i) Using (8) we have

$$v_S(\tau,\mu;\lambda,\mathcal{S}) - v_L(\tau,\mu;\lambda,\mathcal{S}) = (m_L(\bar{c}_L(\tau);\lambda,\mathcal{S}) - m_S(\bar{c}_S(\tau);\lambda,\mathcal{S}))\left(\underline{q}_H - \mu\right). \tag{27}$$

First, if $p_L^C > (\Delta_S/\Delta_L)\, p_S^C$, then we must have that the slopes m_i for $\lambda = 1$ satisfy

$$m_L(\bar{c}_L(\tau);1,\mathcal{S}) = \frac{\Delta_L}{\underline{q}_H - \underline{q}_L} > \frac{\Delta_S}{\underline{q}_H - \underline{q}_S} = m_S(\bar{c}_S(\tau);1,\mathcal{S}). \tag{28}$$

Second, since $\bar{c}_L(\tau) > \bar{c}_S(\tau)$, and

$$\frac{\bar{c}_L(\tau)}{\Delta_L} - \frac{\bar{c}_S(\tau)}{\Delta_S} = (1-\lambda)(1-\tau)\frac{\Delta_S - \Delta_S}{\Delta_L} > 0,$$

it follows that

$$\frac{\lambda}{\overline{F}(\bar{c}_L(\tau))} + \frac{\bar{c}_L(\tau)}{\Delta_L} > \frac{\lambda}{\overline{F}(\bar{c}_S(\tau))} + \frac{\bar{c}_S(\tau)}{\Delta_S}. \tag{29}$$

Using (9), conditions (28) and (29) imply that

$$m_L(\bar{c}_L(\tau);\lambda,\mathcal{S}) > m_S(\bar{c}_S(\tau);\lambda,\mathcal{S}), \tag{30}$$

so that (27) is always positive.

(ii) If (13) holds, then expected tampering costs increase when moving to an experiment in Π_L so we have $\eta(\bar{c}_L(\tau))/\left(\underline{q}_H - \underline{q}_L\right) > \eta(\bar{c}_S(\tau))/\left(\underline{q}_H - \underline{q}_S\right)$. From (9), we have $m_i(\bar{c};\lambda,\mathcal{I}) = m_i(\bar{c};\lambda,\mathcal{S}) + \eta(\bar{c})/\left(\underline{q}_H - \underline{q}_i\right)$. This observation coupled with (30) implies

$$v_S(\tau,\mu;\lambda,\mathcal{I}) - v_L(\tau,\mu;\lambda,\mathcal{I}) = (m_L(\bar{c}_L(\tau);\lambda,\mathcal{I}) - m_S(\bar{c}_S(\tau);\lambda,\mathcal{I}))\left(\underline{q}_H - \mu\right) > 0$$

so that the designer obtains a higher payoff under integration from $\pi_S(\tau)$ than $\pi_L(\tau)$. ■

Proof of Proposition 4: (i-a) Consider experiment $\left\{\underline{q}_i, q(\bar{c})\right\}$ inducing tampering threshold $\bar{c}$. Using (8) and (9), we have

$$\frac{\partial m_i(\bar{c};\lambda,k)}{\partial\lambda} = \frac{\Delta_i}{\overline{F}(\bar{c})}\left(\frac{\mu - \underline{q}_H}{\underline{q}_H - \underline{q}_i}\right) = -\frac{p_i^C \Delta_i}{\overline{F}(\bar{c})},\ \ k \in \{\mathcal{S},\mathcal{I}\},$$

implying

$$\frac{\partial v_i(\tau(\overline{c}), \mu; \lambda, k)}{\partial \lambda} = \Delta_i - p_i^C \frac{\Delta_i}{\overline{F}(\overline{c})} = \Delta_i \Pr[s = q(\overline{c})],$$

which is non-increasing in $\overline{c}$. Therefore, $\partial^2 v_i / \partial(-\lambda) \partial \overline{c}_i \geq 0$.

(i-b) Define the feasible set of tampering thresholds

$$\mathcal{C}_i \equiv [[(1-\lambda)(\Delta_i - \Delta_S), (1-\lambda)\Delta_i] \cup \{0\}] \cap [0, \overline{F}^{-1}(p_i^C / p_i^{FR})], \tag{31}$$

with $p_i^{FR} = \Pr[s = q]$ for experiment $\{\underline{q}_i, 1\}$. To understand $\mathcal{C}_i$, note that $\overline{c} \in \mathcal{C}_i$ must satisfy two conditions. First, it must correspond to some scale-up probability $\tau \in [0,1]$–from (4) this implies that $\overline{c} \in [(1-\lambda)(\Delta_i - \Delta_S), (1-\lambda)\Delta_i] \cup \{0\}$. Second, the experiment $\left\{\underline{q}_i, q(\overline{c})\right\}$ must be feasible–i.e., $q(\overline{c}) \leq 1$–which requires $\bar{F}(\overline{c}) / p_i^C \geq \frac{1-\underline{q}_i}{1-\mu} = 1/p_i^{FR}$ –see (3).

We can write the designer's problem in terms of selecting a tampering threshold $\overline{c}^*$ that solves

$$\max_{\overline{c}} v_i(\tau(\overline{c}), \mu; \lambda, k), \; s.t. \; \overline{c} \in \mathcal{C}_i. \tag{32}$$

The feasible set $\mathcal{C}_i$ is increasing in the strong set order with respect to $-\lambda$ and, from part (i), $v_i(\tau(\overline{c}), \mu; \lambda, k)$ is supermodular in $(\overline{c}, -\lambda)$. Theorem 4' in Milgrom-Shannon (1994) then implies that the set of maximizers of (32) increases in the strong set order sense with $-\lambda$. From (3), for a fixed threshold $\overline{c}$ the experiment $\left\{\underline{q}_i, q(\overline{c})\right\}$ is independent of λ, so the set of optimal experiments $q_i^*(\lambda, k)$ decreases in the strong-order sense with λ.

(ii) We prove simultaneously part (a) and part (b) by appealing to Lemma 7. Let $\tau_i^*(\lambda; k)$ be the scale-up probability of a designer's optimal experiment under a $k-$allocation when restricted to Π_i. Then,

$$V_L(\mu; \lambda, k) = v_L(\tau_L^*(\lambda; k), \mu; \lambda, k) < v_S(\tau_S^*(\lambda; k), \mu; \lambda, k) \leq V_S(\mu; \lambda, k)$$

where the first inequality follows from Lemma 7-i and the last from the definition of $V_S(\mu; \lambda, k)$– see (5). ■

Proof of Lemma 2: If $f(0) > 0$, then whenever $\lambda < 1$ the principal never approves without a conclusive audit if the designer selects $\{\underline{q}_i, \underline{q}_H\}$. In other words, $\tau = 0$ for experiment $\{\underline{q}_i, \underline{q}_H\}$ and

$$v_i(0, \mu; \lambda, k) = v_H - (1-\lambda)\Delta_i + \overline{c}_i(0) - m_i(\overline{c}_i(0); \lambda, k)\left(\underline{q}_H - \mu\right),$$

with $\bar{c}_i(0) = (1-\lambda)(\Delta_i - \Delta_S)$, $i \in \{L, S\}$. Therefore, for experiment $\{\underline{q}_i, \underline{q}_H\}$, the analyst will tamper if $c < \bar{c}_i(0)$—as the principal selects the status quo rather than d_i whenever the audit is inconclusive. We now study conditions such that (a) there exists an experiment that leads to a positive scale-up probability, and (b) the designer's incremental payoff from an experiment that approves with positive probability is positive. These conditions ensure that the designer is responsive to auditing in Π_i.

Consider first (a). The infimum tampering probability among experiments with $\tau > 0$ is $F[\bar{c}_i(0)] = F[(1-\lambda)(\Delta_i - \Delta_S)]$. The experiment that induces the highest posterior if unaudited is $\left\{\underline{q}_i, 1\right\}$, and, for this experiment, $\Pr\left[s = \underline{q}_i\right] = \frac{1-\mu}{1-\underline{q}_i} \equiv p_i^{FR}$. Therefore, there exists an experiment with a positive scale-up probability, iff

$$\frac{p_i^C}{\overline{F}(\bar{c}_i(0))} < p_i^{FR} \Longleftrightarrow \overline{F}((1-\lambda)(\Delta_i - \Delta_S)) > \frac{p_i^C}{p_i^{FR}} (< 1) . \tag{33}$$

If $i = S$, then $\Delta_i - \Delta_S = 0$ and this condition is always satisfied for any $\lambda \in [0, 1]$. If $i = L$, Let $\underline{\Lambda}$ be the the set of $\lambda's$ satisfying (33) if $i = L$, i.e., $\underline{\Lambda} \equiv \left\{\lambda : \bar{F}(\bar{c}_L(0)) \geq p_L^C / p_L^{FR}\right\}$.

Consider now (b). Noting from (4) that $\bar{c}_i(\tau) = \bar{c}_i(0) + \tau(1-\lambda)\Delta_S$, we can differentiate (8)—taking into account (9) and the definition of p_i^C in (1)—to obtain

$$\left.\frac{\partial v_i(\tau, \mu; \lambda, k)}{\partial \tau}\right|_{\tau=0} = (1-\lambda)\Delta_S \left[\left(1 - p_i^C\right) - p_i^C \frac{f(\bar{c}_i(0))}{\left(\overline{F}(\bar{c}_i(0))\right)^2} \left\{\lambda \Delta_i + 1_{\{k=\mathcal{I}\}} \int_0^{\bar{c}_i(0)} \bar{F}(c) d(c)\right\}\right].$$

Note that if $i = S$, then $\bar{c}_S(0) = (1-\lambda)(\Delta_S - \Delta_S) = 0$, so that the condition $\partial v_i(\tau, \mu; \lambda, k)/\partial \tau|_{\tau=0} >$ 0 translates to

$$\frac{1 - p_S^C}{p_S^C \Delta_S} > \lambda \frac{f(0)}{\left(\overline{F}(0)\right)^2},$$

and there is always a $0 < \lambda < 1$ that satisfies this condition.

For $i = L$, define

$$\overline{\Lambda}_L(k) \equiv \left\{\lambda : \frac{f(\bar{c}_L(0))}{\left(\bar{F}(\bar{c}_L(0))\right)^2} \left(\lambda \Delta_L + 1_{\{k=\mathcal{I}\}} \left(\int_0^{\bar{c}_L(0)} \bar{F}(c) dc\right)\right) < \frac{1 - p_L^C}{p_L^C}\right\}.$$

then $\left.\frac{\partial v_L(\tau, \mu; \lambda, k)}{\partial \tau}\right|_{\tau=0} > 0$ for $\lambda \in \overline{\Lambda}_L(k)$ and the designer under a $k-$allocation prefers an experiment that induces a (small) scale-up probability to the commitment experiment.

In summary, if $\Lambda_i(k)$ is the set of auditing intensities such that the designer is responsive to $\lambda-$auditing, then $\Lambda_S(k)$ is always non-empty as long as $f(0) > 0$, while for $i = L$ we have that $\overline{\Lambda}_L(k) \cap \underline{\Lambda} \subseteq \Lambda_L(k) \subseteq \underline{\Lambda}$. ■

Proof of Proposition 5: (i) Applying (15), the principal's equilibrium expected utility when organizing to innovate is

$$U(\lambda, k) = \underline{q}_H + \Pr[s = q_S^*(\lambda, k)]\, \lambda \left(q_S^*(\lambda, k) - \underline{q}_H\right).$$

Proposition 3-i shows that $q_S^*(\lambda, \mathcal{S}) \geq q_S^*(\lambda, I)$ implying $U(\lambda, \mathcal{S}) \geq U(\lambda, \mathcal{I})$.

(ii) Note that for $\lambda = 1$ the designer always selects the commitment experiment, thus, regardless of the task-allocation, $q_S^*(1, k) = \underline{q}_H$. By the definition of designer's responsiveness to auditing in Π_S, there exists $0 < \lambda < 1$ with $q_S^*(\lambda, k) > \underline{q}_H$ and $U(\lambda, k) > U(1, k)$. Therefore, $\lambda^* < 1$. Conversely, if $\lambda^* < 1$ then, for some $k-$allocation, $U(\lambda^*, k) > U(1, k)$ which implies $q_S^*(\lambda^*, k) > \underline{q}_H$ meaning that the designer is responsive to auditing. Lemma 2 then shows that $f(0) > 0$ is sufficient for the designer to be responsive to auditing under separation in Π_S. ■

Proof of Lemma 3: Setting $k = \mathcal{S}$ in (8) and (9) we have

$$\begin{aligned} v_S(\tau(\bar{c}), \mu; \lambda, \mathcal{S}) &= v_H - (1-\lambda)\Delta_S + \bar{c} - \Delta_S \frac{\underline{q}_H - \mu}{\underline{q}_H - \underline{q}_S}\left(\frac{\lambda}{\overline{F}(\bar{c})} + \frac{\bar{c}}{\Delta_S}\right) \\ &= v_H - (1-\lambda)\Delta_S + \left(1 - p_S^C\right)\bar{c} - \lambda \frac{p_S^C \Delta_S}{\overline{F}(\bar{c})}. \end{aligned}$$

Equation (32) defines the designer's problem and the feasible set of tampering thresholds $\mathcal{C}_S = [0, (1-\lambda)\Delta_S] \cap [0, \overline{F}^{-1}(p_S^C / p_S^{FR})]$ is defined in (31). Whenever it exists, the marginal payoff from increased tampering is

$$\frac{\partial v_S(\tau(\bar{c}), \mu; \lambda, \mathcal{S})}{\partial \bar{c}} = \left(1 - p_S^C\right) - \lambda p_S^C \Delta_S \frac{f(\bar{c})}{\left(\overline{F}(\bar{c})\right)^2} = \lambda p_S^C \Delta_S \left(\frac{\phi_S}{\lambda} - L(\bar{c})\right).$$

The single-crossing condition implies that $v_S(\tau(\bar{c}), \mu; \lambda, \mathcal{S})$ is quasiconcave in $\bar{c}$. Suppose first that $\partial v_S(\tau(0), \mu; \lambda, \mathcal{S})/\partial \bar{c} = \lambda p_S^C \Delta_S ((\phi_S/\lambda) - L(0)) \leq 0$, implying $\partial v_S(\tau(\bar{c}), \mu; \lambda, \mathcal{S})/\partial \bar{c} \leq 0$, for $\bar{c} \geq 0$. In this case, we have $\bar{c}^* = 0$, and the designer selects the commitment experiment $\left\{0, \underline{q}_H\right\}$. Suppose now that $\lambda p_S^C \Delta_S ((\phi_S/\lambda) - L(0)) > 0$, and let $\bar{c}_{crit}$ be the minimum threshold that satisfies $\partial v_S(\tau(\bar{c}_{crit}), \mu; \lambda, \mathcal{S})/\partial \bar{c} = 0$ (and set $\bar{c}_{crit} = \infty$ if no such threshold exists). Then, the solution to the designer's problem satisfies

$$\bar{c}^*(\lambda) = \min\left[\bar{c}_{crit}(\lambda), (1-\lambda)\Delta_S, \bar{c}_{FR}\right].$$

■

Proof of Lemma 4: Note first that if the designer selects an "up-or-down" experiment when $\lambda = 1$, so that $p_L^C \Delta_L < p_S^C \Delta_S$, then we must have $\phi_L > \phi_S$. Indeed, since $1 - p_S^C < 1 - p_L^C$ then

$$p_L^C \Delta_L < p_S^C \Delta_S \Rightarrow \frac{p_S^C \Delta_S}{p_L^C \Delta_L} > 1 > \frac{1 - p_S^C}{1 - p_L^C} \Rightarrow \phi_L > \phi_S.$$

We now show that if $L(c) - (\phi_i/\lambda)$ is single-crossing in $[0, \Delta_i]$ and $\phi_L > \phi_S$ then $\bar{c}_S^*(\lambda, \mathcal{S}) \leq \bar{c}_L^*(\lambda, \mathcal{S})$; that is, the tampering threshold under separation is larger when the designer's experiment is restricted to Π_L rather than restricted to Π_S. Then, (15) implies that the principal's is weakly better-off when the designer is restricted to Π_L rather than Π_S.

From lemma 3 and the definition of $\bar{c}_{FR}$ we can write

$$\bar{c}_i^*(\lambda, \mathcal{S}) = \min \left[L^{-1}(\phi_i/\lambda), (1 - \lambda) \Delta_i, \bar{F}^{-1} \left(p_i^C / p_i^{FR} \right) \right].$$

Since $p_i^C / p_i^{FR} = \frac{q_H - \mu}{q_H - \underline{q}_i} \frac{1 - \underline{q}_i}{1 - \mu}$ increases in $\underline{q}_i$, we have $\bar{F}^{-1} \left(p_S^C / (1 - \mu) \right) < \bar{F}^{-1} \left(p_L^C / (1 - \mu) \right)$. Second, if $L(c) - (\phi_i/\lambda)$ is single crossing and $\phi_L > \phi_S$, we must have $L^{-1}(\phi_S/\lambda) < L^{-1}(\phi_L/\lambda)$. Combining both observations with $\Delta_S < \Delta_L$, we must have $\bar{c}_S^*(\lambda, \mathcal{S}) \leq \bar{c}_L^*(\lambda, \mathcal{S})$. ■

Proof of Proposition 6: (i) If Proposition 3-ii holds, then a preference for status-quo experiments is preserved under integration, i.e., $i^* (\mathcal{S}, \lambda) = S \Rightarrow i^* (\mathcal{I}, \lambda) = S$. Proposition 3-i already shows that the principal prefers to separate tasks whenever $i^* (\mathcal{S}, \lambda) = i^* (\mathcal{I}, \lambda)$ as $q_i^*(\lambda, \mathcal{S}) \geq q_i^*(\lambda, \mathcal{I})$ implies $\bar{c}_i^*(\lambda, \mathcal{S}) \geq \bar{c}_i^*(\lambda, \mathcal{I})$. If, additionally, Lemma 4 holds, then the principal prefers that the designer selects an "up-or-down" experiment rather than a status-quo experiment. Since she also prefers the designer's experiment under separation when $i^* (\mathcal{S}, \lambda) = i^* (\mathcal{I}, \lambda) = S$, then we must have that she prefers separation when $i^* (\mathcal{S}, \lambda) = L$ and $i^* (\mathcal{I}, \lambda) = S$.

(ii) From the proof of Proposition 6-i, for integration to be optimal, we must have an adverse switch so that $i^* (\mathcal{S}, \lambda) = S$ and $i^* (\mathcal{I}, \lambda) = L$. We show that if $W (\Delta_S, (1 - \lambda) \Delta_S) > W (\Delta_L, (1 - \lambda) \Delta_L)$ then we can find a range of parameters so that the principal strictly prefers to integrate tasks. To economize on notation, let

$$T (\Delta_i, \tau) \equiv \frac{\lambda \Delta_i}{\bar{F}(\bar{c}_i(\tau))} + \bar{c}_i(\tau),$$

so that $W (\Delta_i, \bar{c})$ defined in (18) simplifies to $W (\Delta_i, \bar{c}) = \eta(\bar{c}) / T (\Delta_i, \tau (\bar{c}))$.

We first derive a sufficient condition for adverse switches when restricted to high scale-up probabilities, i.e., such that for all $\tau \geq \tilde{\tau}$ we have

$$v_S(\tau,\mu;\lambda,\mathcal{S}) \geq v_L(\tau,\mu;\lambda,\mathcal{S}) \text{ and } v_S(\tau,\mu;\lambda,\mathcal{I}) \leq v_L(\tau,\mu;\lambda,\mathcal{I}).$$

Using (8), (9), and the definition of $\overline{c}_i(\tau)$ in (4) these two conditions translate to

$$\begin{aligned} p_S^C T(\Delta_S,\tau) &\leq p_L^C T(\Delta_L,\tau)\,, \text{ and} \\ p_S^C \left(T(\Delta_S,\tau)+\eta\left(\overline{c}_S(\tau)\right)\right) &\geq p_L^C \left(T(\Delta_L,\tau)+\eta\left(\overline{c}_L(\tau)\right)\right), \end{aligned}$$

which simplifies to

$$\frac{T(\Delta_S,\tau)+\eta\left(\overline{c}_S(\tau)\right)}{T(\Delta_L,\tau)+\eta\left(\overline{c}_L(\tau)\right)} \geq \frac{p_L^C}{p_S^C} \geq \frac{T(\Delta_S,\tau)}{T(\Delta_L,\tau)}. \tag{34}$$

Note that a necessary condition for (34) is that $T(\Delta_L,\tau)\,\eta\left(\overline{c}_S(\tau)\right) \geq T(\Delta_S,\tau)\,\eta\left(\overline{c}_L(\tau)\right)$. , which is implied by

$$W\left(\Delta_S,\overline{c}_S(\tau)\right) \geq W\left(\Delta_L,\overline{c}_L(\tau)\right). \tag{35}$$

The condition $W\left(\Delta_S,(1-\lambda)\Delta_S\right) > W\left(\Delta_L,(1-\lambda)\Delta_L\right)$ is equivalent to (35) setting $\tau=1$. Continuity of W and $\overline{c}_i(\tau)$ implies that there is $\tilde{\tau}<1$, so that (35) is satisfied for $\tau>\tilde{\tau}$. In fact, since $p_L^C/p_S^C = \left(\underline{q}_H - \underline{q}_S\right)/\left(\underline{q}_H - \underline{q}_L\right)$ and W does not depend on the prior μ, then we can find $\underline{q}_S$ with $\underline{q}_L < \underline{q}_S < \underline{q}_H$ so that p_L^C/p_S^C is closed to 1 and (34) holds for $\tau>\tilde{\tau}$.

Finally, we have that both under separation and integration, the designer's optimal experiment tends to the maximum scale-up probability as $\mu \to \underline{q}_H$, i.e., $\tau\left(\overline{c}_i^*(\lambda,k)\right) \to 1$ as $\mu \to \underline{q}_H$. Then we can find μ, with $\underline{q}_S < \mu < \underline{q}_H$ so that the principal's optimal experiment satisfies $\tau\left(\overline{c}_i^*(\lambda,k)\right) > \underline{\tilde{\tau}}$, $k=\mathcal{S},\mathcal{I}$. This implies that we have an adverse switch for λ: the designer under separation would select $i^*(\mathcal{S},\lambda)=S$ but under integration he would select $i^*(\mathcal{I},\lambda)=L$. For large α_L, so that "up-or-down" experiments are more valuable to the principal, then she would optimally integrate tasks.

(iii) If $p_L^C > (\Delta_S/\Delta_L)\,p_L^C$, then setting $\lambda=1$ leads to experiment $\{\underline{q}_S,\underline{q}_H\}$ regardless of task allocation. Then, Proposition 5 shows that if $f(0)>0$, then $\lambda^*<1$. If $p_L^C < (\Delta_S/\Delta_L)\,p_L^C$, however, the designer selects $\left\{\underline{q}_L,\underline{q}_H\right\}$ if $\lambda=1$ and setting $\lambda<1$ may trigger an adverse switch. If, however, the designer under a $k-$allocation is $\lambda'-$sensitive in Π_L and $i^*(k,\lambda')=L$, then

$$\lambda' \frac{F\left(\overline{c}_{\mathcal{S}}^*(\lambda;i^*)\right)}{\bar{F}\left(\overline{c}_{\mathcal{S}}^*(\lambda;i^*)\right)} > 0\left(=\frac{F(0)}{\bar{F}(0)}\right)$$

and (15) shows that the principal's utility increases when reducing λ from 1 to λ'. ■

Proof of Corollary 1: Suppose that the designer is responsive to auditing but $\lambda^* = 1$. Then, from Proposition 6-iii we must have that $p_L^C < (\Delta_S/\Delta_L)\, p_S^C$ but for each λ so that he prefers some experiment $\{\underline{q}_L, q\}$, $q > \underline{q}_H$ to $\{\underline{q}_L, \underline{q}_H\}$ it must be that $i^*(k,\lambda) = S$. Since $p_L^C < (\Delta_S/\Delta_L)\, p_S^C$, then Lemma 4 implies that for any λ the designer's choice under separation improves the principal's utility, and she cannot be worse-off by committing to rule out the status-quo. But in this case, the principal's problem converts to a situation in which she organizes to innovate (as she selects from two decisions, d_L and d_H) and Proposition 6-ii implies $\lambda^* < 1$. ■

Proof of Proposition 7: Suppose first that the principal organizes to innovate and, when auditing intensity is λ, the designer under separation selects $\pi_S = \{\underline{q}_S, q\}\, (= \{0, q\})$, with $p = \Pr\left[s = \underline{q}_S\right] = (q - \mu)/\left(q - \underline{q}_S\right)$.[37] We first show that we must have

$$\lambda \leq \frac{1 - p_S^C}{1 - p_S^C + p - p_S^C} \equiv \widetilde{\lambda}_S(p). \tag{36}$$

To see this, we express $v_S(\tau(p), \mu; \lambda, \mathcal{S})$ as a function of p: using (3) we have $\bar{F}(\bar{c}_S(p)) = p_S^C/p$ and we can write (8-9) for $k = \mathcal{S}$ as

$$\begin{aligned} v_S(\tau(p), \mu; \lambda, \mathcal{S}) &= v_H - (1-\lambda)\Delta_S + \bar{c}_S(p) - p_S^C\left(\frac{\lambda\Delta_S}{\overline{F}(\bar{c}_S(p))} + \bar{c}_S(p)\right) \\ &= v_S + \lambda\Delta_S + \bar{c}_S(p) - \lambda\frac{p_S^C\Delta_S}{\overline{F}(\bar{c}_S(p))} - p_S^C\bar{c}_S(p) \\ &= v_S + \lambda\Delta_S(1-p) + \left(1 - p_S^C\right)\overline{F}^{-1}(p_S^C/p). \end{aligned}$$

Designer's optimality of π_S requires

$$\begin{aligned} \left(\overline{F}(\bar{c}_S(p)) =\right)\frac{p_S^C}{p} &\leq \overline{F}((1-\lambda)\,\Delta_S), \\ v_S(\tau(p'), \mu; \lambda, \mathcal{S}) &\leq v_S(\tau(p), \mu; \lambda, \mathcal{S}) \text{ for } p' \in \left[p_S^C, p\right]. \end{aligned}$$

The first condition follows from $\bar{c}_S(p) \leq (1-\lambda)\,\Delta_S$, as the gain from tampering is bounded by $(1-\lambda)\,\Delta_S$, while the second is the designer's incentive compatibility constraint when

[37] Recall that, regardless of the cost distribution, the principal prefers to separate tasks when organizing to innovate–see Proposition 5-i.

comparing π to robust experiments that are less informative than π.[38] Setting $p' = p_S^C$ above, and obviating the common term v_S, incentive compatibility implies

$$\lambda\Delta_S\left(1-p_S^C\right) \;\; \leq \;\; \lambda\Delta_S\left(1-p\right)+\left(1-p_S^C\right)\overline{F}^{-1}(p_S^C/p) \leq \lambda\Delta_S\left(1-p\right)+\left(1-p_S^C\right)\left(1-\lambda\right)\Delta_S,$$

from which we obtain (36). Taking into account $q = (\mu - \underline{q}_S)/(1-p)$, (36) can be written as

$$\widetilde{\lambda}_S(q) = \frac{q-\underline{q}_S}{q-\underline{q}_S+q-\underline{q}_H}.$$

We now derive the cost distributions that would lead the designer to select $\pi_S(p) = \{\underline{q}_S, q(p)\}$ when auditing is $\widetilde{\lambda}_S(p)$. Suppose that experiment $\pi_S(p)$ leads to the principal's rubberstamping, $\tau(p) = 1$, so that $\overline{c}_S(p) = \overline{F}^{-1}(p_S^C/p) = (1-\widetilde{\lambda}_S(p))\Delta_S$. Incentive compatibility requires that for any $p' \in \left[p_S^C, p\right]$,

$$\left(1-p_S^C\right)\overline{F}^{-1}(p_S^C/p') \leq \widetilde{\lambda}_S(p)\Delta_S\left(p'-p\right)+\left(1-p_S^C\right)\left(1-\widetilde{\lambda}_S(p)\right)\Delta_S.$$

Using $\overline{F}^{-1}(p_S^C/p') = F^{-1}(\left(p'-p_S^C\right)/p')$ and simplifying we have

$$\frac{p'-p_S^C}{p'} \leq F\left(\widetilde{\lambda}_S(p)\Delta_S\frac{p'-p_S^C}{1-p_S^C}\right).$$

Alternatively, letting $c = \widetilde{\lambda}_S(p)\Delta_S\left(p'-p_S^C\right)/\left(1-p_S^C\right)$, we have

$$F\left(c\right) \geq \frac{c}{c+\widetilde{\lambda}_S(p)\frac{p_S^C\Delta_S}{1-p_S^C}} = \frac{c}{c+\frac{p_S^C\Delta_S}{1-2p_S^C+p}} \text{ for } c \leq \widetilde{\lambda}_S(p)\Delta_S\frac{p-p_S^C}{1-p_S^C}. \tag{37}$$

That is, the likelihood of low tampering costs must be sufficiently high to allow the principal to approve with low probability if the experiment is not very informative. Note that our argument didn't require the distribution to be smooth or to have a density. One distribution that satisfies (37) is supported only on two cost realizations, 0 and $\widetilde{\lambda}_S(p)\Delta_S\frac{p-p_S^C}{1-p_S^C}$, with

$$\Pr\left[c=0\right] = \left(p-p_S^C\right)/p, \tag{38}$$

and, in equilibrium, the analyst only tampers if $c = 0$ so that expected tampering costs are zero. From (14), for each $\pi_S(q) = \{\underline{q}_S, q\}$, with $p = \Pr[s = \underline{q}_S]$, auditing $\widetilde{\lambda}_S(q)$, and cost distribution satisfying (37), the principal's utility is

$$U\left(\pi_S(q)\right) = \underline{q}_H + \left(1-p\right)\widetilde{\lambda}_S(q)\left(q-\underline{q}_H\right) = \underline{q}_H + \frac{\left(q-\underline{q}_H\right)\left(\mu-\underline{q}_S\right)}{2q-\underline{q}_S-\underline{q}_H}, \tag{39}$$

[38] If $\tau(p) = 1$, the designer cannot improve scale-up probability by switching to a robust experiment that is more informative than $\pi = \{\underline{q}_S, q\}$ so that trivially $v_S(\tau\left(p'\right), \mu; \lambda, \mathcal{S}) \leq v_S(\tau\left(p\right), \mu; \lambda, \mathcal{S})$ for $p' > p$. When implementing $\pi(p)$ with auditing intensity $\lambda_S(p)$ we will look at cost distributions for which $\tau(p) = 1$.

which is increasing in q. Thus, the principal optimally sets $q = 1$—alternatively, $p = (1 - \mu)/(1 - \underline{q}_S)$. Since $\underline{q}_S = 0$ when organizing to innovate, then we must have $\lambda_{opt} = \widetilde{\lambda}_S(1-\mu) = 1/(2 - \underline{q}_H)$. Setting $p = (1-\mu)$ in (38), we obtain $\Pr[c = 0] = \mu(1 - \underline{q}_H)/(\underline{q}_H(1-\mu))$ for the cost distribution supported on 0 and $\widetilde{c} = \lambda_{opt}\Delta_S \frac{1-\mu-p_S^C}{1-\mu-p_S^C} = (1-\underline{q}_H)/(2-\underline{q}_H)$; this distribution minimizes tampering costs among all those inducing a fully informative experiment. As this distribution induces zero costs on the analyst, the designer under integration and separation would select the same experiment.

Suppose now that the principal organizes for scale and $p_L^C < (\Delta_S/\Delta_L)\, p_S^C$, so he selects $\{\underline{q}_L, \underline{q}_H\}$ if $\lambda = 1$. We prove the optimality of (19) and (20) in two steps. First, we show that the cost distributions obtained above satisfy the designer's incentive compatibility for experiments $\{\underline{q}_L, q\} \in \Pi_L$ which now must also account for the possibility of switching to an status-quo experiment. Second, after deriving the optimal auditing and cost distribution, we show that the principal cannot do better by instead inducing the designer to select a status-quo experiment.

Consider $\pi_L(p) = \{\underline{q}_L, q\}\, (= \{0, q\})$ with $p = \Pr\left[s = \underline{q}_L\right] = (q-\mu)/(q - \underline{q}_L)$. Then, by a similar reasoning as above, if the designer under separation selects π_L when auditing is λ then we must have

$$\lambda \leq \frac{1 - p_L^C}{1 - p_L^C + p - p_L^C} \equiv \widetilde{\lambda}_L(p).$$

Consider a cost distribution supported on 0 and $\bar{c}^* \equiv \widetilde{\lambda}_L(p)\Delta_L \frac{p - p_L^C}{1 - p_L^C}$, with

$$\Pr[c = 0] = \left(p - p_L^C\right)/p. \tag{40}$$

The argument above showed that if the auditing intensity is $\widetilde{\lambda}_L(p)$, the designer, when restricted to Π_L, selects experiment $\pi_L(p)$. We now show that the designer does not wish to switch and select instead $\{\underline{q}_S, q'\} \in \Pi_S$. Using (40) we first note that any $\{\underline{q}_S, q'\}$ that leads to a positive scale-up probability must satisfy

$$\Pr\left[s = \underline{q}_S\right] \geq \frac{p_S^C}{1 - \Pr[c = 0]} = p\frac{p_S^C}{p_L^C}.$$

as the analyst always tampers when $c = 0$. Moreover, all experiments satisfying this condition lead to the principal rubberstamping the analyst's recommendation. Recalling that $v_L = 0$, experiments in Π_S are dominated by $\pi_L(p)$ iff

$$\begin{aligned} & v_S + \lambda(1 - \Pr[s = \underline{q}_S])\Delta_S + (1-\lambda)\left(1 - p_S^C\right)\Delta_S \\ \leq \; & \lambda(1-p)\Delta_L + (1-\lambda)\left(1 - p_L^C\right)\Delta_L. \end{aligned}$$

Rearranging and noting that $v_S + \Delta_S - \Delta_L = v_L = 0$, then incentive compatibility requires

$$0 \leq \lambda \Pr[s = \underline{q}_S]\Delta_S - \lambda p \Delta_L + (1-\lambda)\, p_L^C \Delta_L - (1-\lambda)\, p_L^C \Delta_L.$$

Since we assumed that $p_L^C \Delta_L < p_S^C \Delta_S$, then

$$\begin{aligned} & \lambda \Pr\left[s = \underline{q}_S\right]\Delta_S - \lambda p \Delta_L + (1-\lambda)\, p_S^C \Delta_S - (1-\lambda)\, p_L^C \Delta_L \\ \geq \; & \lambda p \frac{p_S^C}{p_L^C}\Delta_S - \lambda p \Delta_L + (1-\lambda)\, p_L^C \Delta_L - (1-\lambda)\, p_L^C \Delta_L \\ = \; & \left(\frac{\lambda p}{p_L^C} + (1-\lambda)\right)\left(p_S^C \Delta_S - p_L^C \Delta_L\right) \geq 0. \end{aligned}$$

Therefore, the designer subject to auditing $\widetilde{\lambda}_L(p)$ and the two-point cost distribution above would optimally select $\pi_L(p)$ with $\Pr[s = \underline{q}_L] = p$.

From (15), the principal's utility is

$$\begin{aligned} U(\pi_L(p)) &= \underline{q}_H + \left(\underline{q}_H - \mu\right)\left(\widetilde{\lambda}_L(p)\frac{p - p_L^C}{p_L^C} + \left(1 + \widetilde{\lambda}_L(p)\frac{p-p_L^C}{p_L^C}\right)\left(\frac{\alpha_L - \underline{q}_H}{\underline{q}_H - \underline{q}_L}\right)\right) \\ &= \underline{q}_H + \left(\underline{q}_H - \mu\right)\left(\frac{1-p_L^C}{p_L^C}\frac{p - p_L^C}{1 - 2p_L^C + p} + \left(1 + \frac{1-p_L^C}{p_L^C}\frac{p - p_L^C}{1 - 2p_L^C + p}\right)\left(\frac{\alpha_L - \underline{q}_H}{\underline{q}_H - \underline{q}_L}\right)\right) \end{aligned}$$

where we have used $F(\bar{c}^*)/\bar{F}(\bar{c}^*) = \Pr[c=0]/(1 - \Pr[c=0]) = \left(p - p_L^C\right)/p_L^C$. This expression is increasing in p, so that the principal sets $p = 1 - \mu$ —and the designer selects in response a fully informative experiment—implying that $\lambda_{opt} = \widetilde{\lambda}_L(1-\mu) = 1/\left(2 - \underline{q}_H\right)$.

To end this proof, we show that the principal cannot improve her payoff by instead inducing the designer to select an experiment in Π_S. To see this, suppose that the designer were restricted to select experiments in Π_S. Then, the principal's maximum expected utility is obtained from (39) by setting $q = 1$. But then we have

$$U(\pi_S(q=1)) = \underline{q}_H + \frac{\left(1 - \underline{q}_H\right)\left(\mu - \underline{q}_S\right)}{2 - \underline{q}_S - \underline{q}_H} \leq \underline{q}_H + \frac{\left(1 - \underline{q}_H\right)\left(\mu - \underline{q}_L\right)}{2 - \underline{q}_L - \underline{q}_H} \leq U(\pi_L(q=1)).$$

■

Proof of Corollary 2: Under the conditions of Proposition 7, the principal does not gain by ruling out a decision(s). Therefore, consider the case not cover by Proposition 7: the principal organizes for scale and $p_L^C > (\Delta_S/\Delta_L)\, p_S^C$, so that the designer selects a status-quo experiment if $\lambda = 1$. This also implies that for any cost distribution under separation the designer would select a status-quo experiment–see Proposition 4-ii(a). The proof of Proposition 7 shows that the principal's maximum payoff when restricted to "up-or-down" experiments is greater than her maximum payoff if the designer selects a status-quo experiment. Therefore, the principal would rule out the status-quo decision d_S, and the optimal auditing and cost distribution would then be given by Proposition 7. ■

Proof of Lemma 5: Suppose that the principal organizes for scale and consider experiment $\{\underline{q}_S, q\}$, $q > \underline{q}_H$. The principal is indifferent between decisions d_L and d_S both after a conclusive audit determines $s = \underline{q}_S$ and after an unaudited message $m = \underline{q}_S$. Let $\tau_I(m)$ be the probability of choosing d_S after message m and a conclusive audit finds $s = \underline{q}_S$ and $\tau_U(m)$ be (i) the probability of selecting d_S after an unaudited $m = \underline{q}_S$, and (ii) the probability of selecting d_H after an unaudited $m = q$. Then,

$$v_I(\underline{q}_S, \underline{q}_S) = v_L + \tau_I(\underline{q}_S)(v_S - v_L)$$
$$v_U(q) = v_S + \tau_U(q)(v_H - v_S)$$

are, respectively, the payoff from truthtelling after a conclusive audit if $s = \underline{q}_S$ and the payoff after an inconclusive audit if $m = q$, and note that the gain from tampering after $s = \underline{q}_S$ is

$$\widetilde{c} \equiv (1-\lambda)[(1-\tau_U(\underline{q}_S))(\Delta_L - \Delta_S) + \tau_U(q)\Delta_S] + \lambda(\tau_I(q) - \tau_I(\underline{q}_S))(\Delta_L - \Delta_S), \qquad (41)$$

so that the probability of truthfull communication conditional on $s = \underline{q}_S$ is $\bar{F}(\widetilde{c})$. The designer's equilibrium payoff from a non-robust status-quo experiment V_S^{NR} can then be written as

$$V_S^{NR} \equiv \Pr[s = \underline{q}_S] \left(v_I(\underline{q}_S, \underline{q}_S) + F(\widetilde{c})\widetilde{c} \right) + \Pr[s = q] \left(\lambda v_H + (1-\lambda) v_U(q) \right). \qquad (42)$$

Recall that $\bar{c}_S$, defined in (3), is the tampering gain that leads the designer's posterior to $\underline{q}_H$ after an unaudited message $m = q$, and it satisfies $\bar{F}(\bar{c}_S) = p_S^C \left(\frac{q - \underline{q}_S}{q - \mu} \right)$. Therefore, principal's sequential rationality requires $\tau_U(q) = 1$ if $\widetilde{c} < \bar{c}_S$ and $\tau_U(q) = 0$ if $\widetilde{c} > \bar{c}_S$.

To find the designer's maximum payoff, fix $\tau_I(\underline{q}_S)$, which pins down the payoff under an audited truthful report $v_I(\underline{q}_S, \underline{q}_S)$. Suppose that $\widetilde{c} < \overline{c}_S$. Then, $\tau_U(q) = 1$ and $v_U(q) = v_H$, so that (42) increases in $\widetilde{c}$. Likewise, if $\widetilde{c} > \overline{c}_S$ then $\tau_U(q) = 0$ and $v_U(q) = v_S$, so that (42) again increases with $\widetilde{c}$. In other words, the designer's payoff increases with the tampering gain whenever it is different from the threshold for a robust experiment. Suppose now that $\widetilde{c} = \overline{c}_S$. Then, to maintain (41) constant we must decrease scale-up probability $\tau_U(q)$ whenever $\tau_I(q)$ increases, but this reduces $v_U(q)$ and so V_S^{NR} decreases with $\tau_I(q)$–see (42). Therefore, for fixed $\tau_I(\underline{q}_S)$ the maximum payoff to the designer is either achieved in an equilibrium in which $\widetilde{c} = \overline{c}_S$ or an equilibrium which maximizes $\widetilde{c}$ –which, if $\max \widetilde{c} > \overline{c}_S$ is attained by setting $\tau_I(q) = 1$.

We now show that the maximum payoff to the designer is obtained when $\widetilde{c} = \overline{c}_S$ and $v_I(\underline{q}_S, \underline{q}_S) = v_U(\underline{q}_S) = v_S$ so that the principal rewards truthtelling. First, note that for $v_I(\underline{q}_S, \underline{q}_S) = v_U(\underline{q}_S) = v_S$ there is always an equilibrium in which the principal scales-up with positive probability after an unaudited $m = q$. This follows as setting $\tau_U(q)$ at the level of a robust experiment and adjusting $\tau_I(q) < 1$ would give $\widetilde{c} \equiv (1-\lambda)[\tau_U(q)\Delta_S] + \lambda(\tau_I(q) - 1)(\Delta_L - \Delta_S) < \overline{c}$. This also implies that the designer's payoff exceeds $v_S + \lambda p v_H$ as the principal scales-up also when the audit is inconclusive. But, if $\widetilde{c} > \overline{c}_S$ then the maximum payoff to the designer cannot be above $v_S + \lambda p v_H$, as the principal scales-up only if a conclusive audit yields $s = q$. It follows that the designer's optimal is achieved when $\widetilde{c} = \overline{c}_S$. Finally, suppose that the optimal is achieved for $v_I(\underline{q}_S, \underline{q}_S) < v_S$ or $v_U(\underline{q}_S) < v_S$ so that the principal (partially) punishes truthtelling when $m = \underline{q}_S$. Increasing either value while maintaining $\widetilde{c} = \overline{c}_S$ raises the payoff $v_U(q)$ and thus increasing V_S^{NR}. Therefore at the maximum we must have $v_I(\underline{q}_S, \underline{q}_S) = v_U(\underline{q}_S) = v_S$

In summary, the designer's maximum payoff from experiment $\{\underline{q}_S, q\}$ is achieved by choosing d_S after the analyst truthfully sends $m = \underline{q}_S$ but punishing tampering with the lowest $\tau_I(q)$ consistent with a tampering gain of $\widetilde{c} = \overline{c}_S$. Note that if the principal rubberstamps the analyst's recommendation for a robust experiment, then the lowest $\tau_I(q)$ consistent with $\widetilde{c} = \overline{c}_S$ is precisely $\tau_I(q) = 1$. Therefore, a robust status-quo experiment $\{\underline{q}_S, q\}$ achieves the designer's maximum payoff in any PBE following $\{\underline{q}_S, q\}$ if the principal rubberstamps the analyst's recommendation. ■

Proof of Proposition 8: In the text. ■