

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/148939>

Copyright and reuse:


This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it.

Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk



Essays on Information Economics

by

Ayush Pant

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Economics

September 2019

Contents

Acknowledgments	v
Declarations	vi
Abstract	vii
Chapter 1 Feedback on Ideas	1
1.1 Introduction	1
1.2 The model	5
1.3 Benchmark: Single agent problem	7
1.3.1 No information (NI) about θ	7
1.3.2 Full information (FI) about θ	8
1.3.3 Comparing β_0^{NI} and β_0^{FI}	11
1.3.4 An important definition	11
1.4 Strategic supervisor	11
1.4.1 Preliminaries	11
1.4.2 Analysis	13
1.4.3 Welfare analysis	20
1.5 Extensions	22
1.5.1 Benevolent supervisor and time-constrained players	23
1.5.2 Perfect recall of previous ideas	24
1.5.3 Alternate interpretations	27
1.6 Conclusion	28
1.7 Figures	29
1.8 Appendix	32
A Proofs from the main text	32
B Additional proofs not in the main text	45
C Committed supervisor	47
Chapter 2 The Newsroom Dilemma	53
2.1 Introduction	53
2.1.1 Stylized facts and new media studies literature	56
2.1.2 Contributions to related economics literature	57
2.2 A model of the newsroom dilemma	59
2.2.1 Solution concept and equilibria selection	62

2.2.2	Discussion of assumptions	62
2.2.3	Preliminary observations and strategies	63
2.3	Competition leads to better reporting	64
2.3.1	Newsroom dilemma with a single firm: Monopoly	64
2.3.2	Newsroom dilemma with two firms: Duopoly	65
2.3.3	Competition may lead to better reporting	67
2.3.4	The role of audience's information	68
2.4	Stories and the effect of better initial information	70
2.5	Information dissemination by a source	73
2.6	Conclusion	78
2.7	Tables	80
2.8	Figures	81
2.9	Appendix	82
A	Proofs from the main text	82
B	Allowing for sitting on information	100
C	Generic prior on the type	105
D	Monopoly with public signal	112
Chapter 3	Optimal Innovation Time-off	113
3.1	Introduction	113
3.2	The general model	117
3.3	Two-period case: model and analysis	119
3.4	Conclusion	125
3.5	Figures	127
3.6	Appendix	128
A	Proofs from the main text	128
Bibliography		1

List of Tables

2.1	Deep (first row) and simple (second row) accountability reporting (as a % of total front-page stories in April) in a sample of 9 newspapers in the US for 1991-2011 in five-year gaps	80
-----	---	----

List of Figures

1.1	Summary of the timing of the game	29
1.2	The optimal belief threshold β_0^{FI} for the complete information about θ case	30
1.3	Comparing β_0^{NI} and β_0^{FI}	30
1.4	Uniqueness of babbling equilibria for priors $\beta_1 < \beta_1^{NI}$	31
1.5	Honest equilibria for different c ranges	31
1.6	Terminal belief possibilities in potential delayed equilibria	31
2.1	Equilibrium σ_D^* and σ_M^* for when $\pi = .6$, $v = .3$, $\bar{c} = 2$ and $\varepsilon = .1$	81
2.2	Equilibria in $\sigma_D - \sigma_M$ space and the behavior of the source	81
2.A.1	Proof of Lemma 2.2: Proving $\frac{\partial^2 A}{\partial \sigma^2} > 0$. Orange plane: $\frac{\partial^2 A}{\partial \sigma^2}$, blue plane: $0.\sigma + 0.\pi$ in $\pi - \sigma$ space.	90
2.A.2	Proof of Proposition 2.7: Proving LHS < RHS. Orange plane: RHS, blue plane: LHS and green plane: $0.\sigma + 0.\pi$ in the $\pi - \sigma$ space.	95
2.B.3	Proof of Proposition 2.10: Proving $\frac{\partial \bar{v}}{\partial \pi} > 0$. Orange plane: $\frac{\partial \bar{v}}{\partial \pi}$, blue plane: $0.\sigma + 0.\pi$ in the $\pi - \sigma$ space.	102
2.B.4	Proof of Proposition 2.10: Proving $\argmax_{\sigma} \bar{v}(\pi) = 1$. Orange plane: $\bar{v}(\pi)$, blue plane: $0.\sigma + 0.\pi$ in the $\pi - \sigma$ space for $\pi \in [0.5, 0.6]$	102
2.B.5	Proof of Proposition 2.10: Proving $c_D(\bar{v}(\sigma = 1)) > c_M$. Orange plane: $c_D(\bar{v}(\sigma = 1))$, blue plane: c_M in the $\pi - \sigma$ space for $\pi \in [0.5, 0.6]$	103
2.B.6	$c_D > c_M$ for $\pi = 0.75$ and $v = 0.7$. Orange line: c_M , blue line: c_D as a function of σ	104
2.B.7	$c_D(v = \bar{v} > c_M)$ for every combination of σ and π . Orange plane: $c_D(v = \bar{v})$, blue plane: c_M in the $\pi - \sigma$ space.	104
2.C.8	Proof of Proposition 2.12: Proving LHS > RHS in (2.9.C.25). Orange plane: LHS–RHS, blue plane: $0 * x + 0 * \pi$ in the $\pi - x$ space.	109
3.1	Optimal delegation mechanism for different c ranges	127

Acknowledgments

I am indebted to my supervisors, Motty Perry, and Robert Akerlof, for their continued support and guidance. They have trained me to think critically, isolate chains of thoughts, and formulate complex ideas in simple mathematical models. They have been a constant source of inspiration, motivation, and knowledge. I am also grateful to Dan Bernhardt, Costas Cavounidis, Kobi Glazer, Ilan Kremer, and Phil Reny for their time, useful advice and valuable inputs.

I gratefully acknowledge the financial and administrative support of the Department of Economics, University of Warwick and ESRC. Margaret Nash, Natalie Deven and Maryanne Heafey have always been kind and supportive.

I would also like to thank my teachers and professors who helped build a scientific temper in me through the years. Bhaskar Dutta, Omer Moav, Thijs Van Rens, Herakles Polemarchakis, Badal Mukhopadhyay, Namrata Gulati, Archana Jain, Bir Singh, Subodh Sharma, Ambika Gulati, Pushpa Tyagi, Sangeeta Gulati, Taruni Soin, V. Madhavi, Nilakshi Chatterji, Ina Chhabra, Nita Ganguly, Rekha Sharma, Anju Sharma, and Indrani Ghoshal are the unsung heroes of this thesis. Their passion and commitment pushed me to pursue my dreams.

A special word of gratitude goes to my friends and co-authors: Federico, Zeinab, Martina, Andy, Horng, Daniel, Michela, Kartik, and Laura. Each has contributed to my happiness in different ways, through the ups and downs of the Ph.D. life.

Finally, I am grateful to my family without whom none of this would have been possible. Thanks, Maa and Babba for always believing in me. Thank you, Akanksha for your unconditional love and support. You took care of everything selflessly all the while telling me how proud you are of me.

Declarations

This thesis is submitted to the University of Warwick in accordance with the requirements of the degree of Doctor of Philosophy. I declare that the thesis is my own and original work. Chapter 1 is joint work with Zeinab Aboutalebi, University of Warwick, and Chapters 2 and 3 are joint works with Federico Trombetta, University of Warwick. Both of them can attest to my significant contribution to these projects in terms of the original idea, modeling, model-solving, and writing. I also declare that any material contained in this thesis has not been submitted for a degree to any other university.

Ayush Pant

September 16, 2019

Abstract

This thesis consists of three essays on information economics. It broadly deals with understanding how and why can players be motivated to research. Researching helps us to achieve better outcomes and make better decisions by seeking better ideas and information. While this is a relevant question for organizations seeking to innovate, it is also applicable for societies in general in the current times. Chapter One deals with supervisor-agent relationships in organizations where supervisors give feedback to their employees on the ideas that they generate. Best implementation requires the supervisor to shoot down more mediocre ideas honestly. However, this potentially discourages the agent. The paper shows that supervisors are only honest with agents who have a high belief in their ability to succeed. Overconfidence is, therefore, potentially welfare improving. Chapter Two studies the effect of increased competition in the modern digital environment on the quality of reporting by media outlets. Two opposing forces determine how media outlets resolve the competing demands for speed vs. accuracy – preemption and reputation. The paper shows that more competitive environments may be more conducive to reputation building. Therefore, it is possible to have better reporting in a more competitive world. Finally, in Chapter Three, I return to the question of innovation in organizations. This paper solves for the optimal delegation mechanism that grants an agent the authority to take time off from his current task to pursue creative endeavors. Driven by a high intrinsic motivation, the agent would like to take time off for any idea. The principal, on the other hand, would like the agent only to pursue ideas that have a high potential to succeed. In the optimal mechanism, the principal is inefficiently harsh on the agent who gets time off. Creativity, therefore, only receives a limited opportunity.

Chapter 1

Feedback on Ideas

1.1 Introduction

Employees are often assigned tasks with two distinct phases: in the first phase, ideas are generated; in the second phase, the best idea is implemented. Furthermore, it is common for supervisors to give feedback to their employees in this process. For instance, a partner in a law firm supervises an associate developing a litigation strategy, a project manager in a technology firm supervises an engineer solving a bug in app development, and a senior designer in an architecture firm supervises a junior designer looking for a design solution. One can trace such examples of feedback and supervision outside of corporate organizations as well; for instance, a professor supervising her grad student in a university.

This paper studies the supervisor's problem. Supervisors face the following trade-off. On the one hand, honest feedback encourages employees to discard bad ideas. On the other hand, such feedback can be demoralizing and discourage both idea generation and effort implementation. We build a model to describe how this trade-off shapes the supervisor's feedback, the employee's effort, and the employee's trust in the supervisor.

We consider a supervisor-agent model with two phases: experimentation and implementation. In the experimentation phase, the agent sequentially generates ideas at a cost, receives feedback from the supervisor on her ideas, and selects an idea to implement. In the implementation phase, the agent decides how much effort to put into completing her chosen idea. The agent's ability is initially unknown, and the agent and supervisor share a common prior. Importantly, we assume the supervisor does not internalize the agent's cost of effort. This misalignment of preferences means that dishonesty is a possibility.¹

Ability plays a central role in our model. We assume a high-ability agent both generates

¹Note that if providing feedback is costly to the supervisor (such as time costs) this could realign the principal's and agent's interests, thereby restoring honesty. We show that the supervisor is more (less) honest when he is more (less) time constrained, and therefore less (more) willing to supervise.

and implements ideas better than a low-ability agent. As a result, the agent's self-opinion (prior belief about her ability) affects both the agent's decision regarding how much to experiment and her choice of implementation's effort. Both of these effects, in turn, impact the supervisor's feedback.

There are three key findings of our model. First, the supervisor never gives a low self-opinion agent honest feedback because doing so is demotivating: it discourages effort in both the experimentation and implementation phases. When negative feedback discourages further experimentation, the supervisor prefers to falsely encourage the agent to induce her to put a higher effort in implementation instead. Therefore, negative feedback is only forthcoming for a high self-opinion agent. Moreover, a high self-opinion agent, independent of her actual ability, is repeatedly informed about her bad ideas and can end up being "treated more harshly".

Second, receiving supervisor feedback magnifies performance differences between high and low self-opinion agents. Because high self-opinion agents receive honest feedback, they have confidence both in their ability and in the quality of their ideas, which leads to high effort. Low self-opinion agents, in contrast, lack confidence, which leads to low effort. Receiving more honest feedback with a higher self-opinion allows the agent not only to experiment more but also to exert an optimal effort in implementing her chosen idea. Such an opportunity might not be available to a slightly lower self-opinion agent because she does not receive honest feedback as often. As a result, she has lower confidence in her idea. Therefore she might end up exerting too much effort on a bad idea, and too little effort on a good idea.

Third, overconfidence can be welfare improving. The discontinuous change in the supervisor's feedback strategy as the agent incorrectly goes from a low self-opinion to a high self-opinion gives rise to this possibility. The cost of overconfidence in ability is that it leads to too much effort exertion. However, the benefit of overconfidence is that it can lead to honest feedback. This benefit may outweigh the cost.

Our results find support in The Sensitivity to Criticism Test from PsychTests which collected responses from more than 3,600 participants.² The study revealed that those who tended to be defensive about negative feedback had lower performance ratings and lower self-esteem. Moreover, managers were skeptical to give feedback to workers who get defensive. "If there was an esteem problem, both men and women seemed to block out the constructive part of the equation and only focus on the criticism", revealed a manager. This further meant that the manager would rather "develop the more (talented and) mature employee," instead of spending time counseling those who easily got defensive. These ideas further find support in

²<https://eu.usatoday.com/story/money/columnist/kay/2013/02/15/at-work-criticism-sensitivity/1921903/>

the situational leadership theory developed by Paul Hersey and Ken Blanchard in mid-1970s. According to Ken Blanchard, “Feedback is the breakfast of champions.”

Related Literature. Our paper relates to two distinct strands of literature: experimentation and dynamic communication games. Within experimentation, our work falls under models of motivating experimentation. Previous research has looked at how information can be optimally delivered to the agents arriving sequentially to experiment (such as [Kremer, Mansour and Perry \(2014\)](#) and [Che and Horner \(2015\)](#)) or at how information should be designed for a single agent to motivate her to experiment (such as [Renault, Solan and Vieille \(2017\)](#) and [Ely \(2017\)](#)).³ Among the two, our setting falls in the latter category. [Ely and Szydlowski \(2017\)](#), [Smolin \(2017\)](#) and [Ali \(2017\)](#) are the closest in this respect.⁴ In each of these papers, a principal must reveal information by balancing the positive effect of good news with the discouraging effect of no or bad news. Nonetheless, these papers do not address situations where ex-ante commitment to a disclosure rule is not possible. How the same tradeoff shapes the honesty in strategic feedback with no commitment is our point of departure from these papers. Thus, our model is one of communication rather than information design. To the best of our knowledge, we are the first to study such settings without commitment.

Another point of departure is how the agent responds to honest feedback. In our setting, the supervisor tries to motivate the agent to exert effort in both the experimentation and implementation phases. As a consequence, honest feedback can discourage the agent at two levels. The first is stopping experimentation too early, and the second is exerting low effort in implementation. Introducing this novel objective makes our setting unique in feedback and experimentation literature.

Some older papers like [Lizzeri, Meyer and Persico \(2002\)](#) and [Fuchs \(2007\)](#) have looked at feedback in dynamic settings without experimentation and show that often it is not optimal to provide feedback.⁵ [Orlov \(2013\)](#) considers a setting in which providing information to the agent might benefit the principal in the short-run but may lead to long term agency costs. There the principal designs an optimal information sharing rule along with a compensation scheme. [Boleslavsky and Lewis \(2016\)](#) also study dynamic settings with commitment in which the agent has new information every period. The principal makes sequential decisions, after which he observes a private signal of the state. These works consider the effect of feedback in settings with commitment but no experimentation. Our paper connects these two types of literature in a

³See [Hörner and Skrzypacz \(2016\)](#) for a survey on the recent advancements in experimentation and information design.

⁴Some other related papers have looked at settings in which a sender commits to dynamic information design to influence a receiver. See, for example, [Bizzotto, Rüdiger and Vigier \(2018\)](#).

⁵Both these papers are also concerned with the issue of dynamic moral hazard, and feedback plays an assistive role to contracting.

no-commitment setting.⁶

The other strand of literature related to our work is dynamic communication games. A few papers like [Aumann and Hart \(2003\)](#), [Krishna and Morgan \(2004\)](#), [Forges and Koessler \(2008\)](#) and [Goltsman, Hörner, Pavlov and Squintani \(2009\)](#) look at repeated communication with an action at the end. Our setup is different in that the receiver should decide after each round whether she wants to experiment again. [Golosov, Skreta, Tsyvinski and Wilson \(2014\)](#) and [Renault, Solan and Vieille \(2013\)](#) are closer in this sense. They look at situations where the receiver decides after every round of communication. However, neither has the above-stated feature of persuasion in two phases.

In this respect, our work relates to dynamic persuasion games. [Morris \(2001\)](#), [Honryo \(2018\)](#) and [Henry and Ottaviani \(2019\)](#) are a few papers that do not assume commitment by the sender of information. The seminal paper by [Morris \(2001\)](#) deals with a potentially biased advisor persuading a decision-maker to choose actions dynamically when reputation matters. [Honryo \(2018\)](#) and [Henry and Ottaviani \(2019\)](#), however, are closer to our setting. In these papers, a sender (entrepreneur or researcher) tries to persuade a receiver (venture capitalist or publisher) to take a favorable action by sequentially disclosing some verifiable or costly information. We instead have a tradeoff with cheap talk communication. In our model, when the supervisor persuades the agent to experiment again, he inadvertently also persuades her to exert lower effort in implementation. It is this feature that creates the main honesty/dishonesty tradeoff in our model.

Finally, our result on the importance of beliefs in final performance is related to some of the older research starting with [Bénabou and Tirole \(2002\)](#). This vast line of economics research is itself based on the original psychology research of [Bandura \(1977\)](#). However, such research usually looks at the importance of belief absent any external supervision. The presence of a supervisor drives our results on the effect of higher self-opinion and overconfidence.⁷

The rest of the paper is structured as follows. In [Section 1.2](#), we describe the basic model. In [Section 1.3](#), we solve two benchmark cases of the model without supervision, which help us build intuition and solve the complete game. Then, in [Section 1.4](#), we present the main analysis of the game with a supervisor without commitment. We move onto presenting how our results are qualitatively the same in a few extensions and offer new interpretations of our model in [Section 1.5](#). Finally, we conclude in [Section 1.6](#).

⁶[Orlov, Skrzypacz and Zryumov \(2018\)](#) is an exception. They look at commitment and no commitment case in a setting in which an agent tries to convince the principal to wait before exercising a real option. Again, however, their model does not have experimentation.

⁷[Koellinger, Minniti and Schade \(2007\)](#) and [Hirshleifer, Low and Teoh \(2012\)](#) are two papers that empirically show the importance of overconfidence in the context of innovation and creativity.

1.2 The model

We consider a setting in which an *agent* (she) works on a project and a *supervisor* (he) is responsible for providing feedback. The project involves two distinct stages that proceed sequentially. The first stage is *planning* or *experimenting with ideas*, and the second stage is *execution* or *implementation of a chosen idea*. The agent is responsible for both experimenting with and implementing ideas for the completion of the project. The supervisor has no commitment power or verifiable signals and provides cheap talk feedback based on what he observes.

Stage 1: Idea generation. The process of idea generation involves multiple rounds $t = 1, 2, \dots$. In each round t , the agent decides whether she wants to draw a new idea. The quality of an idea is determined by its *ex-ante potential to succeed* θ_t which could be either high (\hbar) or low (ℓ). The distribution of θ_t is given by

$$\theta_t = \begin{cases} \hbar & \text{with probability } \alpha, \\ \ell & \text{otherwise} \end{cases}$$

where α is the *ability* of the agent. $\alpha \in \{0, q\}$ where zero is “low”, and $q \in (0, 1)$ is “high”. Therefore, only a high-ability agent can come up with a high potential idea, which happens with probability q . The ability (unlike the idea) remains persistent throughout the play. The agent and the supervisor only know the distribution of the ability; neither observes it. The belief that the agent is high-ability at the beginning of round t is denoted by β_t , with a common prior $\beta_1 \in (0, 1)$ at the beginning of the game in round 1. For much of the text, we use belief and self-opinion interchangeably. We assume that the agent possesses a low potential outside option idea at the beginning in round 1 denoted by $\bar{\theta} = \ell$.

Actions and timing: In each round of experimentation the agent chooses $I_t \in \{0, 1\}$. $I_t = 0$ denotes the agent’s decision to stay in Stage 1 and experiment with another idea in round t , i.e., not implement. There is a cost c of experimentation. It could arise from searching the Internet, looking up for data, reading material, previous works, and seeking inspiration. The agent produces an idea θ after privately incurring c .

Importantly, we assume that only the supervisor can see the potential of the idea generated. The supervisor privately observes θ_t and chooses an announcement about its observed potential, $m_t \in \{\ell, \hbar\}$.⁸ We initially assume limited recall of the agent and the

⁸We can also start with an arbitrary message space M but since we consider a game of cheap talk with binary types and we focus on pure strategy equilibria, what matters are the equilibrium mappings from the supervisor type (what potential of the ideas he observes) to the message space, i.e. what is the meaning of the messages. Here, messages ℓ and \hbar have their natural meaning and are understood as the potential of the idea developed.

supervisor so that they only talk about the last idea produced (and not the entire history of past ideas). We present the analysis of perfect recall in which the supervisor is allowed to make backdated messages in Section 1.5.2.

Alternately, the agent could decide to implement the last idea after the supervisor's message. This is denoted by $I_t = 1$.

Stage 2: Idea implementation. If the agent decides to move to the idea implementation stage in $t + 1$ following the last message of the supervisor m_t , then her idea gets fixed at $\theta \equiv \theta_t$.

Actions and timing: The agent chooses effort $e \in [0, 1]$ at cost $\frac{e^2}{2}$ to complete the project. The final outcome of the project, success or failure, is determined by the following distribution function

$$\Pr(\text{success}) = \begin{cases} e & \text{if } \theta = h, \\ ke & \text{if } \theta = \ell \text{ and agent is high-ability, } k \in (0, 1), \\ 0 & \text{otherwise.} \end{cases}$$

The probability of success is a function of the potential of the chosen idea θ , effort exerted by the agent e and the ability of the agent α . It must be noted that only the high-ability agent is capable of successful completion of the project. Moreover, only she may obtain a success even with a low potential idea. Therefore, when the ability is unknown there is an incentive to implement a low potential idea instead of experimenting again.

We will make the following assumption for mathematical convenience.

$$q \geq (q + (1 - q)k)^2 \geq k \tag{A}$$

Intuitively, this assumption implies that in case the agent has a low potential idea, the supervisor finds it beneficial for the agent to experiment than to implement that idea (with the maximum possible effort of 1). Further, an additional round of experimentation with feedback is preferred to an additional round of experimentation without feedback. We explain these ideas further when presenting the main analysis in Section 1.4.⁹

Payoffs: Completion of the project yields V . If the completed project is successful, it yields a normalized value of 1, and zero otherwise. The payoff of the agent is given by

$$u_A = V - Tc - \frac{e^2}{2}$$

⁹This assumption helps simplify the proofs by providing sufficient conditions. In the absence of this assumption, all our proofs go through but will be belief dependent, which makes them less obvious and more cumbersome.

where T is the number of rounds for which the agent has experimented. The payoff of the supervisor is given by

$$u_S = V.$$

The payoffs highlight the incentive misalignment between the agent and the supervisor. While both players prefer success over failure, the agent alone bears the cost of experimentation and implementation.

Once the payoffs are realized, the game ends. A summary of the timing of the game is provided in Figure 1.1. We provide an alternate interpretation of the model and additional examples in Section 1.5.3.

We now turn to the analysis of the game. Before we describe the behaviour of a strategic supervisor, we describe the benchmark case in the following section without the supervisor. We then introduce the supervisor in Section 1.4 and search for honest equilibrium feedback strategies.

1.3 Benchmark: Single agent problem

In this section we look at a setting in which an agent works on the project without any supervision. This preliminary analysis helps us put bounds on the behavior of the agent and supervisor when they interact with each other. Two cases are possible – the agent does not observe the potential of her idea, or she does so perfectly.

1.3.1 No information (NI) about θ

If the agent does not observe the potential of her idea θ from attempting experimentation at belief β and there is no outside support, then the two alternatives available to her are as follows:

1. The agent can choose to not experiment and directly implement the project using the outside option idea. In this case, the agent $\max_e \beta k e - \frac{e^2}{2}$, which yields a maximized payoff of $\frac{(\beta k)^2}{2}$.
2. The agent can choose to experiment once and then execute the resulting idea. In this case, the agent $\max_e \beta(q + (1 - q)k)e - \frac{e^2}{2} - c$, which gives a maximized payoff of $\frac{\beta^2(q + (1 - q)k)^2}{2} - c$.

Observe that the agent does not want to try experimenting more than once in this setting because experimenting is an additional cost without any added benefit. She will not learn the quality of the new idea and the odds of coming up with a high potential idea remain unchanged. The only reason she might want to experiment once is to take the gamble of coming up with

a high potential idea. She will do so if her belief is high enough. This is illustrated in the following condition:

$$\overbrace{\frac{\beta^2(q + (1-q)k)^2}{2}}^{\text{expected benefit of experimentation}} \geq \underbrace{\frac{(\beta k)^2}{2}}_{\text{opportunity cost}} + \underbrace{c}_{\text{actual cost}}, \quad (\text{C1})$$

which leads to the following lemma:¹⁰

Lemma 1.1. *Let $c < \frac{(q+(1-q)k)^2 - k^2}{2}$. If there is no information about θ , there exists a unique threshold $\beta_0^{NI} := \left(\frac{2c}{(q+(1-q)k)^2 - k^2} \right)^{\frac{1}{2}}$ such that*

1. *if the prior belief $\beta_1 \geq \beta_0^{NI}$ then the agent experiments once before finishing the project by exerting effort $\beta_1(q + (1-q)k)$, and*
2. *if the prior belief $\beta_1 < \beta_0^{NI}$, the agent uses the outside option idea $\bar{\theta} = \ell$ to finish the project by exerting effort $\beta_1 k$.*

In the text we will also be interested in how β_0^{NI} responds to changes in the cost of experimentation c . It is easy to see that a higher cost of experimentation raises this threshold as it reduces the incentives to experiment *ceteris paribus* (see Appendix B for other comparative statics result).

1.3.2 Full information (FI) about θ

When the agent can perfectly observe the outcome of each round of experimentation, then she potentially wants to experiment at least once. This, as before, depends on her belief about her ability. But now she uses Bayes' rule sequentially to update her belief after observing the potential of the idea from the latest round of experimentation in a way that

$$\beta_t = \begin{cases} \frac{(1-q)\beta_{t-1}}{1-\beta_{t-1}q} & \text{if } \theta_{t-1} = \ell, \\ 1 & \text{otherwise.} \end{cases}$$

As is standard in good-news models, the agent revises her belief downwards each time she generates a low potential idea, but her belief jumps to 1 if she generates a high potential one. The agent enters the implementation phase and finishes the project upon observing $\theta_{t-1} = h$.

¹⁰A similar lemma with a belief threshold condition can also be obtained if the agent has no outside option idea. Denote such a cutoff by β_ϕ^{NI} . Then it can be shown that such a cutoff exists and is given by $\beta_\phi^{NI} = \frac{(2c)^{1/2}}{q+(1-q)k}$. Obviously, $\beta_\phi^{NI} < \beta_0^{NI}$. However, we make use of β_0^{NI} in the main analysis – we assume away the possibility of quitting when there is no support from a supervisor.

At this point, she does not have an incentive to experiment further as she only bears an additional cost without any extra benefit. She finalizes the project with the maximum effort of 1 which leads to the project being successful with certainty, and yields a maximized payoff of $\frac{1}{2}$ (the previous cost of experimentation is sunk). Thus, independent of which round of experimentation she is at if $\theta_{t-1} = \ell$ then $I_t^{FI}(\beta_t = 1) = 1$ is optimal with $e^{FI}(\beta_t = 1) = 1$.

On the other hand, after observing $\theta_{t-1} = \ell$ (with the agent observing low potential ideas $\theta_{t'} = \ell$ for all the previous rounds $t' < t - 1$ as well) the agent holds a belief $\beta_t < 1$ about her ability. The agent again faces two choices – to implement the low potential idea or to continue experimenting. If she chooses to implement her low potential idea then she chooses the optimal effort to $\max_e \beta_t k e - \frac{e^2}{2}$. This yields a maximized payoff of $\frac{(\beta_t k)^2}{2}$ where she exerts effort $\beta_t k$ according to her belief β_t . Depending on her belief β_t she might be a high-ability agent with a positive probability of success. If she chooses to experiment once more, then with probability $\beta_t q$ she comes up with a high potential idea and exerts maximal effort of 1 thereafter to finish the project (from above). With probability $1 - \beta_t q$ she comes up with a low potential idea and she faces the same decision problem but with a lower belief $\beta_{t+1} < \beta_t < 1$. Denote the value function of the agent at the beginning of round t with belief β_t when her last observed outcome is $\theta_{t-1} = \ell$ by $\mathcal{V}^\ell(\beta_t)$, such that

$$\mathcal{V}^\ell(\beta_t) = \max \left\{ \frac{(\beta_t k)^2}{2}, -c + \frac{\beta_t q}{2} + (1 - \beta_t q) \mathcal{V}^\ell(\beta_{t+1}) \right\}.$$

Assuming that the agent wants to start experimenting (the condition for which we will outline below), we are interested in if and when the agent stops experimenting with repeated low potential ideas. To do so, let the maximum number of rounds the agent experiments be T . The agent at belief $\beta_T \equiv \beta$ after $T - 1$ rounds will attempt another *final* round of experimentation knowing that irrespective of the outcome she will move to implementing her idea in the following round. So

$$\begin{aligned} \mathcal{V}^\ell(\beta) &= \max \left\{ \frac{(\beta k)^2}{2}, -c + \frac{\beta q}{2} + (1 - \beta q) \mathcal{V}^\ell(\beta') \right\} \\ &= -c + \frac{\beta q}{2} + (1 - \beta q) \mathcal{V}^\ell(\beta') \geq \frac{(\beta k)^2}{2} \end{aligned}$$

where

$$\beta' = \frac{(1 - q)\beta}{1 - \beta q} \text{ and } \mathcal{V}^\ell(\beta') = \frac{(\beta' k)^2}{2},$$

which can be rearranged to

$$\overbrace{\frac{\beta q}{2} + (1 - \beta q) \frac{(\beta' k)^2}{2}}^{\text{expected benefit of experimentation}} \geq \underbrace{\frac{(\beta k)^2}{2}}_{\text{opportunity cost}} + \underbrace{c}_{\text{actual cost}}. \quad (\text{C2})$$

Lemma 1.2 follows from condition (C2) and captures the optimal behaviour of the agent under full information about θ . (All proofs are presented in Appendix A.)

Lemma 1.2. *If there is full information about θ , the optimal decision rule of the agent I_t^{FI} is a unique belief threshold rule such that*

$$I_t^{FI} = \begin{cases} 0 & \text{if } \theta_{t-1} = \ell \text{ and } \beta_t \geq \beta_0^{FI}, \\ 1 & \text{otherwise.} \end{cases}$$

for $c < \frac{q(1-k^2)}{2}$. Further, the optimal effort that the agent exerts to implement her idea is given by

$$e^{FI} = \begin{cases} \beta_{T+1}k & \text{if } \theta_T = \ell, \\ 1 & \text{otherwise.} \end{cases}$$

When $c \geq \frac{q(1-k^2)}{2}$ the agent does not experiment for any belief, and implements her outside option idea with effort $\beta_1 k$.

Figure 1.2 plots the expected benefit from experimentation (LHS plotted in green) and the cost of experimentation (RHS plotted in red) from condition (C2) for different levels of beliefs β . It illustrates the uniqueness result of Lemma 1.2 under the cost condition $c < \frac{q(1-k^2)}{2}$. Note that both the benefit and the costs are declining in belief about ability. A lower belief in ability means that the agent is less likely to get a high potential idea, which reduces the expected benefit of experimentation. At the same time, for the same reason, it induces the agent to exert lower effort when implementing the outside option idea, thereby reducing the opportunity cost of experimentation. However, the fixed component c of the total costs of experimentation ensures that the costs never go down to zero, which in turn guarantees the existence of the unique threshold.

Observe that the optimal decision rule does not depend on t but only on the belief β , which is a function of the potential of the last observed idea. For a given set of parameters, the maximum number of rounds the agent experiments T is only defined by the prior belief β_1 . The agent wants to start experimenting with ideas if $\beta_1 \geq \beta_0^{FI}$, and goes on doing so with repeated

low potential ideas as long as the belief hits β_0^{FI} . T is therefore determined by how far β_1 is from β_0^{FI} .

It only remains to show how β_0^{FI} varies with a change in parameters. Again, we'll be interested in how β_0^{FI} responds to a change in the cost of experimentation. As expected, an increase in the cost of experimentation raises the threshold belief β_0^{FI} as the agent wants to experiment fewer rounds now (for any prior).

1.3.3 Comparing β_0^{NI} and β_0^{FI}

Lemma 1.3. *If $c < \frac{(q+(1-q)k)^2 - k^2}{2}$, then both β_0^{NI} and β_0^{FI} exist and are unique with $\beta_0^{NI} > \beta_0^{FI}$.*

Figure 1.3 illustrates why $\beta_0^{NI} > \beta_0^{FI}$. It shows that for any belief β the value of experimenting is always lower in the case when the agent has no information about her output of experimentation. Experimentation is merely a gamble to try luck without any learning. This makes the threshold for experimentation higher under the no information case.

1.3.4 An important definition

Before moving to the main analysis, we introduce some additional terminology that we will use extensively in the following sections.

Given the no information and the full information belief thresholds β_0^x for $x \in \{NI, FI\}$, define recursively a sequence of belief thresholds $\{\beta_i^x\}_{i=0}^\infty$ such that $0 < \beta_i^x < 1$ and $\beta_{i+1}^x = \frac{\beta_i^x}{1 - q(1 - \beta_i^x)}$. Starting with the threshold β_0^x the sequence identifies β_1^x , the belief that leads to β_0^x when the agent correctly finds out that her idea has a low potential to succeed, and so on. Therefore, β_{i+1}^x is the belief which when updated with the correct information about a low potential outcome leads to the belief β_i^x , and this is recursively defined all the way down to the belief β_0^x .

1.4 Strategic supervisor

1.4.1 Preliminaries

The game between a strategic supervisor and an agent in Stage 1 is one of dynamic cheap talk. The supervisor can costlessly send either of the two messages independent of the true potential of the idea. Our solution concept is (perfect) Bayesian Equilibrium.

To define the strategies of the agent and the supervisor at any time, we would need to define the history for each player when they are called upon to make a decision. Round t begins for the agent after having observed the last message sent by the supervisor m_{t-1} . Accordingly, a realized history for the agent includes the set of all previous messages sent by the supervisor

until and including the last message m_{t-1} and the sequence of past decisions made. Round t begins for the supervisor after observing the last idea of the agent θ_t . Accordingly, a realized history for the supervisor includes, in addition to the history viewed by the agent, the sequence of all the realized idea potential from the past experimentation.¹¹

For most of the paper, we focus on pure strategy equilibria and limited recall, i.e. we are interested in whether the supervisor is honest with the agent when he can only send a message about the last idea generated. A pure strategy for the supervisor in round t is a mapping from the realized history to the message space $\{\ell, \hbar\}$. The supervisor is honest with the agent if for any realization of the history the supervisor sends a message that matches the observed potential of the idea. If the supervisor reveals to the agent the outcome of her last experimentation in round t starting from a prior β_t the agent's updated posterior in round $t + 1$ is as in the full information case:

$$\beta_{t+1}^\ell = \frac{(1-q)\beta_t}{1-q\beta_t} \text{ if } m_t = \ell, \text{ and} \quad (1.1)$$

$$\beta_{t+\tau}^\hbar = 1 \text{ otherwise.} \quad (1.2)$$

If the supervisor uses the same message independent of the realized history the supervisor is said to lie or babble (see footnote 12). In this case the agent's posterior belief is the same as her prior belief. We will assume that when the supervisor is expected to lie the agent does not consult the supervisor. This rules out the possibility of the supervisor privately learning and not revealing to the agent the outcome, and the arising deviations.

Given our focus on pure strategies and that the two players share a common prior, the agent and the supervisor symmetrically update their belief on the agent's ability. If the agent stops experimenting (and implements her last idea) because the supervisor is babbling, neither the agent nor the supervisor have any new information. There is learning only insofar as the supervisor is honest.

¹¹Let $I^t := (I_1, \dots, I_t)$ and $m^t := (m_1, \dots, m_t)$ be the sequence of decisions made by the agent and the public messages given by the supervisor until round t . Define the set of histories for the agent and the supervisor at the beginning of round t by H_t^A and H_t^S respectively. The history for the agent at the beginning of round t is

$$h_t^A = (I^{t-1}, m^{t-1}) \in H_t^A \subset (\{0\}^{t-1} \times \{\ell, \hbar\}^{t-1}).$$

This is also the public history of the play of the game up to round t . In addition to the public history, the supervisor observes $\theta^t := (\theta_1, \dots, \theta_t)$ and an extra decision of the agent to experiment $I_t = 0$. The history for the supervisor at the beginning of round t is

$$h_t^S = (\theta^t, I_t, h_t^A) \in H_t^S \subset (\{\ell, \hbar\}^t \times \{0\}^t \times \{\ell, \hbar\}^{t-1}).$$

1.4.2 Analysis

What feedback strategy the supervisor employs will depend on how he expects the agent will respond to it, both in the experimentation phase and the implementation phase. We begin by discussing the obvious babbling equilibria. Babbling is always an equilibrium for any prior β_1 in the first stage of the game. The agent does not learn about the true potential of the last idea as the supervisor is always expected to send the same message. This is equivalent to the single agent decision-making problem without advice and Lemma 1.1 applies. Thus, the agent experiments once before finishing the project if $\beta_1 \geq \beta_0^{NI}$, otherwise she uses the outside option idea to finish the project. Neither supervisor type can profitably deviate from such an equilibrium given the beliefs. The supervisor sends meaningless messages, the agent correctly believes that there is no information content in the recommendations and she makes her decision only on the basis of her prior belief.¹²

In what follows we determine if there exist any pure strategy equilibria in which the supervisor is honest, and under what conditions. The approach will be to determine the existence for different ranges of beliefs starting with low ones.¹³

Proposition 1.1. *For any belief $\beta < \beta_0^{FI}$, any communication strategy is an equilibrium and none induces the agent to experiment.*

From Lemma 1.3, we know that $\beta_0^{FI} < \beta_0^{NI}$. The region of beliefs $\beta < \beta_0^{FI} < \beta_0^{NI}$ is the one in which the agent does not want to experiment with ideas independent of how much information is provided to her. So all communication strategies are equally informative to the agent and are an equilibrium. The agent does not consult the supervisor in any equilibria as she is very pessimistic about her ability to come up with a high potential idea. She does not want to bear the cost of experimentation at such low beliefs. She simply implements her low potential outside option idea $\bar{\theta} = \ell$ with an effort βk .

A concern when evaluating whether the supervisor can be honest for higher beliefs will

¹² When the supervisor babbles, it might be useful to think of babbling in mixed strategies rather than in pure strategies (see description of mixed strategies in Appendix A). A supervisor babbling in mixed strategies makes use of both the messages in equilibrium, and the posterior β_t after either message remains unchanged. There are also babbling equilibria in pure strategies. Say the agent conjectures that the supervisor only says $m = h$ on-the-equilibrium path. We have that $\Pr(m = h | \theta = h) = 1 - \Pr(m = \ell | \theta = \ell)$ and a potential babbling equilibrium. While there is no update of beliefs on path, the message $m = \ell$ is off path and we would need to specify beliefs in the information set following this message. Such an equilibrium is supported by any belief $\beta^{\text{offpath}} \in [0, \beta_1)$.

¹³ The proofs will be presented in terms of a generic belief β wherever possible. The intuition is the same – whether the agent starts out in the given range with a low potential outside option idea or whether she lands there after continued experimentation (and ending up with a low potential idea that she is aware of), if she finds herself there her behavior is the same. If she finds herself in any of the ranges with the knowledge that her idea was definitely a high potential idea, then she will always immediately implement her idea by exerting effort 1.

be what he thinks is the possibility of the agent experimenting again after a negative message. As in the the full information case outlined in Section 1.3.2, the agent experiences a decline in both the benefit and cost of coming up with a new idea after receiving a truthful negative messages. With continued discouragement the agent must stop experimenting at belief β_0^{FI} . However, the supervisor's payoff is contingent on the agent's success. This implies that the he faces a discontinuous drop in the benefit of being honest at β_0^{FI} , while the cost is that the agent exerts a lower effort in implementation. Our first main result and proposition builds on this intuition. It defines the range of beliefs for which the cost of being honest are higher than the benefits.

Proposition 1.2. *For any belief $\beta_0^{FI} \leq \beta < \beta_1^{NI}$, babbling is the unique equilibrium strategy.*

The intuition for this proposition is illustrated in steps using Figure 1.4.¹⁴

We begin by showing that babbling must be a unique equilibrium strategy of the supervisor in the range of priors $\beta_0^{FI} \leq \beta_1 < \beta_1^{FI}$ (see Step 1 of Figure 1.4). In this range of priors, a message about the idea being low potential if expected in equilibrium must lead to a posterior about ability $\beta_2^\ell < \beta_0^{FI}$. At this point the agent does not want to experiment any more (from Proposition 1.1). Moreover, after experimenting and learning that her idea had a low potential to succeed she reduces her effort when implementing the idea. As a result, the expected probability of success further reduces with the low potential idea. This leads the supervisor observing a low potential idea to deviate from honesty and always send a positive message instead.

A positive message is believed by the agent pushing up the posterior of the agent to 1. The agent best responds by implementing the chosen idea with the maximal effort of 1, which increases the expected probability of success with a low potential idea. The supervisor is at the very least able to extract a higher effort on a low potential idea by deviating. Thus, no equilibria in which the supervisor is honest will survive – babbling is unique in this range of priors. In such a babbling equilibrium, the agent best responds by not experimenting because this is identical to a situation with no supervisor and $\beta_1 < \beta_0^{NI}$ (from Lemma 1.1).

Now, in Step 2 consider the range of priors which when updated with negative messages lead to posteriors below β_1^{FI} . The same argument as the one highlighted above holds because such low posteriors lead the agent to implementing the low potential idea with a lower effort. This time because the supervisor is expected to babble if updated with an honest discouraging message. Therefore, an agent expecting information can be taken advantage of by supervisor

¹⁴Here we discuss the intuition of why honesty cannot be an equilibrium strategy but the proposition is stronger. The argument will also hold to prove that no informative equilibria will survive in this range of beliefs. Our proof in Appendix A presents a general proof that allows for mixed strategies as well.

type who has only observed low potential ideas. This kills honesty and only the babbling equilibria survive. The same logic can now be extended all the way up to all the prior beliefs which when updated with a discouraging message about the idea lead to posteriors below β_0^{NI} . Below β_0^{NI} the agent does not want to experiment when no information is provided by the supervisor. Such is the case for all prior beliefs $\beta_1 < \beta_1^{NI}$ (illustrated in Step 3).

The total communication breakdown between the supervisor and the agent in this range of beliefs is driven by the fear of the supervisor to discourage the agent to the point of no further experimentation. This is why we call this region of beliefs as those in which the agent has a *low self-opinion*. When he sees that the agent has produced a low potential idea the supervisor finds it beneficial to cajole the agent by calling it a high one, so that at the very least the agent exerts a high effort to implement a low potential idea. But lying is counter-productive as the agent expects the supervisor to only provide fake encouragement; neither does she consult the supervisor nor does she experiment.

This region of beliefs $\beta_0^{FI} \leq \beta < \beta_1^{NI}$ where the agent has a low self-opinion reflect pure inefficiencies in the supervisor-agent relationship. From Lemma 1.2 we know that the agent would continue experimenting with ideas until she produces a high potential idea for beliefs $\beta \geq \beta_0^{FI}$ if she receives honest feedback. At the same time, the supervisor is also (always) better off with repeated experimentation until a high potential idea is produced. But neither can achieve this better outcome because the supervisor is unable to commit to honestly revealing the result of the agent's experimentation. Even though the agent is willing to listen to honest feedback, her reaction to negative feedback is too extreme from the supervisor's point of view. If the agent must give up, he prefers she exert the maximum effort instead. Such inefficiency will be a feature of any communication equilibrium we can construct as babbling is unique. The supervisor cannot offer any information in equilibrium.

The extent of babbling and that of the resulting inefficiency is determined by the gap between β_0^{FI} and β_1^{NI} , which is a function of the parameters. An increase in the cost of experimentation (c) increases both these thresholds and causes babbling for even higher beliefs (and also no experimentation for higher beliefs). An increase in the probability of generating a high potential idea (q) reduces the region of babbling. An increase in the success rate from implementing a bad idea (k) can *decrease* the inefficiency by reducing the babbling region as it makes the agent want to experiment more without supervision by reducing β_0^{NI} .

Note, however, the difference in the agent's best response to such an uninformative strategy of the supervisor. Since the supervisor babbles in the entire region of beliefs below β_1^{NI} , from Lemma 1.1 the agent best responds by not experimenting in the region below β_0^{NI} .

and by experimenting once in the region between β_0^{NI} and β_1^{NI} . This produces an added source of inefficiency when she experiments in this region i.e. when the belief is above β_0^{NI} but below β_1^{NI} . In this case, the agent exerts an inefficient level of effort to implement the idea as she is unable to observe the potential of her idea without honest supervision. She exerts more effort on a low potential idea and a lower effort on a high potential idea.

We are now in a position to determine if there are any honest equilibria. The possibility of honesty opens up for beliefs $\beta > \beta_1^{NI}$ because the agent is now willing to experiment at least once without the supervisor's support. This happens in the region of beliefs between β_0^{NI} and β_1^{NI} . The previous threat point for the supervisor now potentially disappears as the supervisor can guarantee that the agent will experiment even when she is discouraged. In this sense, we call this the region of *high self-opinion*. We are now in a position to analyse whether this one extra round of experimentation (without the consultation of the supervisor) and a high self-opinion is sufficient for the supervisor to be honest.

Proposition 1.3. For $c \geq \frac{\kappa k - (\kappa k)^2}{2}$ where $\kappa \equiv \frac{k}{(q + (1-q)k)^2}$ and for all $t \geq 1$,

1. truth-telling is an equilibrium strategy for the supervisor for $\beta_t \geq \beta_1^{NI}$, and
2. babbling is the unique equilibrium strategy for the supervisor for $\beta_t < \beta_1^{NI}$.

The agent's equilibrium strategy is given by

$$I_t^* = \begin{cases} 0 & \text{if } m_{t-1} = \ell \text{ and } \beta_t \geq \beta_1^{NI}, \text{ or } \beta_0^{NI} \leq \beta_t < \beta_1^{NI}, \\ 1 & \text{otherwise.} \end{cases}$$

The agent's optimal effort is given by

$$e^* = \begin{cases} 1 & \text{if } m_{t-1} = h, \\ \beta_t(q + (1-q)k) & \text{otherwise.} \end{cases}$$

Proposition 1.3 identifies the necessary and sufficient condition for an honest equilibrium to arise in the *entire* region above babbling equilibria, i.e. one of high self-opinion. This is shown to be when the agent's cost of experimentation is sufficiently *high*. To see this, let us first look at the supervisor's incentives to be honest in the region of priors $\beta_1^{NI} \leq \beta_1 < \beta_2^{NI}$. Here the agent experiments once even when discouraged. At most the agent's belief can fall down to β_0^{NI} after a negative message. The supervisor is then willing to discourage the agent with a negative message only if he can ensure that even after discouragement the agent does not

reduce her effort significantly. In the absence of further supervision, he can only expect a higher expected probability of success if she exerts a high enough effort in implementation.

A supervisor who has observed a low potential idea expects the project to be successful with probability $(\beta_2^\ell(q + (1 - q)k))^2$ from being honest. After receiving a message $m_1 = \ell$, the agent correctly believes her current idea has a low potential to succeed and experiments once again but does not seek supervision because the supervisor is expected to babble. In this case, the agent then implements the next idea with effort $e = \beta_2^\ell(q + (1 - q)k)$. On the other hand, if such a supervisor deviates from honesty and announces $m_1 = h$, then he expects the probability of success to be $\beta_2^\ell k$. The agent incorrectly believes that her idea had a high potential to succeed and exerts effort of 1 in implementing a low potential idea. For such a conjectured strategy to be an equilibrium, we must have that

$$\begin{aligned} (\beta_2^\ell)^2(q + (1 - q)k)^2 &\geq \beta_2^\ell k \\ \implies \beta_1 &\geq \frac{k}{qk + (1 - q)(q + (1 - q)k)^2} := \beta^{\text{truth}} \end{aligned}$$

Thus, the supervisor requires agent's belief to be sufficiently high even after discouragement, which in turn requires the prior to be large enough. This ensures that the agent exerts a higher effort in implementing her idea of unknown potential. We call this truth-telling threshold on prior β^{truth} .

The truth-telling threshold β^{truth} is a conditional threshold. It identifies how high the prior should be such that the supervisor has an incentive to reveal the truth about the agent's negative outcome *if the agent experiments again without supervision following the negative message*. The supervisor does not directly care about the agent's cost of experimentation in so far as she attempts to experiment again with an idea. So β^{truth} does not depend on c .

Now all we need to do is identify whether the range of priors we are considering delivers honesty by the supervisor, that is we are interested in if $\beta^{\text{truth}} < \beta_2^{NI}$. Specifically, if $\beta^{\text{truth}} \leq \beta_1^{NI}$ then truth-telling is an equilibrium for the full range of beliefs above β_1^{NI} and up to β_2^{NI} . If this condition is satisfied, the supervisor has an incentive to be honest because the prior is sufficiently high given the parameters. As outlined above, β^{truth} does not depend on the cost of experimentation c while β_1^{NI} does. The one free parameter can be used to determine if truth-telling is an equilibrium. The condition $\beta^{\text{truth}} \leq \beta_1^{NI}$ can then be rearranged to

$$c \geq \frac{\kappa k - (\kappa k)^2}{2} \text{ where } \kappa \equiv \frac{k}{(q + (1 - q)k)^2} < 1.$$

Intuitively, a lower bound on the cost of experimentation ensures that the agent's

no information thresholds β_0^{NI} and β_1^{NI} are high enough. Thus, when the agent decides to experiment and consult the supervisor her belief in her ability is already high. The supervisor can then be content with revealing the truth about low potential ideas to the agent. Discouragement does not lead to quitting with low effort; the agent still experiments once more and does so by exerting a sufficiently high effort. While the conditional truth-telling threshold β^{truth} is not a function of the cost of experimentation c , whether truth-telling is an equilibrium depends on it. An increase in the cost of experimentation raises the threshold β_0^{NI} (increasing the region of babbling) but has no effect on β^{truth} , making it easier to satisfy the condition $\beta^{truth} \leq \beta_1^{NI}$ and ensuring truth-telling above β_1^{NI} .

We are now only left with determining why if the supervisor is honest in the range of beliefs $\beta_1^{NI} \leq \beta_1 < \beta_2^{NI}$, then he should be honest in the range of beliefs above β_2^{NI} . For expositional convenience start now with the range of beliefs $\beta_2^{NI} \leq \beta_1 < \beta_3^{NI}$ when it is an equilibrium for the supervisor to be honest in the next lower range of beliefs. Consider whether a conjectured strategy of honesty is an equilibrium for the supervisor. A supervisor who observes a low potential idea can induce another two rounds of experimentation by being honest at this stage, one with supervision and one without. If, however, he deviates he induces the agent to exert maximal effort in a low potential task. Under assumption (A), the payoff from being honest are strictly higher than that from deviating as it is evaluated relative to his private updated belief β_2^ℓ . The same line of reasoning can then be extended to any belief above β_3^{NI} as well so that the supervisor always prefers honestly discouraging the agent and getting her to experiment more often than making her implement a low potential idea.

What happens when $c < \frac{\kappa k - (\kappa k)^2}{2}$? The following corollary identifies the honest equilibrium.

Corollary 1.1. *When $c < \frac{\kappa k - (\kappa k)^2}{2}$, $\beta_j^{NI} \leq \beta^{truth} < \beta_{j+1}^{NI}$ exists such that for all $t > 1$ for $j \geq 1$*

1. *truth-telling is an equilibrium strategy for the supervisor for $\beta_t \geq \beta^{truth}$, and*
2. *babbling is an equilibrium strategy for the supervisor for $\beta_t < \beta^{truth}$.*

The agent's equilibrium strategy is given by

$$I_t^* = \begin{cases} 0 & \text{if } m_{t-1} = \ell \text{ and } \beta_t \geq \beta^{truth}, \text{ or } \beta_{j-1}^{NI} \leq \beta_t < \beta_j^{NI}, \\ 1 & \text{otherwise.} \end{cases}$$

The agent's optimal effort is given by

$$e^* = \begin{cases} 1 & \text{if } m_{t-1} = \ell, \\ \beta_t(q + (1 - q)k) & \text{otherwise.} \end{cases}$$

In this case, $\beta^{\text{truth}} > \beta_1^{NI}$ and can lie between any β_j^{NI} and β_{j+1}^{NI} . We can then again construct an honest equilibrium above β^{truth} and a babbling one below. That all of these beliefs are above β_0^{NI} ensures that the agent experiments once more when a low potential idea is revealed to her in the presence of future babbling and makes such a strategy an equilibrium. The two cases discussed here are depicted in Figure 1.5.

It is worth emphasizing at this stage the key intuition driving the results in Propositions 1.2 and 1.3. What action the agent chooses depends on whether she thinks she is capable of drawing a better idea, and the expected strategy of the supervisor. If the agent has produced a low potential idea, the supervisor needs to incorporate the downwards effect that his negative message has on the belief about her ability. A lower belief discourages the agent at two levels. First is the discouragement to experiment, i.e., stopping experimentation too early. Second is the discouragement to implement, i.e., exerting low effort in implementing the idea. The second effect always exists. However, the low self-opinion arises when it is also matched by the first effect. On the contrary, the possibility of a high self-opinion phase arises when the first effect is not present.

We conclude this section by presenting an important corollary and our second main result.

Corollary 1.2. *The expected performance of the agent is better under a higher self-opinion.*

To see this, first note that the supervisor induces a weakly higher number of rounds of experimentation under a prior $\beta' > \beta$. If $\beta_j^{NI} \leq \beta < \beta_{j+1}^{NI}$, then either $\beta_j^{NI} \leq \beta < \beta' < \beta_{j+1}^{NI}$ or $\beta' > \beta_{j+1}^{NI}$. In the former case, the agent experiments an equal number of rounds under the two beliefs. However, in the latter case, the agent experiments more often under belief β' than under β . The reason is that it is easier to support the mutual expectation of honesty and repeated experimentation under a higher belief so that the agent experiments weakly more often under β' .

However, this has consequences on the agent's overall performance. Honest feedback by the supervisor allows the agent to match her effort more closely to the actual potential maximizing the probability of success. If the agent abandons seeking supervision (and experiments one final round) in the k th round under belief β , then she should still be seeking honest supervision

in the round k under belief β' . While the agent with belief β exerts an inefficiently low amount of effort in a high potential idea in round k , an agent with β' will exert the efficient level of effort of 1. An inefficient level of effort reduces the probability of success in a high potential idea.

Finally, if the idea in round k is a low potential one, then an agent with lower belief exerts an inefficiently high level of effort in its implementation while the agent with a higher belief experiments again. Therefore, there is a magnifying effect of a higher belief that results from the combined effect of better experimentation and better implementation. Conditional on being high-ability, an a priori better agent who has a higher belief in her ability does better in expectation.

It is also worth noting that in this context an a priori better agent (who has a higher self-opinion) will face more “criticism” from the supervisor for the same reason. An agent with a higher belief in her ability receives discouraging messages more often conditional on producing the same number of low potential ideas. However, the agent’s incentive to experiment more often arises precisely out of the supervisor offering honest criticism. In equilibrium, an agent with a higher belief expects to receive honest feedback more often and is therefore willing to experiment more often. In return, the supervisor expecting more experimentation offers more honest feedback to the agent. When the agent’s belief is lower, he fears to discourage the agent with negative messages. In this sense, an agent with a higher belief is more receptive to criticism, and that increases her chances of being successful.

1.4.3 Welfare analysis

The previous result (Corollary 1.2) only talks about the benefit of a higher self-opinion. However, the agent also pays a higher cost under a higher self-opinion owing to the aforementioned magnifying effect. This particularly hurts a low-ability agent who only pays a higher costs of experimentation and/or implementation under a higher belief.

The first part of this section shows that the above is not a concern even when evaluating the agent’s welfare under a higher self-opinion. We show, through a series of lemmas below that the ex-ante expected utility of the agent is always higher under a higher belief.¹⁵ The reason is that under a higher belief the agent places a greater ex-ante weight on being high-ability and believes that she is less likely to find herself in the worst situation.

The second part of the section then analyzes if holding an incorrect higher belief could also be welfare improving. Surprisingly, we show that this is possible. The reason is

¹⁵The supervisor is always better off with a higher self-opinion agent because in expectation such an agent performs better. At the same time, the supervisor doesn’t have to bear any costs.

the discontinuous change in the supervisor's feedback strategy as he goes from babbling to honesty.¹⁶

Welfare effect of a correct increase in self-opinion

Lemma 1.4. *Any increase in the prior from β to β' within the region of beliefs $\beta_0^{FI} \leq \beta < \beta' < \beta_0^{NI}$, $\beta_0^{NI} \leq \beta < \beta' < \beta_1^{NI}$, and $\beta_j^{NI} \leq \beta < \beta' < \beta_{j+1}^{NI}$ for $j > 1$ is welfare improving for the agent.*

This lemma relates to increasing the beliefs of the agent in such a way that only the cost of exerting effort increases in the eventuality that the project is implemented with a low potential idea or after not seeking supervision. In such a situation, welfare may increase on account of better implementation (because of higher effort) but may reduce on account higher costs of implementing.

Lemma 1.5. *An increase in the prior from the region $\beta_0^{FI} \leq \beta < \beta_0^{NI}$ to the region $\beta_0^{NI} \leq \beta' < \beta_1^{NI}$ is welfare improving for the agent.*

When the belief increases in such a manner, the agent is expected to conduct a costly round of experimentation which she did not earlier. Moreover, she is not expected to receive any feedback in this round. At the same time, her optimal effort choice increases unambiguously which is both more costly and more beneficial in expectation. From Lemma 1.4, we know that increasing the effort is always welfare improving when the belief increases. In addition, the increase in belief also makes it worthwhile to conduct experimentation without supervision from Lemma 1.1. This leads to an overall increase in welfare.

Lemma 1.6. *Let $2c < q(1 - (q + (1 - q)k)^2)$. An increase in the prior from $\beta = \beta_{j+1}^{NI} - \epsilon$ to $\beta' = \beta_{j+1}^{NI}$ is welfare improving for the agent.*

Finally, this lemma establishes that just pushing up the belief from an arbitrary region $\beta_j^{NI} \leq \beta < \beta_{j+1}^{NI}$ to the next region $\beta_{j+1}^{NI} \leq \beta' < \beta_{j+2}^{NI}$ is welfare improving. In doing so, the agent is expected to pay not only an additional cost of experimentation c but also that of some minimal increase in effort cost in the event of implementing without supervision.

Proposition 1.4. *Let $2c < q(1 - (q + (1 - q)k)^2)$. An increase in the prior from β to β' is welfare improving for the agent.*

The above proposition combines the information from the three lemmas and concludes that any increase in prior is welfare improving. This highlights the importance of agent's

¹⁶We prove all the statements here assuming that $c \geq \frac{\kappa k - (\kappa k)^2}{2}$ or that the truth-telling threshold $\beta^{\text{truth}} \leq \beta_1^{NI}$. However, this is not required as the proofs go through with a higher β^{truth} as well.

self-opinion – the agent’s confidence in her ability is critical for the overall success of the project.

Welfare effect of overconfidence

Still more interesting is to explain the effect of overconfidence in our environment. To introduce the notion of overconfidence, consider the following. Let the agent and the supervisor hold a common prior belief β about the agent’s ability when the true belief is b .

Definition 1.1. *The agent and the supervisor are overconfident about the agent’s ability if $\beta > b$.*

Under the above definition of overconfidence, we prove the following proposition:

Proposition 1.5. *Overconfidence is sometimes, but not always, welfare improving.*

To understand the intuition, consider the welfare of the agent when the correct belief is $b = \beta_1^{NI} - \epsilon$ but the common prior is β_1^{NI} . In such a situation, her overconfidence will drive her to experiment once with a round of honest feedback by the supervisor (and then potentially once more without any feedback). This would not have been possible under the true belief wherein she would have simply experimented without any feedback. However, the discontinuous benefit that arises from the change in supervisor’s feedback strategy at a higher belief (i.e. receiving honest feedback) outweighs the additional cost that the agent pays for an additional round of experimentation.

In fact, she is able to reduce her inefficient cost of implementation when the supervisor honestly reveals that her idea was a low potential one under the overconfident belief. To see this note that under the true belief she would exert $(\beta_1^{NI} - \epsilon)(q + (1 - q)k)$. Whereas under the overconfident belief she would exert $\beta_0^{NI}(q + (1 - q)k)$. Thus, overconfidence (and holding an incorrect self-opinion) can be welfare improving.

However, the above argument relies on the discontinuous change in behavior of the supervisor at the threshold. It then follows that when the supervisor’s behavior does not change, there might not be a benefit of being overconfident. To illustrate this, we show that overconfidence is welfare reducing when the common prior is β_0^{NI} but the true belief is any $b < \beta_0^{NI}$. In such a situation, holding the incorrect belief only adds to an added cost of experimentation and implementation without any corresponding benefit. Contrasting this with Lemma 1.6, it is immediate to see that overconfidence is different from a correct increase in belief.

1.5 Extensions

1.5.1 Benevolent supervisor and time-constrained players

We start out by discussing what happens when the supervisor also bears the cost of experimentation and implementation. In some situations, it is possible that a benevolent supervisor partially internalizes the costs borne by the agent. Such internalization may arise from the expert's (i.e. the supervisor's) prior experience from when he as an apprentice (agent), or simply because he works on the project with the agent.

For the two players $i \in \{A, S\}$, agent (A) and supervisor (S), let the cost of experimentation be c_i and the cost of implementation be $\frac{\phi_i e^2}{2}$. The difference between these costs for the two players captures any preference conflict between them. In so far as $c_S < c_A$ and $\phi_S < \phi_A = 1$, the preference conflict persists. For a given $(c_S, \phi_S) > 0$, there will be a “full information” threshold for the supervisor as well. Call this threshold β_{S0}^{FI} . This reflects the preferences of the supervisor and determines what are the maximum number of rounds the supervisor desires the agent to experiment (or the belief threshold equivalently) with full information about the potential of the ideas.

In the limiting case of $c_S = \phi_S = 0$ studied in the main text, this threshold did not exist – the supervisor wanted the agent to continue experimenting with complete information until she ended up with a high potential idea. However, when $c_S < c_A$ and $\phi_S < \phi_A$, we have $\beta_{S0}^{FI} < \beta_{A0}^{FI}$ so that the supervisor would still like the agent to experiment more than she would like. In this case, all our results from the main text go through as the fear of discouragement and the agent abandoning experimentation still persists.

One possible interpretation of such a situation are time-constrained players. To keep things simple, let $\phi_S = \phi_A = 1$ so that the supervisor fully internalizes the time cost of implementing to the agent. Now let c_S denote the time cost that the supervisor pays for providing feedback to the agent. This could happen when the supervisor has some alternate tasks to perform or requires time to understand the true potential of the agent's ideas. The following proposition follows from our discussion.

Proposition 1.6. *Let $\phi_S = \phi_A = 1$.*

1. *If $c_S < c_A$ then Propositions 1.1, 1.2 and 1.3 capture the optimal strategies of the agent and the supervisor.*
2. *If $c_S \geq c_A$ then the supervisor offers honest feedback until he reaches the belief β_{S0}^{FI} and the agent experiments with ideas till that point absent a high potential idea.*

The intuition is as follows. When the supervisor is time-constrained, he cares both about the success *and* about costly supervision from the agent experimenting in pursuit of

success. In turn, this eliminates the fear of discouragement. Notably, now it is more costly for the supervisor to keep offering feedback beyond a point over letting the agent implement a low potential idea. We can then get honest equilibria for some additional ranges of beliefs. Thus, a more time-constrained supervisor can potentially offer more honest feedback. The next corollary identifies the condition that makes this possible.

Corollary 1.3. *Let $\phi_S = \phi_A = 1$. If $c_S \geq c_A$ such that $\beta_{S0}^{FI} < \beta_1^{NI}$ then the region of beliefs where honest equilibria exist is larger in the case of $c_S \geq c_A$ than $c_S < c_A$.*

Observe that in the case of $c_S < c_A$ honest equilibria exist in the region of beliefs above β_1^{NI} (depending on c_A). But from the above proposition, honest equilibria in the case of $c_S \geq c_A$ exist starting from β_{S0}^{FI} . Thus, the latter case provides the possibility of more honesty if $\beta_{S0}^{FI} < \beta_1^{NI}$. However, since there is no closed form solution of β_{S0}^{FI} , it is not straightforward to translate this into a condition with only the costs.

Finally, note that if the supervisor does not internalize the cost of exerting effort, there is no benefit (in terms of more honest equilibria) of even partially internalizing the costs of experimentation.

Proposition 1.7. *If $\phi_S = 0$, then the equilibrium strategies are given by Propositions 1.1, 1.2 and 1.3.*

To understand the intuition, let $c_S = c_A$ and consider whether honesty is an equilibrium strategy for $\beta_{A0}^{FI} \leq \beta < \beta_{A1}^{FI}$ (after all, if the supervisor internalizes the full cost of experimentation then the belief thresholds should match). At this belief, if the supervisor is expected to be honest, then following a negative message the agent abandons experimentation and exerts a low effort level on the idea. If instead, she receives a positive message, she exerts 1 on her idea. Now, for a supervisor who has seen a low potential idea and does not internalize the cost of implementation, there is a strictly positive deviation to giving a positive message. This breaks down the honest equilibrium (and the existence of β_{S0}^{FI}).¹⁷

The issue arises here because the supervisor wants the agent to exert the maximal effort independent of the potential of the idea produced. The supervisor fears discouragement leading to lower effort in implementation which precludes honesty.

1.5.2 Perfect recall of previous ideas

Here we describe what happens if the agent and the supervisor have perfect recall of all the previous ideas. In such a situation, the supervisor can potentially make announcements

¹⁷It is possible to derive a belief threshold above which the supervisor is expected to be honest in equilibrium for a generic ϕ_S and given c_S and c_A . This is necessarily different from β_{S0}^{FI} because that is contingent on the equilibrium best response of the agent to the supervisor's strategy.

about each of the previous ideas after each round of experimentation. Given our attention to pure strategies, there are two kinds of honest and informative strategies that a supervisor may employ: immediate honesty and delayed honesty.

In the immediately honest strategy, the supervisor reveals to the agent the outcome of her experimentation immediately after she experiments. This is implicitly what we assumed all throughout Section 1.4. In a strategy of delayed honesty, the supervisor provides uninformative messages for certain rounds and then reveals honestly some or all the previous outcomes. Observe that a variety of delayed honesty strategies are possible – the supervisor may babble for any arbitrary number of rounds and then provide information for any arbitrary number of those rounds, and this may change over time. If the supervisor reveals to the agent the $\tau' \leq \tau$ outcomes of her experimentation after τ rounds starting from a prior β_t the agent's updated posterior in round $t + \tau$ is

$$\beta_{t+\tau}^\ell = \frac{(1-q)^{\tau'} \beta_t}{1 - q \beta_t \sum_{s=0}^{\tau'-1} (1-q)^s} \text{ if } m_t = \ell \text{ for all } \tau' \text{ ideas, and} \quad (1.3)$$

$$\beta_{t+\tau}^h = 1 \text{ otherwise.} \quad (1.4)$$

The case of $\tau = \tau' = 1$ corresponds to immediate honesty where the agent expects the supervisor to reveal the outcome of the experimentation immediately after each round of experimentation. All other cases fall under delayed honesty.

In case the supervisor is expected to babble, the agent's posterior belief is the same as her prior belief. We will assume that when the supervisor is expected to lie about an idea the agent does not consult the supervisor regarding that idea. This rules out the possibility of the supervisor privately learning and not revealing to the agent the outcome, and the arising deviations.¹⁸

Note first that the result of Proposition 1.1 remains unaltered. If the agent does not want to experiment with an immediately honest strategy, she does not want to experiment with a delayed honesty strategy. By experimenting when the supervisor is expected to reveal the outcomes after a delay, the agent only bears a higher cost of experimentation to receive feedback when she is almost convinced that she cannot produce a high potential idea. Thus, implementing the outside option is the best response of the agent, and all strategies of information revelation are an equilibrium.

¹⁸A formal definition of strategies in this case is complicated. But it is easy to describe what a strategy for the two players are in words. A strategy for the supervisor when the agent consults him in round t is a mapping from all the ideas she observes to the set of messages, one for each round of experimentation. A strategy for the agent in round t is a mapping from the observed messages to a decision to experiment again or implement. If she decides to implement, she must also decide which idea to implement given the message history.

Corollary 1.4. *Under perfect recall of ideas, for any belief $\beta_0^{FI} \leq \beta < \beta_1^{NI}$, babbling is the unique equilibrium strategy.*

The result of babbling being a unique equilibrium in the region of beliefs $\beta_0^{FI} \leq \beta < \beta_1^{NI}$ even under perfect recall follows almost directly from Proposition 1.2. To illustrate this point, start again with a prior belief $\beta_0^{FI} \leq \beta_1 < \beta_1^{FI}$. In the absence of commitment, a supervisor who observes only low potential ideas from all the experimentation rounds (after delaying) is tempted to deviate and call any arbitrary idea a high potential one. This is for the same reason as before – when such a message is believed, the agent exerts maximal effort on such an idea assuming it is a high potential one. The supervisor gains from such a deviation because he increases the effort of the agent on a low potential idea in the absence of more experimentation. As a result, babbling is the unique equilibrium and the agent best responds by implementing the low potential outside option idea. The same reasoning can then be extended to all the beliefs which when updated with a negative message lead to the agent abandoning experimentation (as the supervisor is going to babble in the following round). This happens all the way up to the belief β_1^{NI} as before.

For beliefs above β_1^{NI} , we have already identified the condition for *immediate honesty* to arise in Proposition 1.3. It is, however, possible to have other equilibria with some delayed honesty. We identify here a critical feature of such equilibria (if they exist) that allows us to compare it with the immediately honest equilibrium.

Observation 1.1. *In a delayed equilibrium, the supervisor can only induce as many rounds of experimentation as the ones for which he provides honest feedback eventually.*

The above observation merely states that if the supervisor never provides feedback on some rounds of experimentation that the agent performs, then the agent has no incentive to experiment. Since the agent never consults the supervisor for rounds in which he is expected to babble, there is no benefit to the agent from experimenting these extra rounds. This allows us to focus attention on those strategies in which the outcome of all the rounds of experimentation is eventually revealed.

Proposition 1.8. *The number of rounds of experimentation that an equilibrium strategy of delayed honesty induces can be no more than that induced by the equilibrium immediate honesty strategy.*

What matters when evaluating the supervisor's incentive to be honest at the time of final revelation is the belief from truthfully announcing that all the ideas produced are low potential.

Say that the belief after such a revelation at round τ is β_τ^ℓ . This belief can be in one of the following three ranges: $\beta_\tau^\ell \geq \beta_1^{NI}$, $\beta_0^{NI} \leq \beta_\tau^\ell < \beta_1^{NI}$ or $\beta_\tau^\ell < \beta_0^{NI}$ (See Figure 1.6).

Observe that a terminal belief in the first and second range can also be attained by an immediately honest strategy, which is also an equilibrium. For any prior β_1 , for the agent to experiment more rounds than what she does under immediately honest strategy her terminal belief after all the revelations should fall in the third case, i.e. $\beta_\tau^\ell < \beta_0^{NI}$. However, we argue that such a strategy cannot be an equilibrium. This is for the same reason as before – a supervisor who has only observed low potential ideas will prefer to deviate and claim any one of the ideas to be of high potential than inducing the agent to stop experimenting with a lower belief where the supervisor only babbles. Thus, equilibrium experimentation possibilities under perfect recall can be no more than those under limited recall.

1.5.3 Alternate interpretations

Our model more generally speaks to the following type of settings. An informed sender of information (supervisor) communicates with a less informed receiver of information (agent) who needs to take a costly action dynamically. Consider, for instance, an entrepreneur who works on a project experimenting with ideas, *privately observing their potential*, and implementing one of them. However, she relies on the finances of a venture capitalist (VC) who pays for such experimentation and implementation. While the entrepreneur would prefer to continue experimenting until she receives a high potential idea, the VC would like to cut funding for experimentation when he is sufficiently pessimistic.

In such a setting, the entrepreneur is the supervisor, while the VC is the agent.¹⁹ Costs c and $e^2/2$ are the money promised by the agent to the supervisor for experimenting with and implementing ideas. Let $\alpha \in \{0, q\}$ be the state of the project which is determined ex-ante and remains persistent but potentially unknown to both the parties. $\theta \in \{\ell, h\}$ denotes the potential of the idea produced by the entrepreneur. The VC decides in each period, whether to fund experimentation for one extra round or force the entrepreneur to implement the last idea.

We then provide answers to the following questions: When can the entrepreneur credibly release information? How many chances of experimentation can the entrepreneur extract from the VC with her revelation strategy? Notably, our inefficiency result shows that even though the VC would like to continue financing the entrepreneur's experimentation and the entrepreneur would like to continue experimenting, she calls off the project too early. However, there are benefits to be had from the VC both correctly and incorrectly believing that the project is good.

¹⁹Which player is the agent and which one is the supervisor is *not* determined by who is experimenting and implementing, but by who holds the information and who pays for the action.

1.6 Conclusion

In this paper, we showed how an employee responds to criticism influences whether she receives feedback or not. Supervisors may not provide honest feedback to employees who do not believe in their ability. In turn, this hurts their performance and potentially their future careers. Moreover, it also hurts organizations as the supervisors provide inefficiently low levels of honest feedback. In this sense, organizations should seek to hire employees that *believe* in their ability to succeed. In fact, our model shows that overconfidence can sometimes be welfare-improving.

Our results are based on a model of feedback provision in an agent-supervisor environment. The agent experiments with ideas to try to solve a problem at hand and a supervisor offers feedback on whether her ideas have the potential to be successful. We showed the results for when the supervisor has no commitment power and uses cheap talk messages to communicate with the agent. We identified the region of beliefs for which the supervisor could only uniquely babble in equilibrium leading to inefficiency in the relationship. Driven by the fear of discouraging the agent to the point of abandonment of experimentation, the supervisor is not able to offer any credible information to the agent. We then showed if there are possible equilibria in which the supervisor can honestly communicate his information to the agent. A necessary and sufficient condition for honesty above the babbling threshold was found to be the costs of experimentation being sufficiently high.

However, our analysis focused only on pure strategy equilibria. The problem involving mixed strategies is a complicated one that requires determining how the agent responds to the current message when, in the future, there can be more mixing. Our work shows the further scope of looking at mixed communication strategies in such dynamic environments in the absence of commitment. One may also think of introducing new complications in the model such as those involving different priors of the agent and the supervisor.

1.7 Figures




Figure 1.1: Summary of the timing of the game




Figure 1.2: The optimal belief threshold β_0^{FI} for the complete information about θ case






Figure 1.3: Comparing β_0^{NI} and β_0^{FI}

Step 1: Babbling is unique for $\beta_0^{FI} \leq \beta_1 < \beta_1^{FI}$



Step 2: Babbling is unique for $\beta_1^{FI} \leq \beta_1 < \beta_0^{NI}$




Step 3: Babbling is unique for $\beta_0^{NI} \leq \beta_1 < \beta_1^{NI}$



Figure 1.4: Uniqueness of babbling equilibria for priors $\beta_1 < \beta_1^{NI}$

1. Honest equilibria when $c \geq \frac{\kappa k - (\kappa k)^2}{2}$



2. Honest equilibria when $c < \frac{\kappa k - (\kappa k)^2}{2}$




Figure 1.5: Honest equilibria for different c ranges




Figure 1.6: Terminal belief possibilities in potential delayed equilibria

1.8 Appendix

A Proofs from the main text

We present general proofs in mixed strategies, wherever we can. The first section provides some new mathematical notation for this purpose.

Mathematical notation for mixed strategies

We focus attention on limited recall of previous ideas so that when the agent experiments one more round, she does not recall the previous ideas she has worked on. As a result, the supervisor does not need to make back dated messages about all the previous ideas. A strategy for the agent ρ_t in round t is a mapping from the last observed message to a possible mixed decision to continue experimenting with ideas or implementing the last one. We let

$$\rho_t^{m_{t-1}} = \Pr(I_t = 1 \mid m_{t-1})$$

be the probability that the agent decides to implement the project following the last message.

Similarly, when the supervisor is called upon, a strategy for the supervisor σ_t in round t is a mapping from the last idea to a possible mixed message about its potential. We let

$$\sigma_t^{\theta_t} = \Pr(m_t = \theta_t \mid \theta_t)$$

be the probability of the supervisor being honest about the potential of the observed idea. Depending on the expected strategy of the supervisor, the agent conditions her action only on the last message received.

Let the sequence $\hat{\sigma} = \{\hat{\sigma}_t^h, \hat{\sigma}_t^\ell\}_{t=1}^T$ denote the conjectured strategy of the supervisor, and let $\hat{\rho} = \{\hat{\rho}_t^h, \hat{\rho}_t^\ell\}_{t=1}^T$ denote the conjectured strategy of the agent. Given the conjectured strategy of the supervisor, the agent updates beliefs about the two unknowns – her ability and the potential of her previous ideas. The belief about her ability is β_t . Let the belief about whether her idea was as announced by the supervisor be denoted by λ_t . Observe that:

1. the public history h_t^A at the beginning of round t can be summarized by the current public belief β_t about the ability of the agent and by the belief about the true potential of the last idea produced λ_t , while
2. the private history of the supervisor h_t^S at the beginning of round t can be summarized by the current private belief β_t about the ability of the agent.²⁰

²⁰Note that we are currently not making any notational distinction between the private and the public beliefs about

We can now informally describe the notion of equilibrium. We say that a pair of sequences of conjectured strategies σ and ρ constitute an equilibrium if (1) they are both the best responses to each other given the beliefs β_t and λ_t for each t , and (2) the beliefs β_t and λ_t are consistent with what the players are conjectured to do, i.e. σ and ρ . Strategies expressed in the text without a hat constitute an equilibrium.

When both the messages are expected in equilibrium, either one of the messages will lead to a higher and the other to a lower β_t , or β_t remains the same with both the messages. We will call the former *informative* strategy and the latter *babbling* (or lying) strategy. The supervisor is expected to babble in equilibrium in round $t - 1$ if $\hat{\sigma}_{t-1}^h = 1 - \hat{\sigma}_{t-1}^\ell$, i.e. when the probability with which the supervisor is expected to reveal a true high potential idea is the same as the probability with which the supervisor incorrectly calls a low potential idea a high one. Thus, the agent is equally likely to get a positive or a negative message, and in turn does not learn from the messages. When the supervisor is expected to be informative, we will assume without loss of generality that he does so by increasing the posterior after a positive message of $m_{t-1} = h$ (and the posterior beliefs fall after a negative message $m_{t-1} = \ell$). So, we assume that $\hat{\sigma}_{t-1}^h > 1 - \hat{\sigma}_{t-1}^\ell$ for informativeness.

We will restrict attention here to informative strategies in which $\sigma^h = 1$, i.e. the supervisor always truthfully announces that the project has a high potential to succeed when he sees so. The supervisor cannot credibly commit to lying when $\theta_t = h$. In any informative strategy, a positive message $m_t = h$ should increase the posterior belief β_{t+1} of the agent. When the supervisor sees $\theta_t = h$, he has no incentive to discourage the agent. If discouragement leads to another round of experimentation, then the supervisor faces the risk of abandoning the current high potential idea and never getting a new one. Alternately, if discouragement leads to implementation then she will do so with a lower effort. In neither case a supervisor who has observed a high potential idea is better off discouraging the agent. Going forward, we assume $\sigma_t^h = 1$, and with some replace σ_t^ℓ with σ_t . Then the posterior beliefs about ability is

$$\beta_t^\ell = \frac{(1 - q)\beta_{t-1}}{1 - q\beta_{t-1}} \quad (1.8.A.1)$$

$$\beta_t^h = \frac{(1 - \hat{\sigma}_{t-1}(1 - q))\beta_{t-1}}{1 - \hat{\sigma}_{t-1}(1 - q\beta_{t-1})} \quad (1.8.A.2)$$

where $\beta_t^{m_{t-1}} = \Pr(\alpha = q|m_{t-1})$ is the posterior belief of the agent about her ability after

ability. This is to keep things simple. The two will coincide as long as the supervisor is honest. When the supervisor is not honest, the beliefs diverge only when the agent best responds to a dishonest message by experimenting again. This plays a role only in checking for deviations when constructing other informative equilibria.

receiving message m_{t-1} given the conjecture $\hat{\sigma}_{t-1}$. And

$$\lambda_t^\ell = 1 \quad (1.8.A.3)$$

$$\lambda_t^{\hat{n}} = \frac{q}{q + (1 - \hat{\sigma}_{t-1})(1 - q)} \quad (1.8.A.4)$$

where $\lambda_t^{m_{t-1}} = \Pr(\theta_{t-1} = m_{t-1} | m_{t-1})$ is the belief about whether the supervisor's message m_{t-1} matches the true potential of the idea given the conjectured $\hat{\sigma}_{t-1}$.

Thus, the value of a negative message under any informative strategy is the same as in a truth-telling strategy. When an agent receives $m_t = \ell$ then she can be sure that $\theta_t = \ell$ and she revises her belief about her ability downwards to the maximum extent. Under this condition, the agent must decide what to do following a message of $m_t = \hat{n}$ since a positive message cannot be trusted.

Proof of Lemma 1.2

Proof. Part 1: Existence of β_0^{FI}

For a given set of parameters, there is no straightforward closed form solution to the equation in condition C2. We therefore need to establish the existence of belief threshold(s). First, it can be verified that both the LHS and RHS of condition C2 are monotonically increasing and convex in β . We have

$$\begin{aligned} \frac{\partial \text{LHS}}{\partial \beta} &= \frac{q}{2} + \frac{(k\beta')^2}{2} \left(\frac{2}{\beta} - q \right) > 0 \\ \frac{\partial^2 \text{LHS}}{\partial \beta^2} &= \frac{k^2(1-q)^2}{(1-\beta q)^3} > 0 \end{aligned}$$

and

$$\begin{aligned} \frac{\partial \text{RHS}}{\partial \beta} &= k^2 \beta \geq 0 \\ \frac{\partial^2 \text{RHS}}{\partial \beta^2} &= k^2 > 0. \end{aligned}$$

Second, we show that if $2c < q(1 - k^2)$ then the threshold belief β_0^{FI} is unique. Consider the range of beliefs $0 \leq \beta \leq 1$. Since $c > 0$ and LHS at $\beta = 0$ is zero, RHS cuts the LHS from above at least once. Now, under the assumption $2c < q(1 - k^2)$, it can be verified that RHS at $\beta = 1$ is lower than LHS at $\beta = 1$. Since both LHS and RHS are monotonically increasing, they must intersect at exactly one point. Call that belief β_0^{FI} . Thus, β_0^{FI} exists and is unique.

Third, we need to show that if there exists a unique threshold belief β_0^{FI} , then $2c <$

$q(1 - k^2)$. If there is a unique belief threshold then it must be the case that there is a unique point of intersection of LHS and RHS in condition C2. Again, RHS cuts the LHS from above because at $\beta = 0$ $c > 0$. Therefore, given the monotonicity of the two functions, a sufficient condition for uniqueness is $\text{LHS}|_{\beta=1} > \text{RHS}|_{\beta=1}$. This gives $\frac{q}{2} + (1 - q)\frac{k^2}{2} > \frac{k^2}{2} + c$, which can be rearranged to $2c < q(1 - k^2)$.

Lastly, we need to show that the agent does not experiment when $2c \geq q(1 - k^2)$. This is so because then the RHS is always above the LHS, so that even experimentation once is not beneficial. When $2c \geq q(1 - k^2)$ we have that $\text{LHS}|_{\beta=1} \leq \text{RHS}|_{\beta=1}$. Given that both LHS and RHS of condition (C2) are increasing convex functions, a concern is that there might be two points of intersection. However, it is easy to verify that the slope of the RHS is lower than the slope of the LHS at both $\beta = 0$ and $\beta = 1$. This precludes such a possibility. Therefore, the agent does not want to experiment when $2c \geq q(1 - k^2)$ as the RHS is always above the LHS.

Part 2: Optimal decision rule I_t^{FI}

Condition C2 is the condition for experimenting in the worst case scenario, that is when the agent knows she is going to stop after another ℓ idea. Therefore, it follows that $I_t^{FI} = 0$ in $\beta \geq \beta_0^{FI}$ if $\theta_{t-1} = \ell$, i.e the agent continues experimenting.

Next, note that the agent cannot continue experimenting forever after ℓ ideas because at the limit the value of experimentation goes to $-c$. This is so because at the limit the belief about ability goes to zero while the cost of experimentation is a positive constant. Thus, what we need to show is that the agent does not want to experiment even once when condition C2 does not hold, i.e. $I_t^{FI} = 1$ for beliefs $\beta_t < \beta_0^{FI}$ if $\theta_{t-1} = \ell$ is the optimal decision rule.

Suppose not. Say that for some belief $\tilde{\beta} < \beta_0^{FI}$, it does not pay to experiment just once but it pays to experiment at least \tilde{T} times and then stop (Note from above, she does not want to experiment forever). Now at round $\tilde{T} - 1$ when belief is $\tilde{\beta}_{\tilde{T}-1}$ it must be that condition C2 holds i.e.

$$\frac{\tilde{\beta}_{\tilde{T}-1}q}{2} + (1 - \tilde{\beta}_{\tilde{T}-1}q)\frac{(\tilde{\beta}_{\tilde{T}-1}k)^2}{2} \geq \frac{(\tilde{\beta}_{\tilde{T}-1}k)^2}{2} + c$$

But now since $\tilde{\beta}_{\tilde{T}-1} \leq \tilde{\beta} < \beta_0^{FI}$ and we know that for any belief $\beta < \beta_0^{FI}$ condition C2 does not hold, this is a contradiction.

Finally, we have already shown the proof of the choice of e^{FI} in the main text. \square

Proof of Lemma 1.3

Proof. Fix the parameters such that $2c < (q + (1 - q)k)^2 - k^2$. Since, $q(1 - k^2) > (q + (1 - q)k)^2 - k^2$, both β_0^{NI} and β_0^{FI} exist and are unique. To compare β_0^{NI} and β_0^{FI} , we only need

to compare the LHS of the equation that defines condition (C1) with the LHS of the equation that defines condition (C2). We can then compare them with a common RHS.

Observe that the LHS of both the conditions are increasing and convex in β . Further, as $\beta \rightarrow 0$ the LHS in both the conditions also tend to zero. Thus, to establish a relationship between them it is sufficient to look at the behaviour of the LHS as $\beta \rightarrow 1$. This is equal to $\frac{(q+(1-q)k)^2}{2}$ for condition C1 and $\frac{q+(1-q)k^2}{2}$ for condition C2. Again, it can be shown that $\frac{(q+(1-q)k)^2}{2} < \frac{q+(1-q)k^2}{2}$ which is equivalent to $q(1-k^2) > (q+(1-q)k)^2 - k^2$. This implies that the LHS of condition C1 lies below the LHS of condition C2 for all $\beta > 0$. Thus, $\beta_0^{NI} > \beta_0^{FI}$. \square

Proof of Proposition 1.2

Proof. We prove this statement in steps by considering different regions of starting prior β_1 . There exists a $j \geq 0 \in \{0, 1, 2, \dots\}$ where belief β_j^{FI} is such that $\beta_j^{FI} < \beta_0^{NI} \leq \beta_{j+1}^{FI}$. The value that j takes depends on the parameters.

Step 1: Proving babbling is a unique equilibrium for $\beta_0^{FI} \leq \beta_1 < \beta_1^{FI}$

Consider any informative strategy $\hat{\sigma}_1 \in (0, 1]$ including the truth-telling strategy. In any such strategy a message $m_1 = \ell$ is only used when $\theta_1 = \ell$. So the agent believes such a message ($\lambda_2^\ell = 1$) with the posterior about ability $\beta_2^\ell < \beta_0^{FI}$ which makes the agent experiment only once at $t = 1$ and then exert $e = \beta_2^\ell k$ (see Proposition 1.1). A message $m_1 = h$ instead leads to a higher belief $\beta_2^h \in (\beta_1, 1]$, which can either push the agent to implement her idea with a higher effort or to experiment again (depending on $\hat{\sigma}_1$ and $\hat{\sigma}_2$).

If the agent best responds to $m_1 = h$ implementing her idea, she exerts effort $e = \beta_2^h(\lambda_2^h + (1 - \lambda_2^h)k) > \beta_2^\ell k$. In this case, the supervisor type $\theta_1 = \ell$ is better off deviating and sending a message $m_1 = h$ and getting a higher expected probability of success of $\beta_2^h \beta_2^\ell (\lambda_2^h + (1 - \lambda_2^h)k)k$ instead of $(\beta_2^\ell k)^2$. If the agent best responds to $m_1 = h$ by experimenting again, then also the supervisor type $\theta_1 = \ell$ is better off always sending the message $m_1 = h$. This is because the supervisor always prefers experimentation when the current idea is low potential. Thus, the supervisor has an incentive to deviate in either case.

Thus, only the babbling strategy remains which is always an equilibrium. The agent's equilibrium strategy is to implement her outside information idea, i.e. $I_1 = 1$ with $e = \beta_1 k$ since $\beta_1 < \beta_0^{NI}$ (see Lemma 1.1).

Step 2: Proving babbling is a unique equilibrium for $\beta_1^{FI} \leq \beta_1 < \beta_0^{NI}$

If $j = 0$, then either $\beta_0^{FI} \leq \beta_1 < \beta_0^{NI} < \beta_1^{FI}$ or $\beta_0^{FI} < \beta_0^{NI} \leq \beta_1 < \beta_1^{FI}$. In either

case, the scenario highlighted in Step 2 does not exist. Step 1 is sufficient in this case.

If $j = 1$ then it is enough to show that babbling is the unique equilibrium in the range $\beta_1^{FI} \leq \beta_1 < \beta_0^{NI}$ with the knowledge that if the posterior $\beta_2 < \beta_1^{FI}$ then the supervisor babbles (from Step 1 above). Note that any informative messaging strategy conjecture for $t = 1$ with $\hat{\sigma}_1 \in (0, 1]$ must lead to a posterior $\beta_2^\ell < \beta_1 < \beta_2^h$. Now, as before the value of message $m_1 = \ell$ is the same as in truth-telling so that $\beta_2^\ell \in [\beta_0^{FI}, \beta_1^{FI})$. From Step 1 above, the supervisor is then expected to babble in $t = 2$ and the agent best responds by choosing to implement her low potential idea ($I_2 = 1$) from $t = 1$ with effort $e = \beta_2^\ell k$. A message $m_1 = h$ again leads to a higher belief $\beta_2^h \in (\beta_1, 1]$, which can either push the agent to implement her idea with a higher effort or to experiment again (depending on $\hat{\sigma}_1$ and $\hat{\sigma}_2$). As before now, the supervisor type $\theta_1 = \ell$ is better off deviating and sending a message $m_1 = h$. Thus, babbling is the unique equilibrium strategy of the supervisor.

If $j \in \{2, 3, \dots\}$, then it needs to be shown that babbling is a unique equilibrium strategy in the ranges $\beta_1^{FI} \leq \beta_1 < \beta_2^{FI}, \dots, \beta_{j-1}^{FI} \leq \beta_1 < \beta_j^{FI}$ and $\beta_j^{FI} \leq \beta_1 < \beta_0^{NI}$. Consider first the range $\beta_1^{FI} \leq \beta_1 < \beta_2^{FI}$. Any posterior β_2^ℓ for priors $\beta_1^{FI} \leq \beta_1 < \beta_2^{FI}$ must map in to the range of beliefs highlighted in Step 1. This implies that supervisor type $\theta_1 = \ell$ cannot credibly commit to sending a message $m_1 = \ell$. Such a message leads to the agent implementing with effort $e = \beta_2^\ell k$. This makes babbling a unique equilibrium strategy for $\beta_1^{FI} \leq \beta_1 < \beta_2^{FI}$. The same logic applies to all the ranges of prior belief up to β_j^{FI} . Then, in the range $\beta_j^{FI} \leq \beta_1 < \beta_0^{NI}$ the proof is identical to the above described $j = 1$ case.

Therefore, babbling is the unique equilibrium strategy of the supervisor and the agent does not experiment, i.e. $I_1 = 1$ and $a = \beta_1 k$.

Step 3: Proving babbling is a unique equilibrium for $\beta_0^{NI} \leq \beta_1 < \beta_1^{NI}$

For $j = 0$, we have already shown that babbling is a unique equilibrium strategy for $\beta_0^{FI} \leq \beta_1 < \beta_0^{NI} < \beta_1^{FI}$ or $\beta_0^{FI} < \beta_0^{NI} \leq \beta_1 < \beta_1^{FI}$. Note that since $\beta_0^{FI} < \beta_0^{NI}$, it must be the case that $\beta_1^{FI} < \beta_1^{NI} < \beta_2^{FI}$. So, it remains to show that babbling is unique for $\beta_1^{FI} \leq \beta_1 < \beta_1^{NI}$. This argument is the same as the one presented below.

Any informative mixing for $j \geq 1$ leads to $\beta_2^\ell < \beta_0^{NI}$. The supervisor babbles in the range of posteriors $\beta_0^{FI} \leq \beta_2^\ell < \beta_0^{NI}$ from Step 1 and 2 above (and for $j = 0$ case the supervisor babbles in the range $\beta_0^{FI} \leq \beta_2^\ell < \beta_1^{FI}$), and the agent chooses to implement thereafter (from Lemma 1.1). A message $m_1 = h$, on the other hand, is believed and the agent best responds by either implementing with a higher belief or experimenting again. Therefore, the supervisor can do better by lying instead when he observes $\theta_1 = \ell$ when he is expected to be informative. \square

Proof of Proposition 1.3

Proof. We prove the proposition in two parts.

Part 1: To show that if $\sigma_1 = 1$ is an equilibrium for $\beta_1^{NI} \leq \beta_1 < \beta_2^{NI}$, then it must be an equilibrium for all $\beta_2^{NI} \leq \beta_1 < 1$.

Consider the region of priors $\beta_2^{NI} \leq \beta_1 < \beta_3^{NI}$. We check whether $\hat{\sigma}_1 = 1$ is an equilibrium. Here, the supervisor has an incentive to reveal the truth about $\theta_1 = \ell$ if the expected probability of success by sending $m_1 = \ell$ is higher than that from sending the message $m_1 = h$. If he sends a message $m_1 = \ell$, the agent at most experiments two more times - consulting the supervisor after one (which is at β_2^ℓ where the supervisor is again honest given the premise) and not doing so after the other. Therefore, the expected probability of success by sending $m_1 = \ell$ is

$$\beta_2^\ell q + (1 - \beta_2^\ell q)(\beta_3^\ell)^2(q + (1 - q)k)^2$$

By lying the supervisor convinces the agent that her idea has a high potential to succeed ($\hat{\lambda}_2^h = 1$) and that she is of ability q ($\hat{\beta}_2^h = 1$). She then exerts $e = 1$ to implement her idea. However, the supervisor has an updated belief of β_2^ℓ knowing that $\theta_1 = \ell$. Thus, expected probability of success by sending $m_1 = h$ is $\beta_2^\ell k$.

The supervisor has an incentive to reveal the truth at this stage if

$$\beta_2^\ell k \leq \beta_2^\ell q + (1 - \beta_2^\ell q)(\beta_3^\ell)^2(q + (1 - q)k)^2.$$

It is easy to check that the above condition is always holds with a strict inequality sign under Assumption (A). So, $\sigma_1 = 1$ is an equilibrium for $\beta_2^{NI} \leq \beta_1 < \beta_3^{NI}$.

Now, for any $\beta_1 \geq \beta_3^{NI}$ the supervisor can make the agent experiment (and if any of the following ideas has a high potential to succeed make them exert $e = 1$ on it) at least three more times by honestly revealing $\theta = \ell$. Given Assumption (A), this should always be an equilibrium.

Part 2: To show that $\sigma_1 = 1$ is an equilibrium for $\beta_1^{NI} \leq \beta_1 < \beta_2^{NI}$ if and only if $c \geq \frac{\kappa k - (\kappa k)^2}{2}$ where $\kappa \equiv \frac{k}{(q + (1 - q)k)^2}$ and $k < (q + (1 - q)k)^2$.

Suppose $c \geq \frac{\kappa k - (\kappa k)^2}{2}$ where $\kappa \equiv \frac{k}{(q + (1 - q)k)^2}$ and $k < (q + (1 - q)k)^2$. Now consider the conjectured strategy $\hat{\sigma}_1 = 1$ for $\beta_1^{NI} \leq \beta_1 < \beta_2^{NI}$. When the supervisor observes $\theta_1 = \ell$,

his expected probability of success by sending message $m_1 = \ell$ is

$$(\beta_2^\ell)^2(q + (1 - q)k)^2.$$

Following $m_1 = \ell$, the agent experiments once more but does not consult the supervisor thereafter. Thus, she implements her idea of unknown potential by exerting effort $e = \beta_2^\ell(q + (1 - q)k)$. On the other hand by sending a message $m_1 = \hat{\ell}$ when the agent expects supervisor to be honest leads her to exert $e = 1$ in implementing a $\theta_1 = 1$ idea. This is so because she believes in the supervisor's message, $\hat{\lambda}_2^{\hat{\ell}} = 1$ and $\hat{\beta}_2^{\hat{\ell}} = 1$. The expected probability of success is then $\beta_2^\ell k$.

Truth-telling is an equilibrium if

$$\begin{aligned} (\beta_2^\ell)^2(q + (1 - q)k)^2 &\geq \beta_2^\ell k \\ \implies \beta_1 &\geq \frac{k}{qk + (1 - q)(q + (1 - q)k)^2} := \beta^{\text{truth}}. \end{aligned}$$

$\hat{\sigma}_1 = 1$ is an equilibrium if for all $\beta_1 \in [\beta_1^{NI}, \beta_2^{NI})$, it is also the case that $\beta_1 \geq \beta^{\text{truth}}$. This can happen iff $\beta^{\text{truth}} \leq \beta_1^{NI}$. This condition then be rearranged given β^{truth} from above and $\beta_0^{NI} = \left(\frac{2c}{(q+(1-q)k)^2 - k^2}\right)^{\frac{1}{2}}$ (from Lemma 1.1), and using the fact that $\beta_1^{NI} = \frac{\beta_0^{NI}}{1 - q(1 - \beta_0^{NI})}$. This gives us

$$c \geq \frac{\kappa k - (\kappa k)^2}{2}$$

where $\kappa \equiv \frac{k}{(q+(1-q)k)^2}$ and we need that $k < (q + (1 - q)k)^2$. But this is also our premise. Thus, $\sigma_1 = 1$ is an equilibrium.

Alternately, suppose $\sigma_1 = 1$ is an equilibrium for $\beta_1^{NI} \leq \beta_1 < \beta_2^{NI}$. Then it must be the case that $\beta_1 \geq \beta^{\text{truth}}$ for all $\beta_1 \in [\beta_1^{NI}, \beta_2^{NI})$. Specifically, it must be that $\beta_1^{NI} \geq \beta^{\text{truth}}$. This condition can then be rearranged to yield

$$c \geq \frac{\kappa k - (\kappa k)^2}{2}$$

where $\kappa \equiv \frac{k}{(q+(1-q)k)^2}$ and with an added constraint $k < (q + (1 - q)k)^2$. □

Proof of Lemma 1.4

Proof. It is immediate to see that an increase in belief from β to β' such that

1. $\beta_0^{FI} \leq \beta < \beta' < \beta_0^{NI}$ is welfare improving. This is because $\frac{(\beta k)^2}{2} > \frac{(\beta' k)^2}{2}$ which we get by replacing the optimal effort $e = \beta k$ in the expected utility function.
2. $\beta_0^{NI} \leq \beta < \beta' < \beta_1^{NI}$ is welfare improving. This is because $\frac{(\beta(q+(1-q)k))^2}{2} >$

$\frac{(\beta'(q+(1-q)k))^2}{2}$ which we get by replacing the optimal effort $e = \beta(q + (1 - q)k)$ in the expected utility function.

Now consider an increase in belief from β to β' such that $\beta_j^{NI} \leq \beta < \beta' < \beta_{j+1}^{NI}$ such that $j > 1$. Denote the ex-ante expected utility or welfare of the agent at prior β by $W(\beta)$. We have that

$$W(\beta) = \beta \frac{q}{2} [1 + (1 - q) + \dots + (1 - q)^{j-1}] - \beta c [1 + (1 - q) + \dots + (1 - q)^j] \\ + \beta (1 - q)^j [Ke - \frac{e^2}{2}] - (1 - \beta) [(j + 1)c + \frac{e^2}{2}]$$

where $K = q + (1 - q)k$. Similarly, we can write $W(\beta')$ keeping in mind that the maximum number of attempts is still $j + 1$.

Now, comparing term-by-term, it is obvious that everything other than the comparison of $\beta'(1 - q)^j Ke' - ((1 - \beta') + \beta'(1 - q)^j) \frac{e'^2}{2}$ with $\beta(1 - q)^j Ke - ((1 - \beta) + \beta(1 - q)^j) \frac{e^2}{2}$ in $W(\beta')$ is greater than that in $W(\beta)$. Thus it is sufficient to show that

$$\beta'(1 - q)^j Ke' - ((1 - \beta') + \beta'(1 - q)^j) \frac{e'^2}{2} > \beta(1 - q)^j Ke - ((1 - \beta) + \beta(1 - q)^j) \frac{e^2}{2}$$

which can be rearranged to

$$\beta'(1 - q)^j Ke' - (1 - \beta'(1 - (1 - q)^j)) \frac{e'^2}{2} > \beta(1 - q)^j Ke - (1 - \beta(1 - (1 - q)^j)) \frac{e^2}{2}$$

where $e = K\beta_{j+1}^\ell$ and $e' = K\beta_{j+1}'^\ell$.

Now it is easy to check that $Ke - \frac{e^2}{2}$ is increasing in beliefs. So that

$$Ke' - \frac{e'^2}{2} > Ke - \frac{e^2}{2} \\ \implies Ke' - (1 - \beta'(1 - (1 - q)^j)) \frac{e'^2}{2} > Ke - (1 - \beta(1 - (1 - q)^j)) \frac{e^2}{2} \\ \implies \beta'(1 - q)^j Ke' - (1 - \beta'(1 - (1 - q)^j)) \frac{e'^2}{2} > \beta(1 - q)^j Ke - (1 - \beta(1 - (1 - q)^j)) \frac{e^2}{2}$$

where in the second step the inequality is preserved because a greater number is added to the LHS than the RHS. And in the third step the inequality is again preserved because Ke' (which is greater than Ke) on the LHS is multiplied with a greater number than Ke in the RHS. Hence, the welfare has increased. \square

Proof of Lemma 1.5

Proof. Using the language introduced in Lemma 1.4, we can write $W(\beta)$ and $W(\beta')$ where $\beta < \beta_0^{NI}$ and $\beta_0^{NI} \leq \beta < \beta_1^{NI}$ as

$$W(\beta) = \frac{(\beta k)^2}{2} \text{ and } W(\beta') = \frac{(\beta' K)^2}{2} - c$$

Now, if the agent finds herself in $[\beta_0^{NI}, \beta_1^{NI})$ then Condition (C1) must be slack. This means

$$\frac{(\beta' K)^2}{2} - c > \frac{(\beta' k)^2}{2} > \frac{(\beta k)^2}{2}$$

where the second inequality follows from the fact that $\beta' > \beta$. Hence, the welfare has increased. \square

Proof of Lemma 1.6

Proof. We show here the proof of how an increase in belief from $\beta = \beta_1^{NI} - \varepsilon$ to $\beta' = \beta_1^{NI}$ is welfare improving. The general proof of an increase in the prior from $\beta_{j+1}^{NI} - \varepsilon$ to β_{j+1}^{NI} follows the same argument.

We can write the ex-ante expected welfare in the two cases as follows:

$$\begin{aligned} W(\beta_1^{NI} - \varepsilon) &= (\beta_1^{NI} - \varepsilon)Ke - \frac{e^2}{2} - c \\ W(\beta_1^{NI}) &= \beta_1^{NI} \frac{q}{2} + \beta_1^{NI}(1-q)Ke' - (c + \frac{e'^2}{2})(1 - \beta_1^{NI}q) - c \end{aligned}$$

where $e = (\beta_1^{NI} - \varepsilon)K$ and $e' = \beta_1^{NI}K$.

Now, if $W(\beta_1^{NI}) > W(\beta_1^{NI} - \varepsilon)$, then substituting for e and e' , letting $\varepsilon \rightarrow 0$, and simplifying the inequality by using $\beta_0^{NI} = \frac{(1-q)\beta_1^{NI}}{1-q\beta_1^{NI}}$ gives

$$\beta_1^{NI} \frac{q}{2} - c(1 - \beta_1^{NI}q) - \frac{(\beta_1^{NI}K)^2}{2} > -\frac{K^2}{2}(1-q)\beta_1^{NI}\beta_0^{NI}.$$

If the above inequality holds, then we are done.

Let $2c < q(1 - K^2)$. Under this assumption, Condition (C2) must hold in a way that k is replaced with K as

$$\beta_1^{NI} \frac{q}{2} - c > \frac{(\beta_1^{NI}K)^2}{2} - (1 - \beta_1^{NI}q) \frac{(\beta_0^{NI}K)^2}{2}.$$

Now, the inequality is preserved if the c on the LHS is reduced. Then rearranging gives

$$\beta_1^{NI} \frac{q}{2} - (1 - \beta_1^{NI} q) c - \frac{(\beta_1^{NI} K)^2}{2} > -(1 - \beta_1^{NI} q) \frac{(\beta_0^{NI} K)^2}{2}.$$

It is now straightforward to verify that $(1 - \beta_1^{NI} q) \frac{(\beta_0^{NI} K)^2}{2} = \frac{K^2}{2} (1 - q) \beta_1^{NI} \beta_0^{NI}$, so that our original hypothesis on welfare comparison holds. \square

Proof of Proposition 1.4

Proof. From Lemma 1.4 and 1.5, it is immediate that an increase in belief of up to, but not including the level β_1^{NI} is welfare improving. Now, from Lemma 1.6, an epsilon increase in belief that pushes the agent in to experimentation with supervision is also welfare improving. Finally, from Lemma 1.4, any increase in belief of up to but not including the level β_2^{NI} is welfare improving. This reasoning can then be extended for any increase in belief. \square

Proof of Proposition 1.5

Proof. To prove the statement, we consider two particular situations, and show how in each the welfare at the correct and overconfident beliefs differ. Let $W(\beta; b)$ be the ex-ante expected utility of the agent when the common prior is β but the correct belief is b .

Part I: Showing that overconfidence can be welfare improving

Let $\beta = \beta_1^{NI}$ but $b = \beta_1^{NI} - \varepsilon$. The two expected utility functions can be written as

$$\begin{aligned} W(b; b) &= \frac{(bK)^2}{2} - c \\ W(\beta_1^{NI}; b) &= \frac{bq}{2} + b(1 - q)\beta_0^{NI}K^2 - (1 - bq)\left(c + \frac{(\beta_0^{NI}K)^2}{2}\right) - c \end{aligned}$$

We need to show if $W(\beta_1^{NI}; b) > W(b; b)$. In order to do so, first observe that $\frac{bq}{2} > \frac{(bK)^2}{2}$. This follows immediately from Assumption (A). So if we are able to show that

$$b(1 - q)\beta_0^{NI}K^2 - (1 - bq)\left(c + \frac{(\beta_0^{NI}K)^2}{2}\right) \geq 0$$

then we are done. Rearranging the above and recognizing that $\frac{b(1-q)}{1-bq} = \beta_0^{NI} - \varepsilon'$ where $\varepsilon' \neq \varepsilon$, we need that

$$\frac{(\beta_0^{NI}K)^2}{2} \geq \varepsilon' \beta_0^{NI}K^2 + c$$

But we know from Condition (C1) that

$$\frac{(\beta_0^{NI}K)^2}{2} = \frac{(\beta_0^{NI}k)^2}{2} + c.$$

Therefore, it is possible to find an ε' (and consequently ε) such that welfare improves under overconfidence. This requires $\varepsilon' \leq \beta_0^{NI} \frac{k^2}{2K^2}$.

Part 2: Showing that overconfidence can be welfare reducing

Let $\beta = \beta_0^{NI}$ but $b < \beta_0^{NI}$. The two expected utility functions can be written as

$$W(b; b) = \frac{(bk)^2}{2}$$

$$W(\beta_0^{NI}; b) = b\beta_0^{NI}K^2 - \frac{(\beta_0^{NI}K)^2}{2} - c$$

This time we need to show that $W(\beta_0^{NI}; b) < W(b; b)$. Again using Condition (C1) to substitute for $-\frac{(\beta_0^{NI}K)^2}{2} - c$ in $W(\beta_0^{NI}; b)$, we can reduce the above to

$$b < \beta_0^{NI} \left(\frac{2K^2}{k^2} - 1 \right),$$

which must always be true because $\frac{2K^2}{k^2} - 1 > 1$. □

Proof of Proposition 1.6

Proof. Let $\phi_S = \phi_A = 1$. Consider the supervisor who has seen a $\theta_{t-1} = \ell$ and reveals it honestly to the agent. His value function is given by

$$\mathcal{V}_S^\ell(\beta_t) = \max \left\{ \frac{(\beta_t k)^2}{2}, \frac{\beta_t q}{2} - c_S + (1 - \beta_t q) \mathcal{V}^\ell(\beta_{t+1}) \right\}.$$

where the first term is the value that the supervisor would get if he gets the low idea implemented and the second term is what he would get if he gets experimentation again. Given his costs, he would then like the agent to continue experimenting for as long as

$$\frac{\beta q}{2} + (1 - \beta q) \frac{(\beta' k)^2}{2} \geq \frac{(\beta k)^2}{2} + c_S,$$

which gives a belief threshold β_{S0}^{FI} . However, under an honest strategy, the agent would like to continue experimenting for as long as

$$\frac{\beta q}{2} + (1 - \beta q) \frac{(\beta' k)^2}{2} \geq \frac{(\beta k)^2}{2} + c_A,$$

which gives a belief threshold β_{A0}^{FI} .

Now, if $c_S < c_A$ then $\beta_{S0}^{FI} < \beta_{A0}^{FI}$ so that the supervisor would like the agent to experiment beyond β_{A0}^{FI} . The supervisor then fears discouraging the agent through honest

revelation for any prior belief that leads the agent to a belief lower than β_{A0}^{FI} . Therefore, the results of Propositions 1.1, 1.2 and 1.3 hold.

On the other hand if $c_S \geq c_A$, then $\beta_{S0}^{FI} > \beta_{A0}^{FI}$. The agent would like to experiment more than β_{S0}^{FI} . Consider a belief $\beta_{S0}^{FI} \leq \beta_1 < \beta_{S1}^{FI}$ and consider the expected strategy of honesty. When the supervisor has seen a low potential idea, then by announcing it truthfully he gets an effort of $e = \beta_2^\ell k$ which is also optimal from the point of view of the supervisor because $\phi_S = 1$. This is so because it is an equilibrium strategy for the supervisor to babble tomorrow. So, even though the agent at this stage would like to experiment again but in the absence of honesty tomorrow, and $\beta_2^\ell < \beta_{A0}^{NI}$ she prefers to implement. By deviating and calling it a high potential idea, he induces an effort of 1 on a low-potential idea. However, this is suboptimal from his perspective, since he would also the full cost of implementation. Thus, there is no incentive to lie and honesty is an equilibrium. \square

Proof of Proposition 1.8

Proof. Consider a prior $\beta_j^{NI} \leq \beta_1 < \beta_{j+1}^{NI}$. In an immediately honest equilibrium strategy, the agent experiments for j rounds with subsequent messages $m = \ell$ before reaching the babbling region so that $\beta_0^{NI} \leq \beta_j^\ell < \beta_1^{NI}$. In addition, the agent experiments one extra round without supervision. Now, consider any strategy that reveals some $j' \leq j$ outcomes together. Let the round of eventual revelation be denoted by τ . Now, the agent is induced to experiment a higher number of rounds in this strategy iff $\beta_\tau^\ell < \beta_0^{NI} \leq \beta_j^\ell$. Say that this is the case. We determine whether such a strategy is an equilibrium.

Observe that at $\beta_\tau^\ell < \beta_0^{NI}$ the agent best responds by abandoning experimentation and implementing any one of her low potential ideas with an effort $\beta_\tau^\ell k$. If the supervisor is honest, his expected payoff is $(\beta_\tau^\ell k)^2$. By deviating, and calling any one of the low potential ideas a high one, the supervisor is able to induce an effort of 1 by the agent on that idea. This gives the supervisor an expected payoff of $\beta_\tau^\ell k$. Since the latter is greater than the former, such an eventually honest strategy cannot be an equilibrium. \square

B Additional proofs not in the main text

Comparative statics of β_0^{FI}

Lemma 1.7. β_0^{FI} is increasing in e , increasing in k , and decreasing in q .

Proof. Consider, first, an exogenous increase in e . It is easy to verify that an increase in e raises the value of the RHS (i.e. of implementing the idea) in condition C2 for every belief level β . This raises the β_0^{FI} .

Second, consider the effect of an exogenous increase in k .

$$\frac{\partial \text{LHS}}{\partial k^2} = (1 - \beta q) \frac{(\beta')^2}{2}$$

$$\frac{\partial \text{RHS}}{\partial k^2} = \frac{\beta^2}{2}.$$

Now, since $\beta > \beta'$ and $1 > \beta q$, $\frac{\partial \text{LHS}}{\partial k^2} < \frac{\partial \text{RHS}}{\partial k^2}$. Thus, the value from implementing increases by more than the value from experimenting, which leads to a higher β_0^{FI} .

Finally, consider an exogenous increase in q . The RHS remains unchanged with an increase in q . For the LHS,

$$\frac{\partial \text{LHS}}{\partial q} = \frac{\beta}{2} - k^2 \beta \beta' \left(1 - \frac{\beta'}{2}\right).$$

This is positive if $\frac{1}{2} > k^2 \beta' \left(1 - \frac{\beta'}{2}\right)$, which is true since $\frac{\partial k^2 \beta' \left(1 - \frac{\beta'}{2}\right)}{\partial \beta'} = k^2 (1 - \beta') > 0$ and at the limits the inequality holds. As $\beta' \rightarrow 0$, we have that $k^2 \beta' \left(1 - \frac{\beta'}{2}\right) \rightarrow 0$ and as $\beta' \rightarrow 1$, $k^2 \beta' \left(1 - \frac{\beta'}{2}\right) \rightarrow \frac{k^2}{2}$. \square

An exogenous increase in k makes executing a low potential idea more attractive and therefore, leads to a higher β_0^{FI} and reduces the incentives to experiment for long. The agent desires to finish the project with a sufficiently high belief so that he can exert a higher effort in implementing a low potential idea (if need be), thereby maximizing the probability of success even with a poor idea. Finally, an increase in q lowers the belief threshold. This is so because conditional on being of high-ability, a higher q increases the chances of coming up with a high potential idea. Therefore, in a world in which ability is unknown it makes experimentation more attractive and pushes the agent to experiment for longer.

Comparative statics of β_0^{NI}

It is straightforward to derive how β_0^{NI} behaves with a change in parameters. A decrease in the probability of coming up with a high potential idea q or an increase in the cost of experimentation c has the effect of increasing the threshold. Finally, an increase in k can

have a non-monotonic effect on β_0^{NI} depending on the initial value. For $k < \frac{1-q}{2-q}$, an increase in k decreases β_0^{NI} . For $k > \frac{1-q}{2-q}$, an increase in k increases β_0^{NI} . The intuition behind a non-monotonic relation between k and β_0^{NI} is as follows. k measures the success rate (for any given effort level) from a bad idea when the agent is of high-ability. When the agent does not observe the value of θ from experimentation, then she experiments only as a gamble (and this gamble is not worth taking more than once). When k increases from a sufficiently low k to begin with, it makes this gamble more attractive – the agent reasons that even if the gamble fails (i.e. $\theta = \ell$ is the outcome of the gamble), she is more likely to succeed because of a higher k . On the other hand, when k increases further from an already high level, then the gamble becomes less attractive. This is so because the agent already has an outside option $\bar{\theta} = \ell$ available which then becomes relatively more attractive to finish.

C Committed supervisor

A note on the enforcement of commitment

Here we present the case of the supervisor committing to an information policy before the agent starts experimenting with ideas. Before we do so, we should understand how such a commitment may be enforced. An information disclosure policy is a sequence of revelation strategies about the observed potential of ideas produced by the agent to which the supervisor is committed. One may imagine the policy as a sequence of public tests - the supervisor may or may not observe the true potential of the idea but he designs tests that will reveal to the agent (and to the supervisor) the true potential of the idea. Thus, commitment to information disclosure policy is akin to commitment to test designs. This interpretation is in the spirit of Kamenica and Gentzkow (2011) and Smolin (2017).

Another way in which such a commitment may be enforced is through “presentation” of ideas to multiple supervisors. Many co-supervisors rather than one main supervisor may work to discipline each other. This requires that if the optimal disclosure policy involves mixing by the supervisors then they all should agree on such a mixing and then enforce it (say by punishing deviations with full disclosure). Alternately, one supervisor’s recommendation may be cross-examined by another supervisor who has also observed the agent’s idea. However, these interpretations are not immediate and might not be realistic in many settings. An apprentice working on a project might only be assigned one expert due to cost concerns. It is also not obvious how a supervisor might commit to a test design that reveals his private information to the agent. Because of this limitation, we present the commitment case as an extension of the model in Section 1.4. We consider here only the flavour of an optimal policy by discussing the incentives of the supervisor and the agent, and showing how the supervisor can achieve better outcomes (relative to the equilibrium outcome) for both himself and the agent by committing to information disclosure policies.

Immediate honesty

Consider first the policy in which the supervisor is committed to revealing the true potential of the idea after each round of experimentation. We call this a policy of *immediate honesty*. As illustrated in Lemma 1.2 such a policy induces the agent to experiment with continued low potential ideas all the way down to the belief β_0^{FI} . It is immediate that the agent prefers to experiment more under this policy relative to the equilibrium outlined in Proposition 1.3. Immediate honesty guarantees maximum possible learning to the agent and in the least cost, which allows the agent to match effort to the true potential of the idea. This helps retain the

attractiveness of experimentation insofar as condition (C2) holds. The prior β_1 determines how many more rounds the agent ends up experimenting under this policy relative to the equilibrium.

That the supervisor prefers such a policy is not immediate in the region of beliefs in which the supervisor is honest in equilibrium as well. While on the one hand such a policy induces more experimentation (and therefore, a higher probability of the agent producing a high potential idea), it also depresses the effort of the agent when she does not ever produce a high potential idea. The agent exerts a higher effort in equilibrium on an idea of unknown potential (see Proposition 1.3) because of a higher belief. Let $\beta_1 > \beta_1^{NI}$ such that under both the equilibrium and the immediately honest policy the agent experiments for t rounds until β_1^{NI} , then in equilibrium the agent experiments for one additional round (without supervision) while under the immediately honest policy she does so for t' additional rounds with supervision. Note that t and t' are functions of β_1 . The supervisor prefers the immediately honest policy over the equilibrium policy iff

$$\begin{aligned} (\beta_{t+1}^\ell)^2(q + (1 - q)k)^2 &< \beta_{t+1}^\ell q + (1 - \beta_{t+1}^\ell q)\beta_{t+2}^\ell q + \\ &+ (1 - \beta_{t+1}^\ell q)(1 - \beta_{t+2}^\ell q)\beta_{t+3}^\ell q + \\ &+ \dots + (1 - \beta_{t+1}^\ell q)(1 - \beta_{t+2}^\ell q) \dots (1 - \beta_{t+t'}^\ell q)(\beta_{t+t'+1}^\ell k)^2. \end{aligned}$$

Until round t both policies yield the same payoff to the supervisor. The LHS captures the additional payoff from one more round of experimentation in $t + 1$. The RHS captures increase in the payoff from t' additional rounds of experimentation with the agent implementing a low potential idea in round $t + t' + 1$. A sufficient condition for the above to be satisfied is $q > (q + (1 - q)k)^2$, which we know is satisfied from Assumption (A). Lemma 1.8 follows from the above discussion.

Lemma 1.8. *The immediately honest policy is pareto superior to the equilibrium policy.*

Thus, both the supervisor and the agent stand to gain if the supervisor commits to honesty. However, as we show below, the supervisor can do better than immediate honesty.

Delayed honesty

The supervisor's preferred policy is driven by the desire to make the agent experiment more when she has low potential ideas but implement immediately if she gets a high potential idea. Thus, while on the one hand he wants to be honest with the agent, he also wants the agent to experiment as often as possible. We show how the supervisor can fulfil these two objectives through a delayed disclosure policy which we call *delayed honesty* and quantify the

gain attainable over immediate honesty.²¹

A policy is a combination of a disclosure time and what to recommend at that disclosure time. A disclosure timing rule is a mapping from the current belief β_t to a choice of round τ at which the supervisor requires the agent to show her ideas to him (or equivalently the number of rounds the agent is required to experiment). He then makes a comment about each of the τ ideas according to a recommendation policy which is a mapping of $\{\ell, h\}^\tau$ onto itself. A recommendation policy is honest if the supervisor honestly reveals the type of all the ideas that the agent has produced. We restrict attention to honest recommendation policies for the time being and analyse what is the optimal disclosure time τ^* . At the disclosure time τ , the agent and the supervisor update their belief about the ability sequentially according to Bayes' rule. Thus, if the supervisor reveals that any of the ideas are high potential they both update their belief to 1 and otherwise revise their belief downwards by τ times

$$\beta_\tau^\ell = \frac{(1-q)^\tau \beta_1}{1 - q\beta_1 \sum_{t=0}^{\tau-1} (1-q)^t}.$$

Fix a prior $\beta_1 \geq \beta_0^{FI}$ and consider a disclosure policy that requires the agent to experiment at least τ times to receive feedback from the supervisor. We are interested in finding out the *maximum* number of rounds of delay. Let the disclosure policy be such that after the agent discovers all her ideas were of low potential she quits experimentation and implements any one her ideas, i.e. $\beta_{\tau+1}^\ell < \beta_0^{FI}$.²² We say that such a policy is *implementable* if the agent prefers to experiment τ times and receiving feedback to not experimenting and implementing her outside option idea.²³ This yields the following implementability constraint (IC)

$$\underbrace{\frac{1}{2}\beta_1[1 - (1-q)^\tau(1 - (\beta_{\tau+1}^\ell k)^2)]}_{\text{expected benefit of experimentation}} \geq \underbrace{\frac{(\beta_1 k)^2}{2}}_{\text{opportunity cost}} + \underbrace{\tau c}_{\text{actual cost}}. \quad (\text{IC})$$

Observe that since the agent is expected to carry out multiple rounds of experimentation without knowing their outcome, she evaluates the possibility of attaining a high potential idea relative to β_1 . Conditional on being high-ability, with probability $(1-q)^\tau$ she expects to attain

²¹Since we are not focussing on delayed partial disclosure, we will omit any mathematical complexity that comes with it such as that of defining mixed strategies. We will focus on the supervisor using pure strategies.

²²If there is any implementable delayed policy that leads to a posterior above β_0^{FI} , then the same can be achieved by an immediately honest policy by inducing the same number of rounds of experimentation. We will refer to delayed honesty policy as the one which leads to posteriors below β_0^{FI} so that more number of rounds are induced than in immediately honest policy.

²³There is no expected benefit to the agent by experimenting less than τ times since given the policy the supervisor does not reveal any information to the agent when this is the case.

only low potential ideas to implement, and with the remaining probability she expects at least one high potential idea. Therefore, with probability $\beta_1(1 - (1 - q)^\tau)$ she receives $1/2$ and with probability $\beta_1(1 - q)^\tau$ she will revise her belief down to $\beta_{\tau+1}^\ell$ after the supervisor honestly reveals all her τ ideas are low potential. At this point, she will implement any one of her low potential ideas to obtain an expected benefit of $\frac{(\beta_{\tau+1}^\ell k)^2}{2}$. Finally, there is no benefit of experimentation if the agent is of low-ability type. This is captured in the LHS of IC condition as the expected benefit of experimentation.

If the agent instead opts for implementing her low potential outside option idea, she expects to receive a payoff of $\frac{(\beta_1 k)^2}{2}$. As illustrated in the RHS, she must forego this expected benefit when she decides to experiment, in addition to paying the cost of experimentation c for τ rounds. The IC condition thus puts a limit on the maximum number of rounds the agent is willing to experiment when she is at a belief β_1 and the supervisor is committed to revealing all the information after those rounds.

We next analyse the supervisor's incentives under such a policy. The supervisor's ex-ante expected payoff from a τ -implementable policy is

$$\beta_1[1 - (1 - q)^\tau(1 - (\beta_{\tau+1}^\ell k)^2)].$$

The supervisor, like the agent, only sees the potential of the ideas once they are presented to him – he evaluates the probability of at least one high potential idea among the τ attempts according to β_1 . Does the supervisor benefit from a higher or a lower τ ? While on the one hand a higher τ reduces the probability of the agent only producing low potential ideas, but on the other hand it also depresses the effort of the agent in case of such event. The following lemma shows that the first order effect of reduced probability dominates the second order effect of reduced effort so that the supervisor is always better off inducing a higher τ .

Lemma 1.9. *Under assumption (A), the supervisor's payoffs are increasing in the number of rounds the agent experiments τ .*

Proof. Consider the expected probability of success from a τ -implementable policy:

$$\beta_1[1 - (1 - q)^\tau(1 - (\beta_{\tau+1}^\ell k)^2)] \tag{1.8.C.5}$$

Now consider the expected probability of success from a $\tau + 1$ -implementable policy:

$$\beta_1[1 - (1 - q)^{\tau+1}(1 - (\beta_{\tau+2}^\ell k)^2)] \tag{1.8.C.6}$$

Subtracting equation (1.8.C.5) from (1.8.C.6) and looking at the condition for it being positive, we get

$$q + (1 - q)(\beta_{\tau+2}^\ell k)^2 - (\beta_{\tau+1}^\ell k)^2 > 0$$

This always the case since $q > k$ from Assumption (A), which implies $q > (\beta_{\tau+1}^\ell k)^2$. Therefore, the payoff of the supervisor is increasing in the number of rounds of experimentation. \square

Supervisor's maximization problem therefore reduces to getting the agent to experiment as many rounds as possible. This is solely determined by the **IC** condition. It is immediate that the expected benefit of experimentation to the agent under such a policy, although increasing in β_1 , is bounded above by $1/2$. Consequently, for a higher β_1 the agent should want to experiment more number of rounds but up to a limit. This limit is imposed by the bounded benefits on the one hand, and the increasing cost of experimentation on the other. Our objective is to determine the maximum (β_1, τ) combination that is implementable with such a policy.

For this purpose, fix τ . Now, if there exists a prior belief that makes the **IC** condition bind, then it must be the *minimum* prior that does so. Define this minimum prior belief by $\bar{\beta}^\tau$. So for any belief $\beta_1 \geq \bar{\beta}^\tau$ the agent finds it optimal to at least experiment τ times. Observe that $\bar{\beta}^\tau$ must be increasing in τ since the agent must have a higher belief to induce him to experiment more often by paying a higher cost. Let $\bar{\beta}^{\bar{\tau}}$ be the maximum of this increasing sequence so that $\bar{\tau}$ gives the maximum number of rounds that are implementable and $\bar{\beta}^{\bar{\tau}}$ is the minimum prior that can induce those many rounds. Proposition 1.9 follows from the above discussion.

Proposition 1.9. *The maximum number of rounds τ^* the supervisor can delay honestly revealing the outcomes and therefore induce experimentation at prior β_1 is given by*

$$\bar{\beta}^{\tau^*} \leq \beta_1 < \bar{\beta}^{\tau^*+1} \text{ if } \beta_1 \leq \bar{\beta}^{\bar{\tau}},$$

and is equal to $\bar{\tau}$ if $\beta_1 > \bar{\beta}^{\bar{\tau}}$.

We end this section with the following observation.

Observation 1.2. *The supervisor weakly prefers a policy of delayed honesty to immediate honesty when delayed honesty is implementable, i.e. when $\beta_1 \leq \bar{\beta}^{\bar{\tau}}$.*

Ali (2017) derives the same result when determining the optimal dynamic disclosure policy in a slightly different environment. In his setting, the agent needs two consecutive successes in order to be successful in the project. The experiments yield success with a positive probability only if the project is of a good type. Ali shows that the more informed party always has an incentive to delay information revelation while the less informed party would prefer early

revelation. While we do not solve for the optimal policy here, we showed here delaying may be preferred by the supervisor to immediately revealing the outcome.

For priors above $\bar{\beta}^T$, a combination of immediate honesty and delayed honesty may be preferred by the supervisor. The prospect of finding out the outcome of experimentation immediately after experimenting makes the agent assess future costs probabilistically. Since it might be determined immediately that the last idea had a high potential to succeed, the agent then does not have to bear future costs of experimenting. This reduces the expected cost of experimentation to the agent and makes her willing to experiment. So for higher beliefs, where the agent is not willing to pay a lump sum cost for experimenting with delayed honesty, the supervisor can induce experimentation with immediate honesty. The supervisor can then commit to delayed honesty when the agent reaches a lower belief. However, immediate honesty might provide too much incentive to the agent and the supervisor might do better by committing to a mixed revelation for high beliefs.²⁴

²⁴We do not consider these policies in this paper as our primary objective is to highlight the tensions when the supervisor does not have commitment power. We merely want to show how the supervisor can do better when there is commitment in the relationship, and what incentives shape a “preferred” policy.

Chapter 2

The Newsroom Dilemma

2.1 Introduction

On April 18, 2013, the *New York Post* plastered its cover page with a picture of two men under the headline “BAG MEN: Feds seek these two pictured at Boston Marathon.” The Post was hinting that the duo was responsible for the Boston Marathon bombings and had carried the bombs in their bags. They were innocent, and the Post was wrong. 16-year-old Salaheddin Barhoum and 24-year-old Yassine Zaimi later filed a lawsuit, and the New York Post’s credibility was damaged. Similarly, in September 2008, *Bloomberg* incorrectly reported that United Airlines was filing for bankruptcy. Before Bloomberg issuing a correction, United Airlines’ stock price nosedived 75 percent.

Media critics often cite such examples to argue that competitive pressures in the modern digital environment have pushed outlets towards early release of less accurate information (Cairncross, 2019).¹ Matt Murray, Editor-in-Chief of the *Wall Street Journal*, acknowledged in a recent interview that the Internet had created both time and competitive pressures. However, part of the pressure, he noted, “is to stay true to what has worked and works (really) well, which is reporting verified facts.” In a similar vein, some media scholars argue that the fears surrounding the effect of competition may be overblown (Knobel, 2018; Carson, 2019).

In this paper, we discuss why competition among media outlets might not privilege speed over accuracy. We consider the implications of competition on audience welfare and information dissemination. We argue that two opposing forces determine the resolution of the speed-accuracy tradeoff: preemption and reputation. While preemption pushes outlets towards speed, reputation gives media outlets a reason to engage in careful, detailed reporting.

¹This, of course, is a cause of concern for modern democracies. Media outlets, through fact-checking and investigative journalism, deliver revelations that have a profound impact on the society and its institutions. For instance, *The Hindu*’s Bofors scam exposé in India in 1987 brought the topic of political corruption to center stage and lead to the defeat of the government in power in 1989. More recently, the *New York Times*’ exposé on sexual abuse in Hollywood and corporate America has reignited discussions on gender discrimination in the workplace.

We build a two-period model in which two career-concerned media outlets compete against one another and fear preemption. There is a topic on which the outlets may publish stories. Both outlets receive an initial informative signal about the topic. They may research the topic further at a cost, which depends upon their ability. We model research as generating a perfectly-informative signal about the topic. There is a scoop value associated with being the first outlet to publish a story on the topic. In addition to valuing scoops, outlets also care about their reputations. Reputation depends upon an audience's inference about the outlet's ability to research.

Our model yields three main results. The first two speak to the changes in the media landscape brought about by the Internet. The last result deals with how a source disseminates information to media outlets facing the speed-accuracy tradeoff.

Effect of the Internet. One effect of the Internet has been to increase competitive pressures. The Internet has reduced barriers to entry and contributed to a 24-hour news cycle where reporters are always on deadline. Consequently, pressure on media outlets to be the first to publish have increased.

In our model, while competition can push media outlets to publish more quickly, it can also have the opposite effect – to push outlets to research stories more thoroughly. We find that in more competitive environments, it is easier for outlets to build reputation. This effect increases outlets' willingness to hold back on stories and research them thoroughly. Importantly, our argument relies upon the assumption that the audience does not observe the amount of time outlets spend researching stories but they do observe which outlet publishes first. Knowing the sequence of publication rather than the amount of research, allows for additional observational learning with competition. Consequently, it gives better outlets a reason to differentiate when facing competition.²

We show that when there is a high scoop value, competition drives media outlets to publish more quickly; in contrast, when there is a low scoop value, competition drives media outlets to research stories more. Therefore, the model suggests that breaking news-type stories such as those on terrorist attacks, malfeasance of senior government officials or adverse economic shocks, will suffer particularly from problems of accuracy in the Internet age. In contrast, outlets do better research on non-urgent stories that do not influence immediate decision-making. Examples include: revelations of sexual abuse by Hollywood executives, how terrorist organizations work, and illegal data hacking that is used to influence public opinion.³

²We discuss in detail the new media studies literature in Section 2.1.1 and show some anecdotal support for our main finding.

³The first story was published in both the *New York Times* and the *New Yorker*. <https://www.newyorker.com/news/news-desk/from-aggressive-overtures-to-sexual-assault-harvey-weinsteins-accusers-tell-their-stories>.

A second effect of the Internet has been to improve what quickly-released stories look like. Journalists can quickly “contact people, access government records, file Freedom of Information Act requests, and do searches” (Knobel, 2018). Similarly, Chan (2014) argues that “digitization brings better access to sources and data.” At the same time, however, the cost of doing in-depth research has not changed much. For instance, one would not expect the cost of conducting interviews and building trustworthy sources to have changed significantly. We model such an effect as improving the quality of the initial signal without changing the cost of research.

We find that a *better* initial signal can *reduce* the welfare of the audience. When initial signal becomes better, the audience is less able to attribute correct information by the media outlets to their ability to conduct in-depth research. The audience instead assign it to better initial signal of the outlets that is due to better technology. Thus, reputational concerns get diluted and timing pressures become more salient, making the media outlets move towards speed. Moving towards speed reduces overall welfare only if a significant proportion of audience values better reporting. However, it improves welfare if the audience does value early reporting. It is easy to map the above examples from the previous paragraphs to the relevant situation for audience welfare.

Information dissemination by a source. Our model is also useful for determining how a strategic source shares its information with competing media outlets. Notably, it helps explain why politically-motivated sources may share rumours with multiple outlets to get “unverified facts” out to the audience.

Our model predicts that a source who is merely interested in getting potentially incorrect information out without further research can exploit the time pressures that competing media outlets face. We show that when media outlets are intrinsically driven to explore issues, it is better to share information with all the media outlets to get the information out quickly. More intrinsically motivated media outlets are more likely to do further research independent of the competition. However, by sharing with all the media outlets and creating competitive pressures, additional time pressure can be created. Thus, politicians with propaganda may still hold media outlets hostage even without explicitly capturing or buying them off.

There are, however, situations when such a source shares information only with one media outlet for the quick release of information. The source is likely to do so when media outlets are not intrinsically motivated, and the information is not urgent. When the information

The second story appeared on the *New York Times* after the reporters researched for more than a year and a half. <https://www.nytimes.com/interactive/2018/04/04/world/middleeast/isis-documents-mosul-iraq.html>. The third story broke out in *The Guardian*. <https://www.theguardian.com/us-news/2015/dec/11/senator-ted-cruz-president-campaign-facebook-user-data>.

is not urgent, there is a general tendency for competing media outlets to investigate further independent of their intrinsic motivation. In this situation, sharing information with just one outlet gets the information out more quickly.

It is worth emphasizing that our model generally covers settings that have elements of preemption and career concerns. For instance, competing researchers working to solve similar problems and hoping to convince a market about their ability face a similar newsroom dilemma. Technology firms face a speed-accuracy tradeoff as they build products and technology to match consumer preferences. Our main results have a natural interpretation in these situations. Notably, better research in competitive environments requires that the initial research idea is not too well-developed.

2.1.1 Stylized facts and new media studies literature

The speed vs. accuracy tradeoff is commonly recognized in the media studies literature. The BBC Academy website observes that “every journalist has to resolve the conflicting demands of speed and accuracy. [...] If you are working on a breaking news story, it is important to remember that first reports may often be confused and misleading. [...] That is why it is important to weight the facts you have.”

The terms of this tradeoff hinge on the surrounding environment. The literature highlights two critical determinants of the rise of “speed-driven journalism” in the modern digital environment. The first one is increasing competitive pressure. Lionel Barber, the Editor of *Financial Times*, points out, “Technology has (also) flattened the digital plain, creating the illusion that all content is equal. It has made it possible for everyone to produce and distribute content that looks equally credible”. Thus, outlets cannot only count on their pre-existing reputation to attract readers, and being the first to break the news is increasingly important. Rosenberg and Feldman (2008) note, “Why do experienced journalists telecast unscreened material in volatile situations? Because they can, and because they are driven by powerful, rush-to-report hard instinct, the one commanding them to beat or at least keep astride of the competition and not be left behind”.

The second is the 24-hour news cycle (Lee, 2014; Starbird, Dailey, Mohamed, Lee and Spiro, 2018), which leads to the possibility of being preempted at any point in time. Newspapers used to have editions making it possible to verify information up until the night before publication, almost without fear of someone else breaking the news. That is no longer the case. As Howard Kurtz from *Washington Post* describes, “In the last year, the pendulum has swung in our newsroom to putting things on the Web almost immediately [...] everybody wants it now-now-now. [...] But the sacrifice (clearly) is in the extra phone calls and the chance to

briefly reflect on the story that you are slapping together” (Rosenberg and Feldman, 2008).

Importantly, however, reputational concerns remain relevant. *Reuters Handbook of Journalism* states “Reuters aims to report facts, not rumors. Clients rely on us to differentiate between fact and rumor, and our reputation rests partly on that”. Note that reputation is based on the ability to check the facts before releasing them, which is also how we model it. Knobel (2018) summarizes her interviews with the editors by saying that they realize that readers can be induced to pay for quality journalism. She quotes Rex Smith, editor of the *Albany Times Union*, “What can separate great journalism from everything else is our commitment to the journalism of verification and watchdog reporting. It will give us credibility that other organizations do not have.”

Some new literature from media studies paints a more positive image of the future of watchdog reporting. While not exactly the same as reporting accurate stories, watchdog reporting, which includes investigative journalism and fact-checking, takes time. We show here the data from Knobel’s study in support of our theoretical results in Table 2.1. The table shows an increasing share of accountability reporting among a sample of 9 US newspapers for 1991-2011. 2001 in her sample marks the year that the Internet and social media took off in a big way, and became an essential source of news for the audience. We can see how almost all newspapers have increased their accountability reporting since then. The increase is visible for both deep and simple accountability reporting, and across newspaper groups. While the increase may be due to several reasons, her data together with the interviews hint at similar path to that which we outline in this paper.

The model we build tries to combine these insights into a unified analysis of the speed-accuracy tradeoff and the competing forces that determine its direction.

2.1.2 Contributions to related economics literature

We primarily contribute to the literature on media competition and quality of news by explicitly modeling the newsroom dilemma. The newsroom dilemma, or the speed-accuracy tradeoff, is surprisingly understudied in the field despite agreement among media scholars on its importance. One exception is Andreottola and De Moragas (2017). They look at the political economy impact of a similar speed-accuracy tradeoff and find that competition leads to a release of less accurate information. Our paper differs because we explicitly model the reputational concerns of media outlets. We identify conditions where the additional information transmitted by the presence of competitors overcomes the preemption concerns.

Some new literature has also started exploring theoretically the effect of the Internet on the media landscape. In Angelucci and Cagé (2019), for instance, the authors show that an

Internet-driven drop in the advertising revenues leads to a smaller newsroom, decrease in prices and a move towards “soft” information. Similarly, [Armstrong \(2005\)](#) looks at the relative effect of advertising-only with a subscription-based funding mechanism on journalistic quality. All of these papers and others ([Ellman and Germano, 2009](#); [Gentzkow, 2014](#)) build on two-sided market models ([Rochet and Tirole, 2003, 2006](#)) and are concerned with pricing decisions. We do not explicitly model advertising and pricing concerns. We instead subsume them under either preemption or reputational concerns.

Some recent papers that do not look at pricing explicitly but explore the political consequences of new media or of media competition are [Sobbrío \(2014\)](#), [Allcott and Gentzkow \(2017\)](#), [Barrera, Guriev, Henry and Zhuravskaya \(2017\)](#), [Chen and Suen \(2016\)](#), [Perego and Yuksel \(2018\)](#) and [Vaccari \(2018\)](#). For instance, [Chen and Suen \(2016\)](#) look at media competition and endogenous attention allocation. They do not have the speed-accuracy tradeoff, but they show that increased competition reduces outlets’ investment in reporting quality, increasing the overall influence of the media industry. [Perego and Yuksel \(2018\)](#) and [Vaccari \(2018\)](#) look at the distortive effect of competition on information provision by biased outlets, and as a consequence on the level of information voters can acquire. In both cases, competition can increase distortions. In this paper we abstract from outlets’ political motivations to focus on their incentives to provide good quality journalism.

Focusing more on the reputation-building and signaling in media markets, [Gentzkow and Shapiro \(2006\)](#) model media bias and reputation building, showing that competition reduces bias. The model explores an entirely different tradeoff looking at the content of the reporting directly, rather than the timing. Also, the “positive” effect of reputation comes from a different channel – with competition the reader is more and more likely to learn the actual state eventually. Our model does not have this feature as the audience eventually learn the actual story, and competition does not affect the revelation incentives in this way. [Gentzkow and Shapiro \(2008\)](#) later provide an outline of a model that may incorporate reputation-building incentives like ours but they do not consider preemption. [Shapiro \(2016\)](#) shows that reputational concern for unbiasedness may induce journalists to report evidence as ambiguous even when it is not. Preemption concerns and endogenous choice of research are not considered there.

Our modelling strategy shares some features with [Hafer, Landa and Le Bihan \(2018\)](#) and [Hafer, Landa and Le Bihan \(2019\)](#). Like us, they have a two period model where competing outlets can acquire information about a politically relevant state of the world and choose when to release it. However, we do not focus on media bias and on the possibility of claiming credit for a story, but rather on the trade off between time pressure and quality of journalism. See [Prat](#)

and Strömberg (2013) and Strömberg (2015) for recent developments in the political economy of media literature, and other related papers.

We also contribute to the literature on strategic information release. We differentiate from Guttman (2010) and Guttman, Kremer and Skrzypacz (2014) by adding reputational concerns and endogenizing the information acquisition choice. Therefore, our results are driven by completely different incentives. Relatedly, Aghamolla (2016) looks at a model of (anti-)herding between financial analysts with endogenous information acquisition. While observational learning is critical in such herding models, reputation building drives such incentives in our model. Observational learning is relevant for the audience in our model because it signals the type of the outlet. Gratton, Holden and Kolotilin (2017) look at a model in which a sender strategically releases a stream of information to influence perceptions about herself. They show that better sender types release the information earlier and expose themselves to scrutiny. This is in contrast with our model, where better outlets release information later. Preemption concerns drive the incentives in our model, which produces our different result.

Finally, we also contribute to the literature on preemption games and R&D races by adding reputational concerns. Preemption games have long been studied in economics (Fudenberg, Gilbert, Stiglitz and Tirole, 1983; Fudenberg and Tirole, 1985), but our paper contributes to the more recent literature on preemption games with private information (Hopenhayn and Squintani, 2011, 2015; Bobtcheff, Bolte and Mariotti, 2016). It is worth noting that Bobtcheff et al. (2016) have a similar “separating” result for different types of firms, but in a set up without reputation. Here we point out that reputation, combined with actions that partially reveal an opponent’s type, can be a different force leading to separating strategies in preemption games.

2.2 A model of the newsroom dilemma

We build a simple two-period model indexed by $t = 1, 2$ featuring three players: two strategic media outlets i, j and a fixed mass of audiences. We also consider a version with just one media outlet.

State of the world. The state of the world ω is binary and unknown to the players. Formally, $\omega \in \Omega := \{a, b\}$ with common prior $\Pr(\omega = a) = \frac{1}{2}$. Ω pertains to the topic on which the media outlets are digging a story, and the relevant information for the audience. This could be, for instance, who is responsible for a terrorist attack, whether a senior government official is involved in corruption or not, who is an appropriate candidate to vote for in the election, etc.

Media outlets. Initially, each outlet privately observes a signal s^i about the state of the world in $t = 1$. We call this the story that the outlets have. We assume that s^i is free and i.i.d. conditional on the state. Its precision is $\Pr(s = \omega | \omega) = \pi > \frac{1}{2}$. The two outlets decide simultaneously

at this stage whether to publish their signals, or conduct further research. The decision d^i for outlet i in $t = 1$ is, therefore, to choose from $\{pub, res\}$ where pub is publish immediately, i.e. in $t = 1$, and res is do more research and then publish in $t = 2$.

Publishing is equivalent to endorsing a particular state of the world (independent of whether published in $t = 1$ or 2). When an outlet publishes its story it sends a message $m \in M = \{\tilde{a}, \tilde{b}\}$ where $\tilde{\omega}$ means endorsing state ω .

Conducting further research (and then publishing in $t = 2$) is costly. In particular, there is a type specific cost of research that perfectly reveals the true state of the world in $t = 2$. Outlets can be of two types, high or low quality, depending on how efficient they are at digging into stories, and this is the private information of each individual outlet. Formally, the type of outlet i is $\theta^i \in \{h, l\}$ with a common prior $\Pr(\theta^i = h) = q = \frac{1}{2}$. The types are independent.

$\theta = l$ faces an infinite cost of conducting research. The low quality outlet never digs stories further and chooses $d = pub$ in $t = 1$. The cost c for the high quality outlet is private information of that outlet, and is story-specific. It comes from a uniform distribution F with support $[-\varepsilon, \bar{c}]$ and is drawn independently for each high quality outlet. ε is greater than zero but small to capture the idea that some high quality outlets may still want to conduct research even in the absence of other rewards.⁴ We assume $\bar{c} \geq 2$ so that the support of the distribution F is sufficiently large.

Finally, the assumption on q is just for analytic convenience. A generic $q \in (0, 1)$ would not alter the results, qualitatively. We show this case in Appendix C.

Audience. The audience enters the game when one or both of the outlets publish their story, and their story is revealed (i.e. m). They only rationally form beliefs about the types of the outlets. They enter with the knowledge of the priors and an understanding of the competition between the outlets. Other than this, the precise information of the audience at the time of belief formation is denoted by the set \mathcal{I} .

We assume that the audience observes the sequence of publication but not the actual time of publication, or whether the outlets conducted research. The sequence, as distinguished from the timing, shows whether the outlets moved sequentially or simultaneously. Under this assumption, the audience will be able to determine the actual time of publication (i.e. $t = 1, 2$) only if the outlets moved sequentially. It can be summarized by $\tilde{t}^i \in \{I, II, \emptyset\}$, which shows whether outlet i was first, second, or it moved simultaneously with j . This assumption is

⁴Interviews with editors often confirm such motivations; often they feel a sense of responsibility in their positions. For instance, Knobel quotes Marcus Brauchli, *Washington Post*'s former editor, "Doing investigative journalism is in the *Post*'s DNA and has been as long as any of us have been around in journalism." Similarly, Kevin Riley, the Editor of the *Atlanta Journal-Constitution* explains, "People want us to do this. They don't think anyone else will if we don't."

discussed in more details in Section 2.2.2 and its implications are described in the main analysis 2.3.

In addition, after both the outlets publish their stories, the state is revealed exogenously. If $m^i = \omega$, then outlet i is said to be right, or R . Otherwise, the outlet is wrong, denoted by W . We call this the outcome O of verification. The audience sees the outcome. Therefore, the information of the audience \mathcal{I} at the end of the game is denoted by a tuple (O_t^i, O_t^j) that consists of four pieces of information, i.e. the position of each outlet in the sequence of publication and their outcomes. Using \mathcal{I} , the audience updates its beliefs about each outlet's type. Denote the posterior belief about $\theta = h$ by $\gamma(\mathcal{I})$ when the information held by the audience is \mathcal{I} .

Payoffs. Currently, we do not illustrate the payoffs of the audience as they only form beliefs. We will, however, place more structure on its preferences at a later stage and explain the source of outlets' payoffs. For the time being, we only focus on the outlets' payoffs, which are composed of three elements.

1. The first is a scoop value v to the first outlet publishing the story. It captures the preemptive nature of the media market, highlighted for example in Besley and Prat (2006). v can be interpreted as the mass of audience that is drawn to the first media outlet breaking the story.
2. The second is a reputation value of γ^i or the audience's posterior on the quality of outlet i calculated after revelation of the true state. This captures the extent to which the outlets care about their reputation. For instance, future audience of the outlets might depend on their reputations. We assume that reputation enters linearly in the outlets' payoffs. Importantly, the audience cares about whether the outlet is high or low type, not about c . A new c is drawn for every new story and only the high type has the ability to conduct further research.
3. The third is the cost c that the high type outlet chooses to pay if it does research in period 1.

Timing. The timing of the game can now be summarized as follows:

0. Nature draws ω , θ^i and θ^j . θ is privately observed by each outlet. ω is unobserved.
1. At $t = 1$ each outlet privately observes s^i . A cost c of digging into the story is drawn from a uniform distribution $F[-\varepsilon, \bar{c}]$ for the high type.
2. The outlets simultaneously decide $d^i \in \{pub, res\}$ and if $d^i = pub$ then also choose m . As stated before, this is a relevant decision only for the high type. The low type always

chooses *pub*.

3. If both outlets publish, the game ends. Otherwise, the game goes to period 2.
4. At $t = 2$, the state is revealed to every outlet that chose $d^i = res$. Those who did not publish in $t = 1$, publish now by choosing m .
5. Once both the outlets have published, the state ω is revealed to the audience. They observe \mathcal{I} and update beliefs on the type of each outlet. Payoffs are realized.

2.2.1 Solution concept and equilibria selection

The solution concept we use is the Perfect Bayesian Nash Equilibrium in pure strategies. We focus on equilibria where outlets optimally follow the signal they receive, i.e they endorse the state that is more likely to be the true one given their signal. We call such equilibria *signal-based equilibria*.⁵ For the rest of the paper, we use “equilibrium” and “signal-based equilibrium” interchangeably.

2.2.2 Discussion of assumptions

Before proceeding to the analysis, it is worth discussing our assumptions in detail.

The first assumption we make is regarding what the audience observes about the timing of the game. The fact that the audience only observes the content of what was published (i.e, m) and the sequence of publication (i.e., \tilde{t} but not the actual t) captures the idea that it is unaware of how much the outlets researched story. We believe this is a realistic assumption in that the amount of research is hardly observable from outside the newsroom. Of course, the amount of research conducted maps in a probabilistic way into the accuracy of a story, which the audience can check more easily. We allow for such a possibility by letting the audience observe whether the story is true or false.

The important consequence of this assumption is that player i 's decision to publish/not publish can potentially convey information about player j 's type. For example, if the two outlets move sequentially and only a high type is expected to conduct research, moving later is a signal of the first outlet being a low type. We show how relaxing this assumption changes our result in Section 2.3.4.

The second assumption we make is about who possesses stories on a topic. In reality, competing media outlets are often unaware of whether their competitors are also exploring the same story. We assume that both of the media outlets are aware that their competitor also

⁵This means that we ignore equilibria where outlets choose to endorse one particular state to signal their type. Those equilibria may exist, but we argue that they do not make much sense given the environment we are considering. Alternatively, we can assume that signals are hard information, but the reader cannot infer the level of precision: the result would be exactly the same.

possesses the story. Doing so pushes the incentives of the outlets the most towards speed. Still, we show that more research is possible under competition. Including such a possibility further adds to the complications of the model.

The third assumption we make is that outlets build a reputation on their consistent types, and not on the cost of digging into each independent story. Given that different outlets usually have different expertise, it is reasonable to assume that they face different costs when exploring different stories. For instance, *The Wall Street Journal* is a business-centric daily and has invested in building sources and methods for dealing with business stories (such as avoiding lawsuits when potentially sensitive corporate information is published). However, in general, some outlets have a culture of researching while others do not. Their type captures this.

We also make a few assumptions for tractability reasons. First, we do not allow for the outlets to “sit on information” or wait for a period before publishing.⁶ Second, we assume that the audience correctly finds out the state at the end of the game. Third, we assume that the media outlet correctly finds out the state upon choosing to research. Almost all of these assumptions can be relaxed to some degree without altering our predictions.

Finally, it is worth emphasizing that there is an informational value of the news to the audience in our model. They want to know the actual state, which allows them to make decisions or form opinions. Our model, therefore, does not deal with the entertainment value of news where the audience enjoys getting the information.

2.2.3 Preliminary observations and strategies

We start with a few simplifying observations. All the proofs are in Appendix A.

Observation 2.1. *If an outlet decides to publish in $t = 1$, it follows its signal s , i.e. sends $m = s$. If an outlet decides to do research and then publish in $t = 2$, it follows the outcome of research.*

Observation 2.1 follows from the fact that in $t = 1$ the most informative signal is s . Therefore, the most likely state is the one given by the signal. This is a standard result in this type of environment and follows from the flat priors on the state. Moreover, in $t = 2$ the outlet choosing to research has learned the actual state and therefore, publishes it (independent of what the original signal s stated). Thus, as long as there is a gain in matching the state, each outlet follows its last signal, which is also the most informative one.

There is also a useful result arising from our particular signal structure and flat prior over the state.

Lemma 2.1. *If each outlet follows its last signal when publishing, the following results hold:*

⁶We can show that for a sufficiently high v and relevant off-path beliefs, the outlets never choose to wait.

1. *The probability of matching the state after only s is π .*
2. *Regardless of whether i decides to publish or research, from its point of view the expected probability of player j matching the state without research is π .*

Lemma 2.1 will be helpful in writing the incentive compatibility conditions for the players. Doing so will require each outlet to consider whether the other will do research and the subsequent probability of matching the state.

It is useful to define precisely the strategies we will focus our attention on. Note first that the only relevant and meaningful decision that deserves our attention is the one of the high type outlet in period 1. From the the outlet's point of view, this will be a threshold strategy where the threshold is defined on the cost c of research. The high outlet conducts research if the realized cost c is less than some threshold c_D (where subscript D represents the case of a two firm duopoly).⁷ From the other outlet's (and the audience's) point of view, define σ^i , the conjectured probability that outlet i chooses to research further in $t = 1$, conditional on outlet i being a high type. Therefore,

$$\sigma^i = \Pr(c \leq c_D) = F(c_D) = \begin{cases} 0 & c_D < -\varepsilon \\ \frac{c_D + \varepsilon}{\bar{c} + \varepsilon} & -\varepsilon \leq c_D \leq \bar{c} \\ 1 & c_D > \bar{c} \end{cases}$$

We are now ready to move to the equilibrium analysis arising in different market configurations.

2.3 Competition leads to better reporting

2.3.1 Newsroom dilemma with a single firm: Monopoly

Let us start with the simplest case: there is a single media outlet and its type is known to the audience.

Proposition 2.1. *If there is one media outlet and θ is known to the audience, then the high quality outlet conducts research with probability $F(0) = \frac{\varepsilon}{\bar{c} + \varepsilon}$.*

In this case, none of the aforementioned incentives are at play. There is obviously no preemption risk and there is nothing to do in terms of reputation. Every type of outlet gets $v + \mathbb{1}\{\theta = h\}$ so it is pointless to pay any cost for researching. The outlet is driven to research only because of its intrinsic motivation.

⁷Similarly, the case of single firm monopoly is denoted by a threshold c_M and in general, by a subscript M .

The case of monopoly with unknown type is more interesting. Proposition 2.2 summarizes the main result.

Proposition 2.2. *If there is one media outlet and θ is not known to the audience, there exists a unique equilibrium in which the high quality outlet conducts research in $t = 1$ iff*

$$c \leq (1 - \pi)(\gamma(R) - \gamma(W)) := c_M$$

where $\gamma(R)$ and $\gamma(W)$ are the audience beliefs about the outlet's quality after it gets the state right and wrong respectively. As a consequence, $\sigma^* = F(c_M) = \frac{c_M(\sigma^*) + \varepsilon}{\bar{c} + \varepsilon}$.

Note that the preemption risk is absent in this case as there is only one outlet. v does not play any role in the threshold above. But a high outlet is incentivized to do research to build a reputation for being a high quality. However, this reputation cannot be based on the observation of sequence or timing. The only relevant thing that the audience observes is whether the outlet is right or wrong, i.e. whether $m = \omega$ or not after the state is verified. Therefore, if the high outlet is expected to choose to research with probability σ , the two relevant belief updates are

$$\gamma(R) = \frac{\sigma + (1 - \sigma)\pi}{\sigma + (1 - \sigma)\pi + \pi} \text{ and } \gamma(W) = \frac{(1 - \sigma)(1 - \pi)}{(1 - \sigma)(1 - \pi) + (1 - \pi)} = \frac{1 - \sigma}{2 - \sigma}$$

from Bayes' rule. Now, it must be that the additional cost c of choosing $d = res$ is more than compensated by the expected reputational gains from endorsing the correct state. This is captured in c_M of Proposition 2.2.

2.3.2 Newsroom dilemma with two firms: Duopoly

The main effect of competition is the introduction of preemption risk. When preemption is relevant and reputation building is not, then the equilibrium where the high quality outlet conducts research becomes even rarer than in Proposition 2.1. Proposition 2.3 below highlights this.

Proposition 2.3. *If there are two media outlets and θ is known to the audience, there exists a unique symmetric equilibrium in which the high quality outlets conduct research with probability $\sigma_D^* = F\left(-\frac{v}{2}\right)$.*

Intuitively, there is nothing to gain from conducting research in terms of reputation as θ is known. The only reason to investigate further is if there is an intrinsic motivation to do so. But now there is a preemption risk that reduces the incentives to investigate. However, if v is sufficiently small relative to the intrinsic motivation (i.e. if $v < 2\varepsilon$), there will still be some high outlets willing to investigate.

The case of competition plus hidden types is the most interesting one. In this case, both the preemption and reputation building concerns are simultaneously relevant and interact with each other. Before we present the key proposition, we discuss how the audience updates beliefs in this environment. Recall that the audience observes both the outcome of verification $\mathcal{O} \in \{R, W\}$ and the sequence of publication $\tilde{t} \in \{I, II, \emptyset\}$ for both i and j . Suppose now that a high quality outlet chooses to research with probability σ^i . Then, for a given conjectured level of σ^i and σ^j , the relevant audience's on-path beliefs need to be defined for the following events:

$$\{(R_\emptyset, R_\emptyset), (R_\emptyset, W_\emptyset), (W_\emptyset, W_\emptyset), (W_\emptyset, R_\emptyset), (R_I, R_{II}), (W_I, R_{II}), (R_{II}, R_I), (R_{II}, W_I)\}$$

where the first outcome-sequence element in each information set is outlet i 's and the second is outlet j 's.⁸

It can be shown that there are three relevant set of events for belief updating. The first is when both the outlets get the state correct and they publish simultaneously.

$$\gamma^i(R_\emptyset, R_\emptyset) = \frac{\sigma^i \sigma^j + (1 - \sigma^i)(2 - \sigma^j)\pi^2}{\sigma^i \sigma^j + (2 - \sigma^i)(2 - \sigma^j)\pi^2} := \gamma^i(\emptyset)$$

Here the audience is unable to determine the actual timing of publication. It cannot distinguish as to whether both conducted research (which happens only if both are high types) or both published immediately (either because they are both low types, or because there is only one high type and it faced a high cost, or because both are high types but they faced high costs). With some abuse of notation, we denote the updated belief under “no information about timing” event by $\gamma(\emptyset)$.

The second is when the audience is able to determine that outlet i moved in $t = 1$.

$$\gamma^i(R_\emptyset, W_\emptyset) = \gamma^i(W_\emptyset, W_\emptyset) = \gamma^i(W_\emptyset, R_\emptyset) = \gamma^i(R_I, R_{II}) = \gamma^i(W_I, R_{II}) = \frac{1 - \sigma^i}{2 - \sigma^i} := \gamma^i(1)$$

This, of course happens when i moves first and j moves second (independent of whether i gets the state correct or not). But the audience is also able to understand it when the outlets move simultaneously and at least one of them gets the state incorrect (since researching further perfectly reveals the state). Here the only uncertainty for the audience is whether the outlet is a high quality one that faced a high cost or a low quality one. We denote the updated belief under the “published in period 1” event by $\gamma(1)$. Observe how in these events the presence of a

⁸Note that it never happens that an outlet moves second in the sequence and gets the state incorrect. Any outlet that moves second has conducted research and matches the state perfectly. Therefore, any event with W_{II} does not occur on-path.

competitor conveys to the reader some additional information about the type of each outlet.

Finally, the third is when the audience is able to determine that outlet i moved in $t = 2$.

$$\gamma^i(R_{II}, R_I) = \gamma^i(R_{II}, W_I) = 1 := \gamma^i(2)$$

This only happens when outlet i moves second and gets the state correct, which in turn is only possible if it is a high quality outlet. Therefore, the updated belief under “published in $t = 2$ ” event is $\gamma(2) = 1$.

Using these updated beliefs, a high quality outlet’s incentive compatibility can be written as follows. For any given conjectured σ^j and audience’s beliefs, a high quality outlet i with cost c^i chooses to research further if

$$\overbrace{\frac{1}{2} \left[\sigma^j \left(\frac{v}{2} + \gamma^i(\emptyset) \right) + (1 - \sigma^j) \gamma^i(2) \right] + \frac{1}{2} \gamma^i(2) - c^i}_{\text{expected payoff from research}} \geq \underbrace{\frac{1}{2} \left[\sigma^j (v + \gamma^i(I)) + (1 - \sigma^j) \left(\frac{v}{2} + \pi^2 \gamma^i(\emptyset) + (1 - \pi^2) \gamma^i(1) \right) \right] + \frac{1}{2} \left(\frac{v}{2} + \pi^2 \gamma^i(\emptyset) + (1 - \pi^2) \gamma^i(1) \right)}_{\text{expected payoff from publication}}$$

which further simplifies to

$$c^i \leq \frac{1}{2} \left[(\gamma^i(\emptyset) - \gamma^i(1)) (\sigma^j - (2 - \sigma^j) \pi^2) + (2 - \sigma^j) (1 - \gamma^i(I)) \right] - \frac{1}{2} v := c_D^i \quad (2.1)$$

Proposition 2.4 then follows:

Proposition 2.4. *If there are two media outlets and θ is not known to the audience, there exists a unique and symmetric equilibrium where $\sigma^{i*} = \sigma^{j*} := \sigma^* = F(c_D)$ such that*

$$c_D = \frac{1}{2} \left[(\gamma(\emptyset) - \gamma(1)) (\sigma^* - (2 - \sigma^*) \pi^2) + 1 \right] - \frac{1}{2} v$$

where $\gamma(\emptyset) = \frac{(\sigma^*)^2 + (1 - \sigma^*)(2 - \sigma^*) \pi^2}{(\sigma^*)^2 + (2 - \sigma^*)^2 \pi^2}$ and $\gamma(1) = \frac{1 - \sigma^*}{2 - \sigma^*}$.

Looking now at the cost threshold c_D of Proposition 2.4, we can see the negative effect of v . If preemption concerns are very salient (i.e. v is high), then separation happens for a smaller range of c making research less likely. On the other hand, the positive side of the condition is given by the expected reputational gains of matching the state (and publishing second).

2.3.3 Competition may lead to better reporting

The comparison between monopoly and duopoly when reputation building is relevant (Propositions 2.2 and 2.4) provides interesting insights.

Lemma 2.2. *The reputational gains are always higher in duopoly than in monopoly.*

The reason lies in the availability of additional information in the case of duopoly. First, the presence of two outlets allows the audience to compare their contents, i.e. which states outlets i and j endorsed. Second, it allows the audience to observe the sequence of publication of the two outlets. Together, these two factors allow outlet i to publish after outlet j , match the state correctly, and signal its type more easily. In turn, this makes outlet i more willing to pay the cost of research. However, the additional preemption concerns in duopoly counterbalance this positive information effect, and makes c_D decreasing in v . The two effects combined yield our first main result pertaining to the effect of Internet-driven competition on reporting.

Proposition 2.5. *There exists a nonempty interval of v values where $\sigma_D^* > \sigma_M^*$.*

Basically, what Proposition 2.5 says is that there is a nonempty set of parameters where research is more likely in duopoly than in monopoly. Therefore, competition may lead to better reporting.

A good way to illustrate Proposition 2.5 is Figure 2.1. The orange line is $F(c_D)$, the green line is $F(c_M)$ and the blue one is the 45° line. The equilibrium probability of research is given by the point of intersection of $F_c(c_D)$ and $F_c(c_M)$ with the 45° line. It is clear that $\sigma_D^* > \sigma_M^*$ for sufficiently small v .

Intuitively, reputational gains in monopoly are given by the increased probability of getting the state right. In duopoly, the audience can use one extra piece of information – the action of the other outlet, which includes the outcome of verification and the sequence of publication. Hence, competition induces a trade off between those two forces pushing in opposite directions. Importantly, this trade off is not obvious. The main point of Proposition 2.5 is precisely to point out that, contrary to the wisdom of the crowd in media studies literature, competition does not necessarily lead to a faster release of less accurate information.

2.3.4 The role of audience's information

The previous results relied critically on what the audience observes from the competition, or simply the “transparency”. To build further intuition, here we analyze how changing the transparency affects our result. In general, the effect of transparency on the possibility that competition induces better reporting is non-monotonic. To see why, consider the two other possibilities – nothing about the timing is observable and the timing of research is fully observable. Our original assumption lies in the middle of this increasing transparency spectrum. Of course, the content of publication is always visible to the audience, i.e. the audience observes m .

Unobservable timing or zero transparency. Suppose the audience observes neither the timing of publication nor the sequence of publication. It simply consumes the content of the outlet publishing the story. In this case, the behavior of the monopolist is exactly as before. Hence, $c_M = (1-\pi)(\gamma(R)-\gamma(W))$ does not change. In the case of duopoly, however, the endorsement of the other outlet does not matter anymore in the updating. The audience considers each outlet separately because nothing about the timing is observed. Therefore, $\gamma(R, \cdot) = \gamma(R)$ and $\gamma(W, \cdot) = \gamma(W)$. The consequence is summarized in the following corollary.

Corollary 2.1. *If neither time nor the sequence of publication are observable,*

$$c'_D = c_M - \frac{1}{2}v$$

and therefore, $c'_D < c_M$ for every strictly positive v .

Intuitively, there are no additional reputational gains because it is not easier to “look good” in the presence of a competitor. In fact, the reputational part of the cost threshold is exactly the same. But the additional risk of preemption pushes c_D down.

Observable timing or full transparency. If the timing of publication is observable, the monopolist can fully differentiate itself by publishing in period 2. This is possible because the audience can now perfectly distinguish between period 1 and 2, and therefore, is fully aware of whether research was conducted or not. Moreover, this is true in duopoly as well. In fact, the actual content of the publication does not matter for the reputation-building, and differentiation is driven entirely by the timing. As a consequence, the logic applies as before. The reputational part of the threshold is the same, but preemption concerns reduce the incentives to investigate and conduct research.

Corollary 2.2. *If the timing of publication is observable,*

$$c''_M = 1 - \gamma(1) \text{ and } c''_D = 1 - \gamma(1) - \frac{1}{2}v$$

where $\gamma(1) = \frac{1-\sigma}{2-\sigma}$. Therefore, $c''_D < c''_M$ for every strictly positive v .

Note that now the cost thresholds are bigger than in the previous information environments. This is so because now maximum distinction is possible between the two outlets. Therefore, the actual levels of reputational benefits are also higher. This is captured in the belief updating,

$$\gamma(1) = \frac{1-\sigma}{2-\sigma} \text{ but now } \gamma(2) = 1.$$

It is worth emphasizing that both of these extreme transparency assumptions are somehow problematic. Completely unobservable timing clashes with the idea of a scoop value, or more generally with the preemptive nature of the media market. If the audience has no understanding of when the publication happened, there is nothing to gain from being first. There are only gains from ultimate publication. This is obviously not true in reality. Completely observable timing, on the other hand, implies that the reader perfectly understands exactly how much research went into an article. Therefore, the whole differentiation happens on the time dimension, rather than on the truthfulness of the story. Again, this hardly seems true in reality.

2.4 Stories and the effect of better initial information

We are now in a position to discuss what kinds of stories are susceptible to more speed-driven journalism and what aren't. To do so, we place more restrictions on audience preferences.

Let there be a unit mass of audience. For any given story, a fraction u of this audience requires the information urgently. The audience seeks out the information because it has to take an action (for example, vote or form opinions). Let this action be denoted by $\alpha \in \{a, b\}$. Formally, if $\alpha = \omega$, then the urgent audience gets a payoff of 1, and 0 otherwise. u is story-specific and when the outlets get a story they also learn perfectly the value of u . The idea is that those stories with a relatively high u are more urgent than others. These could include, for example, information about whether a company has gone bankrupt, or whether the police caught the terrorists, etc. Non-urgent readers, instead, derive utility from high quality journalism. For simplicity, we assume that they derive utility of 1 from reading a well researched article, and that they acquire the content only if they are sure that it has been researched.

Given these preferences, the audience can pick its most preferred outlet. Type u audience always picks up the first outlet to publish the story, and $1 - u$ picks the second one. Therefore, u is akin to v , or the scoop value from the previous analysis. When both of the outlets publish simultaneously, then only u types are available. This audience chooses one of the outlets randomly. In addition, we assume that the entire mass of audience is available for reputation building.

First, observe that nothing changes relative to the monopoly case discussed in Proposition 2.2. As there is no sense of time order, the audience preference for urgency does not alter the equilibrium. However, now the duopoly case looks different. Noting that the belief updating

remains the same, the new condition for outlet i conducting research becomes

$$\overbrace{\frac{1}{2} \left[\sigma^j \left(\frac{u}{2} + \gamma^i(\emptyset) \right) + (1 - \sigma^j)(1 - u + \gamma^i(2)) \right] + \frac{1}{2}(1 - u + \gamma^i(2)) - c^i}_{\text{expected payoff from research}} \geq \underbrace{\frac{1}{2} \left[\sigma^j(u + \gamma^i(1)) + (1 - \sigma^j) \left(\frac{u}{2} + \pi^2 \gamma^i(\emptyset) + (1 - \pi^2) \gamma^i(1) \right) \right] + \frac{1}{2} \left(\frac{u}{2} + \pi^2 \gamma^i(\emptyset) + (1 - \pi^2) \gamma^i(1) \right)}_{\text{expected payoff from publication}},$$

which simplifies to

$$c^i \leq \frac{1}{2} \left[(\gamma^i(\emptyset) - \gamma^i(1)) (\sigma^j - (2 - \sigma^j)\pi^2) + (2 - \sigma^j) (1 - \gamma^i(1)) - \sigma^j(1 - u) \right] + 1 - \frac{3}{2}u := \bar{c}_D^i. \quad (2.2)$$

Observe in condition (2.4) that the incentives to research have increased. By researching and being the second one to publish the story, the outlet gets an additional $1 - u$ readers on top of building a perfect reputation. Said another way, this dilutes preemption concerns as both the first and the second mover have their respective markets. Therefore, we first need to check if a symmetric and unique equilibrium \bar{c}_D exists à la Proposition 2.4.

Proposition 2.6. *Let there be a fraction u of audience available to the first outlet publishing and let $\bar{c} \geq 2.5$. If there are two media outlets and θ is not known to the audience, there exists a unique and symmetric equilibrium where $\bar{\sigma}^{i*} = \bar{\sigma}^{j*} := \bar{\sigma}^* = F(\bar{c}_D)$ such that*

$$\bar{c}_D = \frac{1}{2} \left[(\gamma(\emptyset) - \gamma(1)) (\bar{\sigma}^* - (2 - \bar{\sigma}^*)\pi^2) - \bar{\sigma}^*(1 - u) \right] + \frac{3}{2}(1 - u)$$

where $\gamma(\emptyset) = \frac{(\bar{\sigma}^*)^2 + (1 - \bar{\sigma}^*)(2 - \bar{\sigma}^*)\pi^2}{(\bar{\sigma}^*)^2 + (2 - \bar{\sigma}^*)^2\pi^2}$ and $\gamma(1) = \frac{1 - \bar{\sigma}^*}{2 - \bar{\sigma}^*}$.

Note that while v was unbounded, $u \in [0, 1]$. But an increase in the fraction of urgent audience u still has a negative effect on \bar{c}_D and decreases $\bar{\sigma}^*$. Therefore, a high fraction of impatient audience pushes the outlets towards speed. The next proposition compares the probabilities of research in the no-competition monopoly case with the duopoly case on the basis of u .

Proposition 2.7. *There exists an interior $u, \bar{u} \in (0, 1)$ such that*

- for stories with $u < \bar{u}$, $\bar{c}_D > c_M$ so that research by high outlets in duopoly is more likely than in monopoly ($\bar{\sigma}_D > \sigma_M$);
- for stories with $u > \bar{u}$, $\bar{c}_D < c_M$ so that research by high outlets in duopoly is less likely than in monopoly ($\bar{\sigma}_D < \sigma_M$); and

- for stories with $u = \bar{u}$, $\bar{c}_D = c_M$ so that research by high outlets in duopoly is equally likely as in monopoly ($\bar{\sigma}_D = \sigma_M$).

We can therefore see that competitive environments are better for research on non-urgent topics. A good example is the recent *New York Times* exposé on sexual abuse in Hollywood. It is reasonable to believe that sexual abuse in the movie industry does not directly impact a large fraction of society. Yet, it was an important finding that will have a long-run impact as women come forward and demand justice, and organizations respond. On the flip side, investigations and research on urgent topics is less likely in competitive environments. The example of terrorist attacks fits perfectly in this setting. In fact, after the Boston Marathon Bombings in April 2013 there was much confusion in the media and articles were published without fact-checking. The intuition is simple: when a large fraction of the audience seeks information quickly, outlets compete to be the first one to publish the news.

We can now also make assessments about the audience's welfare.⁹ The audience's welfare V is defined as follows

$$\begin{aligned} V &= \left[\left(\frac{1}{2} \right)^2 + 2 \frac{1}{4} (1 - \bar{\sigma}^*) + \left(\frac{1}{2} \right)^2 (1 - \bar{\sigma}^*)^2 \right] \pi u + 2 \frac{1}{4} \bar{\sigma}^* (1 - u) + \left(\frac{1}{2} \right)^2 (1 - (1 - \bar{\sigma}^*)^2) u \\ &= \frac{(2 - \bar{\sigma}^*)}{4} \pi u + \frac{1}{2} \bar{\sigma}^* (1 - u) + \left(\frac{1}{2} \right)^2 (1 - (1 - \bar{\sigma}^*)^2) u. \end{aligned}$$

The first term is the probability that the two outlets move together but do not research further, i.e. they publish in $t = 1$. As a result, the probability of matching the state is π and only fraction u of the audience gets this payoff. The second term is the probability that there is a second mover outlet that does research and therefore, the remaining fraction $1 - u$ receives this payoff. Finally, the third is when both outlets move together in $t = 2$ after researching further. In this case, they match the state perfectly but fraction $1 - u$ does not receive this payoff.

As discussed in Section 2.1, another important effect of the Internet has been to make it easier to conduct preliminary research. Emails and social media make it particularly easy to share pictures, video and text from any part of the world. One way to interpret it is as an increase in π or the precision of s . This, [Knobel \(2018\)](#) argues, should lead to better reporting. We show below that that is not necessarily true. Our next proposition shows that the overall effect of an increase in π on V is dependent on the kind of story u being explored.

Proposition 2.8. *There exists an interior u , $\bar{u}^V \in (0, 1)$, such that if $u < \bar{u}^V$ an increase in precision π of initial signal s decreases the overall welfare V .*

⁹Note that even if the audience knows that outlets may be publishing without research, it is still better to listen to the outlets rather than to follow the priors in decision-making.

The intuition for this somewhat surprising result is easy. The equilibrium probability of research falls as precision π increases. This is because a higher π reduces the reputational gain that comes with separation. The audience attributes correctly matching the state more to better initial information that comes costlessly due to better technology rather than actual research. Preemption concerns, therefore, become more salient and push the outlets towards speed. Of course, this is not a concern if the outlets were inclined towards speed-driven journalism to begin with. However, it reduces the welfare and hurts the audience when the outlets are more accuracy driven. This happens when the proportion of urgent audience type seeking the information is low and then π increases.

Formally, the welfare of the urgent audience increases with an increase in π .

$$\frac{\partial V_u}{\partial \pi} = \frac{(2 - \sigma)^2}{4} + \left[-\pi \frac{2 - \bar{\sigma}}{2} + \frac{1 - \bar{\sigma}}{2} \right] \frac{\partial \bar{\sigma}}{\partial \pi} > 0$$

because $\frac{(2 - \bar{\sigma})^2}{4} > \frac{1 - \bar{\sigma}}{2}$. And the welfare of the patient audience reduces due to an increase in π ,

$$\frac{\partial V_{1-u}}{\partial \pi} = \frac{1}{2} \frac{\partial \bar{\sigma}}{\partial \pi} < 0.$$

When the fraction of urgent audience is low enough, an increase in π hurts an average audience member. Better preliminary research is good news for the audience only if separation does not happen and is not desired. However it also discourages separation, which hurts the audience when it is desired.

2.5 Information dissemination by a source

We now turn back to our original model and discuss the case of a strategic source. We can use our model to determine how a source can share information with media outlets.

In general, our strategic source's preferences are summarized by the following objective function,

$$\mathbb{1}\{\text{publication in } t = 1\} + \mu \Pr(\text{matching the state}).$$

Therefore, the source has a preference for speed vs. accuracy. The parameter $\mu \geq 0$ captures the weight that the source places on accurate information from at least one outlet *vis-à-vis* having at least one outlet publishing in period 1. For instance, a concerned citizen or an employee in a firm witnessing some wrongdoing might have a high preference for accuracy. On the flip side, a politically-motivated source who merely wants to get some potentially incorrect information out quickly will have a low preference for accuracy. We want to determine whether a source wants to share information with one or both the outlets to fulfill her objective.

In line with our model, we will assume that if the source shares a story with both of the outlets, both are aware that the other also possesses the same story. Therefore, the information is shared “publicly”.¹⁰ But when the source shares information with just one outlet, we will assume that the other is unaware. This allows the outlet with a story to effectively behave as a monopolist from our analysis in Section 2.3.1. In addition, we assume that the source possesses a story of a fixed precision π . She makes her decision about who to share the story with at the beginning of the game before time 0. The type of the outlet is still each outlet’s private information; the source does not have this information when making her decision.

First, we make a simple observation that follows from our analysis of monopoly and duopoly. (In what follows, we drop the star notation for convenience with an understanding that we are talking about equilibrium values.)

Corollary 2.3. *The equilibrium probability of research by a high outlet in monopoly $\sigma_M > 0$ while in duopoly is $\sigma_D \geq 0$.*

Corollary 2.3 is an important one. It highlights that while in monopoly the probability of research is always positive; in duopoly it might be zero if v is sufficiently high. This corollary will help us outline the behavior of a source who is aware of how high v is associated with her story.

Second, we write down the expected utility of the source for the equilibrium research probabilities that will be induced in the following subgame. The expected payoff from sharing information with one outlet is

$$\frac{1}{2}(1 + \mu\pi) + \frac{1}{2}[\sigma_M\mu + (1 - \sigma_M)(1 + \mu\pi)] \quad (2.3)$$

The first term reflects what the source gets if she gives the story to a low quality outlet, and the second term is for giving it to a high quality outlet. Similarly, the expected payoff from sharing information with both the outlets is

$$\frac{1}{4}(1 + \mu\pi) + \frac{1}{4}[1 + \mu(\sigma_D + (1 - \sigma_D)\pi)]2 + \frac{1}{4}[(1 - \sigma_D)^2(1 + \mu\pi) + 2\sigma_D(1 - \sigma_D)(1 + \mu) + \sigma_D^2\mu]. \quad (2.4)$$

Again, the first term reflects the source’s payoff from facing two low type outlets. The second is the payoff from facing one high type and one low type outlet. Note that in this case the story is always published in the first period, but the high outlet matches the state only if it does research.

¹⁰One may imagine a politician revealing some negative evidence about a competitor on Twitter as an example. It is common for news outlets to pick up this information and relay it, either as is or after further fact-checking and investigations.

The third term is the payoff from facing two high type outlets. Here, the possible situations are that neither researches; one researches, or both research. The following lemma helps simplify the source's optimal response for a given σ_M and σ_D .

Lemma 2.3. *The source's best response can be summarized as follows:*

- *The source prefers to share the story with both the outlets unambiguously for any $\mu > 0$ if $\frac{\sigma_D^2}{2} \leq \sigma_M \leq \frac{\sigma_D(4-\sigma_D)}{2}$.*
- *Otherwise, the source prefers to share the story with both outlets if*

$$\mu(1 - \pi)(2\sigma_M - \sigma_D(4 - \sigma_D)) \leq 2\sigma_M - \sigma_D^2$$

The lemma shows that there is a range of equilibrium σ_M and σ_D for which the source always prefers to send information to both the outlets independent of μ . Interestingly, this region lies around the $\sigma_D = \sigma_M$ line. Therefore, the lemma shows that for σ_M and σ_D close to each other there is reason to prefer both outlets. To understand why, let us break this down into two further statements.

First, there are parameters where one outlet alone is more likely to research than when it is competing with another (i.e. $\sigma_M > \sigma_D$) and μ is very large, and yet the source prefers to share the story with two outlets. This happens because a lower σ_D is compensated by a higher probability of investigation from more firms. But this requires σ_D and σ_M to be close to each other. To see this, let us compare the total probability of research (and matching the state) from sharing the story with one vs. both the outlets. When shared with one it is equal to $\frac{1}{2}\sigma_M$. When shared with both it is given by

$$\frac{1}{4}[\sigma_D^2 + 2\sigma_D(1 - \sigma_D)] + \frac{2}{4}\sigma_D = \sigma_D - \frac{\sigma_D^2}{4}.$$

Therefore, despite $\sigma_M > \sigma_D$ the source shares the story with both outlets if $\sigma_D - \frac{\sigma_D^2}{4} \geq \frac{1}{2}\sigma_M$. This condition simplifies to give us our upper bound

$$\sigma_M \leq \frac{\sigma_D(4 - \sigma_D)}{2}.$$

Second, there are parameters where one firm is less likely to research than two (i.e. $\sigma_M < \sigma_D$) and μ is very low, and yet the source prefers to share the story with two outlets. This, on the other hand, happens because competition between two firms ensures the story comes out quicker despite each independent outlet researching with a higher probability. To see this,

we now compare the total probabilities of the story being published in $t = 1$ under the two scenarios. When shared with one, this probability is equal to $\frac{1}{2} + \frac{1}{2}(1 - \sigma_M) = 1 - \frac{\sigma_M}{2}$. When shared with both it is given by

$$\frac{1}{4} + \frac{2}{4} + \frac{1}{4}[(1 - \sigma_D)^2 + 2\sigma_D(1 - \sigma_D)] = 1 - \frac{\sigma_D^2}{4}.$$

So, now despite $\sigma_M < \sigma_D$ the source shares the story with both outlets if $1 - \frac{\sigma_D^2}{4} \geq 1 - \frac{\sigma_M}{2}$. This condition simplifies to

$$\sigma_M \geq \frac{\sigma_D^2}{2},$$

giving us our lower bound. But note again that for this argument to work σ_D and σ_M should not be too different from each other. When this is the case, then what the source does depends on her preference μ (captured in the second bullet point of Lemma 2.3).

In Figure 2.2, the shaded gray region shows the combinations of σ_D and σ_M where the source always prefers to share stories with both outlets. The region is enclosed between $\sigma = \frac{\sigma_D(4 - \sigma_D)}{2}$ (green) and $\sigma_M = \frac{\sigma_D^2}{2}$ (orange), which includes $\sigma_M = \sigma_D$ (blue).

We now look at possible equilibria that can arise in the $\sigma_D - \sigma_M$ space relative to the source's preferences. We begin by plotting an equilibrium frontier for a given \bar{c} and ε .

Definition 2.1 (Equilibrium frontier). *The equilibrium frontier is given by the combination of equilibrium σ_D and σ_M generated by varying $\pi \in [.5, 1]$ for $v = 0$ and a fixed \bar{c} and ε .*

The equilibrium frontier, therefore, shows the maximum equilibrium value that σ_D can take for any equilibrium σ_M (since σ_D is decreasing in v from Corollary 2.3 and we are setting $v = 0$). As proved in Lemma 2.2, when $v = 0$, $\sigma_D > \sigma_M$. Therefore, the frontier lies to the right of the 45° line. In addition, note that it is upwards sloping. The positive slope is a result of the fact that both σ_M and σ_D are decreasing functions of π .¹¹ A north-east movement along the frontier arises due to a decrease in π . Figure 2.2 plots the equilibrium frontier for $\bar{c} = 2$ and $\varepsilon = 1$ in red.¹²

Once we have the equilibrium frontier, it is easy to see the set of all possible equilibrium values that might arise for different parameter ranges. Particularly, increasing v is a leftward movement from the frontier along the same σ_M . For v sufficiently high, $\sigma_D = 0$ while $\sigma_M > 0$ (see Corollary 2.3). We are now left with comparing these equilibrium values with what the source wants.

¹¹The proofs have been omitted from the main text for the sake of brevity.

¹²We choose a high value of ε for graphical representation only. When ε is low, the range of σ_M and σ_D is also small, and it becomes difficult to clearly see the equilibria graphically.

Proposition 2.9. *For a source with preferences given by $\mu \geq 0$,*

- *there exists an $\varepsilon > 0$ small enough and two thresholds $\bar{v} > \underline{v}$ such that if $v < \underline{v}$ the source sends the story to both if $\mu \geq \frac{\sigma_D^2 - 2\sigma_M}{\sigma_D(4 - \sigma_D) - 2\sigma_M} \frac{1}{(1 - \pi)}$, if $\bar{v} \geq v \geq \underline{v}$ the source always sends the story both, and if $v > \bar{v}$ the source sends the information to both if $\mu \leq \frac{2\sigma_M - \sigma_D^2}{2\sigma_M - \sigma_D(4 - \sigma_D)} \frac{1}{(1 - \pi)}$;*
- *there exists an $\varepsilon > 0$ large enough and a threshold $\bar{\bar{v}}$ such that if $v \leq \bar{\bar{v}}$ the source sends the story to both, and if $v > \bar{\bar{v}}$ the source sends the information to both if $\mu \leq \frac{2\sigma_M - \sigma_D^2}{2\sigma_M - \sigma_D(4 - \sigma_D)} \frac{1}{(1 - \pi)}$.*

Our third main result follows by setting $\mu = 0$ in the above proposition. It pertains to the situation where the source only cares about getting the story out quickly independent of whether it is accurate or not. Political actors are often interested in doing so to highlight their achievements or to bring out potentially damaging information about their competitors. Twitter and other social media platforms are one way to communicate such stories, which are then picked up by media outlets and relayed to the public without further research.

Corollary 2.4. *When the source does not care about accuracy, i.e. $\mu = 0$,*

- *there exists an $\varepsilon > 0$ small enough and \bar{v} such that for $v < \bar{v}$, the source sends the story to one outlet, and sends to two in all other cases, and*
- *there exists an $\varepsilon > 0$ large enough such that the source sends the story to both outlets.*

The proof of both Proposition 2.9 and its corollary is by construction. The idea is that when ε is small, (at least a part of) the frontier lies below the orange line in Figure 2.2. Therefore, there arise two thresholds on v where only the middle part lies between the two curves. For a ε high enough, there is only threshold on v as depicted in the figure. One can easily get the result of Corollary 2.4 by setting $\mu = 0$ in Proposition 2.9.

Consider the intuition for the case of $\mu = 0$. When the intrinsic motivation to conduct research is high then independent of whether only one outlet has the story or both, the outlets are more likely to conduct research. This is, however, not something a $\mu = 0$ source desires. By sending to both, she is able to create preemption risk as well (even for a low v). This improves on the situation of sending to one as the outlets are more driven towards speed. On the other hand, when intrinsic motivation is low, outlets are less likely to research. Now, the source does not always want to share the story with both. Notably, when v is low the source wants to share information with just one. Sending to both risks the outlets trying to separate by doing research, thereby increasing the overall probability of research. However, again when v is high,

the source is happy to share the story with both as preemption concerns will become salient for the outlets.¹³

2.6 Conclusion

There have been increasing concerns in the past decade about how the Internet has altered the incentives of media outlets. Notably, media critics have argued that increasing competition in the Internet era has pushed outlets towards speed-driven journalism. Our model shows that conventional wisdom about the effect of competition and the modern digital environment on the media market should be taken *cum grano salis*. We prove that competition in itself may make it easier for high quality outlets to engage in more research-driven journalism to separate themselves from the low quality outlets. For this to happen, it must be that the action of one of the outlets is somehow informative about the type of the other. This result and intuition finds support in some of the new media studies literature such as in [Knobel \(2018\)](#) and [Carson \(2019\)](#).

It is, however, worth emphasizing the importance of a “sophisticated” audience in generating the better-reporting result. We need the audience to place importance on the accuracy of stories, and not always seek quick information. [Gentzkow and Shapiro \(2008\)](#) suggest that scoop value is usually not too high in the media markets. But at the same time, some media scholars have argued that the audience usually seeks information earlier on social media. Similarly, our model shows the importance of the audience observing the *sequence* of publication. This might also be an issue if technology perfectly “flattens the digital plain” (see Section 2.1.1).

Our paper is one of the first to incorporate preemption and reputation concerns in a single model by thinking of a natural setting where both incentives play a role. However, there is further scope for research here. For instance, one may expand the model to include news media bias. Bias and the speed-accuracy tradeoff can interact in interesting ways. If bias makes reputational gains less salient (e.g. because future readership does not depend on reputation) then it should push toward speed. On the other hand, if bias implies a less informative publication and hence a smaller “scoop value”, then it may actually push toward accuracy.

Our model also produces important testable predictions about how the modern digital environment has altered the media landscape. First, we should see better reporting of non-urgent issues in the Internet-age as the outlets try to build a reputation on such stories. Second, the

¹³The intuition for the general case presented in Proposition 2.9 is similar but it is not easy to make sharp predictions like we could with $\mu = 0$ case. However, some additional predictions can be made by choosing specific μ values. For instance, when v is very high so that $\sigma_D = 0$, the source prefers to send to one outlet only if $\mu > 2$.

effect of the Internet on the reporting of breaking-news type stories is ambiguous. It might improve because of better source information but might deteriorate because of more time pressure.

2.7 Tables

Table 2.1: Deep (first row) and simple (second row) accountability reporting (as a % of total front-page stories in April) in a sample of 9 newspapers in the US for 1991-2011 in five-year gaps

Newspaper group	Newspaper	1991	1996	2001	2006	2011	Average
Large	<i>Wall Street Journal</i>	1.28	2.33	5.88	5.26	4.85	4.03
		30.77	22.09	23.53	22.11	27.18	25.06
	<i>Washington Post</i>	1.51	3.55	4.23	2.72	7.74	3.80
		25.63	27.41	31.92	37.50	36.13	31.43
	<i>New York Times</i>	0.34	0.93	4.35	5.43	3.19	2.46
Metropolitan dailies		10.51	9.29	18.26	19.57	28.72	15.82
	<i>Albany Times Union (NY)</i>	6.35	1.22	3.45	4.12	3.61	3.64
		47.62	23.17	28.74	17.53	36.14	26.37
	<i>Denver Post</i>	0.00	4.85	1.80	3.06	5.13	2.92
		23.33	22.33	28.83	29.59	43.59	28.96
	<i>Minneapolis Star Tribune</i>	2.46	1.15	1.83	2.86	5.00	2.68
		31.97	36.78	22.02	34.29	41.00	32.89
	<i>Atlanta Journal-Constitution</i>	1.20	0.00	1.06	1.75	11.84	2.30
		14.97	11.11	13.30	30.70	48.68	20.52
Small	<i>Bradenton Herald (FL)</i>	0.93	1.61	1.14	1.27	1.44	1.26
		19.44	33.87	32.95	21.52	19.42	24.16
	<i>Lewiston Tribune (ID)</i>	0.00	0.00	0.00	0.00	1.45	0.32
		22.22	15.25	40.74	28.33	23.19	25.80
Average		1.26	1.81	2.92	3.25	4.46	2.69
		21.52	19.78	24.46	27.26	32.59	24.94

Source: *The Watchdog Still Barks: How Accountability Reporting Evolved for the Digital Age*. [Knobel \(2018\)](#). The author analyzed the content of every front-page story that was published in the month of April (randomly selected) in five-year gaps starting 1991 in a select sample of 9 newspapers. The stories chosen for deep and simple categories involved the following procedure. First, the author eliminated stories that were breaking news. Second, she eliminated stories that had no relation to public policy or politics. In all, she analyzed 1,491 stories in depth using content analysis. Simple accountability reports/stories are those that took a few hours or days to complete, relying on straightforward reporting such as interviews or reviewing published documents. Deep accountability reports/stories are those that took weeks or months to develop and would have remained secret without the journalists' work.

2.8 Figures




Figure 2.1: Equilibrium σ_D^* and σ_M^* for when $\pi = .6$, $v = .3$, $\bar{c} = 2$ and $\varepsilon = .1$




Figure 2.2: Equilibria in $\sigma_D - \sigma_M$ space and the behavior of the source

2.9 Appendix

A Proofs from the main text

Proof of Observation 2.1

Proof. Suppose that the outlet chooses $d = \text{pub}$. Without loss of generality, suppose that $s^i = a$. It is easy to see that $\Pr(\omega = a | s^i = a) > \Pr(\omega = b | s^i = a)$ because

$$\frac{\pi^{\frac{1}{2}}}{\pi^{\frac{1}{2}} + (1 - \pi)^{\frac{1}{2}}} > \frac{(1 - \pi)^{\frac{1}{2}}}{\pi^{\frac{1}{2}} + (1 - \pi)^{\frac{1}{2}}}$$

which is true because $\pi > \frac{1}{2}$. □

Proof of Lemma 2.1

Proof. First part. Without loss of generality, suppose that $s^i = a$. Then, if i chooses to publish, it will endorse state a , i.e. send message $m = a$. by Bayes' rule,

$$\Pr(\omega = a | s^i = a) = \frac{\pi^{\frac{1}{2}}}{\pi^{\frac{1}{2}} + (1 - \pi)^{\frac{1}{2}}} = \pi$$

as claimed.

Second part. We are interested in the probability that j matches the state from choosing $d = \text{pub}$ when i has received a signal s^i . This is equal to

$$\Pr(s^j = a | s^i) \Pr(\omega = a | s^j = a \text{ and } s^i) + \Pr(s^j = b | s^i) \Pr(\omega = b | s^j = b \text{ and } s^i) \quad (2.9.A.1)$$

Note that, for a generic s^j , by Bayes' rule we have that $\Pr(s^j | s^i) = \frac{\Pr(s^j \text{ and } s^i)}{\Pr(s^i)}$ and

$$\Pr(\omega = s^j | s^j \text{ and } s^i) = \frac{\Pr(s^j \text{ and } s^i | \omega = s^j) \Pr(\omega = s^j)}{\Pr(s^j \text{ and } s^i)}$$

As a consequence, (2.9.A.1) can be simplified to

$$\frac{\Pr(s^j = a \text{ and } s^i | \omega = a) \Pr(\omega = a)}{\Pr(s^i)} + \frac{\Pr(s^j = b \text{ and } s^i | \omega = b) \Pr(\omega = b)}{\Pr(s^i)} \quad (2.9.A.2)$$

However, since signals are independent conditional on the state,

$$\Pr(s^j \text{ and } s^i | \omega = s^j) = \Pr(s^j | \omega = s^j) \Pr(s^i | \omega = s^j)$$

Moreover, $\Pr(s^j | \omega = s^j) = \pi$. Hence, (2.9.A.2) becomes

$$\pi \frac{\Pr(s^i|\omega = a)\Pr(\omega = a) + \Pr(s^i|\omega = b)\Pr(\omega = b)}{\Pr(s^i)} = \pi$$

as claimed. □

Proof of Proposition 2.2

Proof. Suppose that a high type outlet chooses $d = res$ with probability σ . Reminding ourselves from the main text that

$$\gamma(R) = \frac{\sigma + (1 - \sigma)\pi}{\sigma + (1 - \sigma)\pi + \pi}$$

$$\gamma(W) = \frac{(1 - \sigma)(1 - \pi)}{(1 - \sigma)(1 - \pi) + (1 - \pi)} = \frac{1 - \sigma}{2 - \sigma}$$

from Bayes' rule and using the fact that a low type outlet always chooses *pub*.

A high type outlet optimally chooses *res* if

$$\gamma(R) - c \geq \pi\gamma(R) + (1 - \pi)\gamma(W) \implies c \leq (1 - \pi)(\gamma(R) - \gamma(W)) := c_M$$

In equilibrium the conjectured σ must be equal to the actual one, hence it must be that

$$\sigma^* = F(c_M(\sigma^*)) = \frac{c_M(\sigma^*) + \varepsilon}{\bar{c} + \varepsilon}. \quad (2.9.A.3)$$

We need to check if such a fixed point exists. To do so, three observations are in order. First, note that both the LHS and RHS of the above are continuous in σ^* . Second, $\text{LHS}(\sigma^* = 0) = 0 < \text{RHS}(\sigma^* = 0) = \frac{\varepsilon}{\bar{c} + \varepsilon}$ (as $c_M = 0$ at $\sigma^* = 0$). Third, $\text{LHS}(\sigma^* = 1) = 1 > \text{RHS}(\sigma^* = 1) = F(\frac{1 - \pi}{1 + \pi})$. Therefore, the above is true.

Finally, we need to check for the uniqueness of the fixed point. Note that

$$\frac{\partial \text{RHS}}{\partial \sigma^*} = \frac{1 - \pi}{\bar{c} + \varepsilon} \left[\frac{\pi(1 - \pi)}{(\sigma^* + (1 - \sigma^*)\pi + \pi)^2} + \frac{1}{(2 - \sigma^*)^2} \right] > 0,$$

but the sign of

$$\frac{\partial^2 \text{RHS}}{\partial (\sigma^*)^2} = \frac{2(1 - \pi)}{\bar{c} + \varepsilon} \left[-\frac{\pi(1 - \pi)^2}{(\sigma^* + (1 - \sigma^*)\pi + \pi)^3} + \frac{1}{(2 - \sigma^*)^3} \right]$$

is not clear immediately. $\frac{\partial^2 \text{RHS}}{\partial (\sigma^*)^2} > 0$ requires

$$-\pi(1 - \pi)^2(2 - \sigma^*)^3 + (\sigma^* + (1 - \sigma^*)\pi + \pi)^3 > 0 \quad (2.9.A.4)$$

It is easy to see that the LHS of (2.9.A.4) is strictly increasing in σ^* for all $\pi \in (0.5, 1]$. Moreover, the LHS of (2.9.A.4) when we substitute $\sigma^* = 0$ is $-1 + 2\pi > 0$. As a consequence, the RHS of (2.9.A.3) is strictly increasing and convex. Combined with the above, it means that there is only one fixed point in the $[0, 1]$ interval. \square

Proof of Proposition 2.3

Proof. If θ is known, then by choosing *pub* in $t = 1$ a high quality outlet receives a payoff of

$$\frac{1}{2} \frac{v}{2} + \frac{1}{2} \left[v\sigma + \frac{v}{2}(1 - \sigma) \right] + \mathbb{1}\{\theta = h\},$$

where σ is the (symmetric) probability that the high quality competitor engages in more research. By instead choosing *res* and publishing in $t = 2$ a high type outlet gets a payoff of $\frac{1}{2}\sigma\frac{v}{2} + \mathbb{1}\{\theta = h\} - c$. Comparing the two, each outlet is willing to investigate iff $c \leq -\frac{v}{2}$. As a consequence, $\sigma_D^* = F\left(-\frac{v}{2}\right)$ in symmetric equilibrium. Research happens with positive probability when $-\frac{v}{2} > -\varepsilon$, which can be rearranged to $v < 2\varepsilon$. \square

Proof of Proposition 2.4

Proof. We complete this proof in several steps. To begin with, we conjecture that whenever an outlet chooses to publish, it is optimal to endorse the state suggested by the signal. This will be verified at the end of the proof.

Step 1: We begin by showing that in any signal-based equilibria outlets' period 1 decisions on whether to research or publish is described by a threshold on c . This follows from the discussion in the text. Let σ^i and σ^j be the conjectured strategies. Then equation (2.1) defines the threshold c_D^i for outlet i .

$$c^i \leq \frac{1}{2} \left[(\gamma^i(\emptyset) - \gamma^i(1)) (\sigma^j - (2 - \sigma^j)\pi^2) + (2 - \sigma^j) (1 - \gamma^i(1)) \right] - \frac{1}{2}v := c_D^i \quad (2.1)$$

where $\gamma(\emptyset) = \frac{\sigma^i\sigma^j + (1-\sigma^i)(2-\sigma^j)\pi^2}{\sigma^i\sigma^j + (2-\sigma^i)(2-\sigma^j)\pi^2}$ and $\gamma(1) = \frac{1-\sigma^i}{2-\sigma^i}$. The problem is identical for player j .

Step 2: Next, we show that for any σ^j there is only one σ^i that solves the equilibrium fixed point for player i .

Given that cost is uniformly distributed in $[-\varepsilon, \bar{c}]$ and that, in equilibrium the conjectured probability of investigation must be equal to the actual probability, the equilibrium levels of σ^i

and σ^j must be the solutions of

$$\sigma^i = F(c_D^i(\sigma^i, \sigma^j)) \text{ and } \sigma^j = F(c_D^j(\sigma^j, \sigma^i))$$

where

$$F(c_D^i(\sigma^i, \sigma^j)) = \begin{cases} 0 & c_D^i(\sigma^i, \sigma^j) < -\varepsilon \\ \frac{c_D^i(\sigma^i, \sigma^j) + \varepsilon}{\bar{c} + \varepsilon} & -\varepsilon \leq c_D^i(\sigma^i, \sigma^j) \leq \bar{c} \\ 1 & c_D^i(\sigma^i, \sigma^j) > \bar{c} \end{cases}$$

and

$$f(c_D^i(\sigma^i, \sigma^j)) = \begin{cases} 0 & c_D^i(\sigma^i, \sigma^j) < -\varepsilon \\ \frac{1}{\bar{c} + \varepsilon} & -\varepsilon \leq c_D^i(\sigma^i, \sigma^j) \leq \bar{c} \\ 0 & c_D^i(\sigma^i, \sigma^j) > \bar{c} \end{cases}$$

We want to show that, for every σ^j , there is only one σ^i that solves $\sigma^i = F(c_D^i(\sigma^i, \sigma^j))$.

1. The LHS is linear, with slope equal to 1, starting at 0 and ending at 1.
2. As $c_D^i(\sigma^i = 1, \sigma^j) < 1 < \bar{c}$, the RHS evaluated at $\sigma^i = 1 < 1 = \text{LHS at } \sigma^i = 1$;
3. The RHS evaluated at $\sigma^i = 0$ is greater than or equal to zero.
4. For any σ^j , both LHS and RHS are continuous in σ^i .

Hence, they cross at least once and there is at least one solution to this fixed point problem.

To show that they cross only once, we need to show that the slope of the RHS is never above 1. First, note that the slope of the RHS is either 0 or $f(c_D^i) \frac{\partial c_D^i}{\partial \sigma^i}$. Second, $\frac{\partial \gamma^i(\emptyset)}{\partial \sigma^i} = \frac{(\sigma^j - (2 - \sigma^j)\pi^2)\pi^2(2 - \sigma^j)}{(\sigma^i \sigma^j + (2 - \sigma^i)(2 - \sigma^j)\pi^2)^2}$, whose sign depends on the sign of $(\sigma^j - (2 - \sigma^j)\pi^2)$ and $\frac{\partial \gamma^i(1)}{\partial \sigma^i} = \frac{-1}{(2 - \sigma^i)^2} < 0$. Using these we can write $\frac{\partial c_D^i}{\partial \sigma^i} = \frac{1}{2} \left[\frac{(\sigma^j - (2 - \sigma^j)\pi^2)^2 \pi^2 (2 - \sigma^j)}{(\sigma^i \sigma^j + (2 - \sigma^i)(2 - \sigma^j)\pi^2)^2} + \frac{2 - \pi^2(2 - \sigma^j)}{(2 - \sigma^i)^2} \right]$ where both terms are always positive. Third, we can show that the sign of $\frac{\partial^2 c_D^i}{\partial (\sigma^i)^2}$ is ambiguous, but $\frac{\partial^3 c_D^i}{\partial (\sigma^i)^3} \geq 0$. As a consequence, the second derivative is always increasing in σ^i and the first derivative is convex in σ^i . So, $\frac{\partial c_D^i}{\partial \sigma^i}|_{\sigma^i=1} > \frac{\partial c_D^i}{\partial \sigma^i}|_{\sigma^i=0}$, and c_D^i reaches its steepest point around $\sigma^i = 1$. Therefore, it is enough to show that $\frac{\partial c_D^i}{\partial \sigma^i}|_{\sigma^i=1} \leq 1$. This requires

$$2(\sigma^j + (2 - \sigma^j)\pi^2)^2 \geq (\sigma^j - (2 - \sigma^j)\pi^2)^2 \pi^2 (2 - \sigma^j) + (\sigma^j + (2 - \sigma^j)(1 - \pi^2))(\sigma^j + (2 - \sigma^j)\pi^2)$$

which further simplifies to

$$(\sigma^j + (2 - \sigma^j)\pi^2)^2 (2 - \sigma^j - 2 + \sigma^j) \geq -4\sigma^j(2 - \sigma^j)^2 \pi^4.$$

This latter condition is always verified (strictly for positive σ^j , weakly when $\sigma^j = 0$).

Now, combining the above with the fact that $c_D^i(\sigma^i = 1, \sigma^j) < 1$, implies that they cannot cross more than once.

Step 3: Third, we show that if an equilibrium exists, it is unique for $\bar{c} \geq 2$.

Define $\hat{\sigma}^i(\sigma^j)$ the optimal σ^i for a given σ^j . In equilibrium, it must be that

$$\hat{\sigma}^i(\hat{\sigma}^j(\sigma^i)) = \sigma^i \quad (2.9.A.5)$$

Rearranging, the equilibrium is the solution of $\hat{\sigma}^i(\hat{\sigma}^j(\sigma^i)) - \sigma^i = 0$. Differentiating with respect to σ^i , we obtain $\frac{\partial \hat{\sigma}^i}{\partial \hat{\sigma}^j} \frac{\partial \hat{\sigma}^j}{\partial \sigma^i} - 1 = 0$. For the equilibrium to be unique (conditional on its existence), it is now sufficient to show that the LHS is negative. This implies that only one fixed point of (2.9.A.5) can be found. This happens when $\frac{\partial \hat{\sigma}^i}{\partial \hat{\sigma}^j}$ and $\frac{\partial \hat{\sigma}^j}{\partial \sigma^i}$ are between -1 and 1 . As the players are identical, it is enough to show that this holds for one of them.

To show the above, begin by noting that $\sigma^i(\sigma^j)$ is implicitly defined by the unique solution of $\sigma^i - F(c_D^i(\sigma^i, \sigma^j)) = 0$. (Going forward we drop the $\hat{\cdot}$ notation with an understanding that we are concerned with optimal responses.) As $\frac{\partial c_D^i}{\partial \sigma^i}|_{\sigma^i=1} \leq 1$, we can use implicit function theorem. Therefore,

$$\frac{\partial \sigma^i}{\partial \sigma^j} = \frac{\frac{\partial F(c_D^i)}{\partial \sigma^j}}{1 - \frac{\partial F(c_D^i)}{\partial \sigma^i}} \quad (2.9.A.6)$$

Consider first the denominator of (2.9.A.6). From Step 2, we know that it is always positive. Moreover, it will be smaller the bigger is $\frac{\partial F(c_D^i)}{\partial \sigma^i}$. On the other hand, it is the biggest when $\frac{\partial F(c_D^i)}{\partial \sigma^i}$ is zero. When $\frac{\partial F(c_D^i)}{\partial \sigma^i}$ is non-zero, it is linear and increasing in $\frac{\partial c_D^i}{\partial \sigma^i}$. As this reaches its maximum for $\sigma^i = 1$, we simply replace it and look for a maximum with respect to σ^j .

$$\begin{aligned} \max_{\sigma^j} \frac{\partial c_D^i}{\partial \sigma^i}|_{\sigma^i=1} &= \frac{1}{2} \left[\frac{(\sigma^j - (2 - \sigma^j)\pi^2)^2 \pi^2 (2 - \sigma^j)}{(\sigma^j + (2 - \sigma^j)\pi^2)^2} + 2 - \pi^2 (2 - \sigma^j) \right] \\ &= \frac{1}{2} \left[2 - \frac{4\sigma^j(2 - \sigma^j)^2 \pi^4}{(\sigma^j + (2 - \sigma^j)\pi^2)^2} \right] \\ &= 1 \end{aligned}$$

where the second equality is a rearrangement and the third one follows from the fact that this is maximized for $\sigma^j = 0$.

As a consequence, $\max_{\sigma^i, \sigma^j} \frac{\partial F(c_D^i)}{\partial \sigma^i} = \frac{1}{\bar{c} + \varepsilon}$ and the smallest the denominator can be is

$$\frac{1}{\bar{c}+\varepsilon}.$$

Second, consider the numerator. $\frac{\partial F(c_D^i)}{\partial \sigma^j}$ is either zero or $\frac{1}{\bar{c}+\varepsilon} \frac{\partial c_D^i}{\partial \sigma^j}$. Further, note that

$$\frac{\partial c_D^i}{\partial \sigma^j} = \frac{1}{2} \left[\frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} (\sigma^j - (2 - \sigma^j)\pi^2) + \gamma^i(\emptyset) + \pi^2(\gamma^i(\emptyset) - \gamma^i(1)) - 1 \right]. \quad (2.9.A.7)$$

Finding the overall maximum and minimum is complicated, so we look for sufficient conditions. We start out by looking at $\frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j}$. After few algebraic manipulations, we derive

$$\frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} = \frac{2\sigma^i\pi^2}{(\sigma^i\sigma^j + (2 - \sigma^i)(2 - \sigma^j)\pi^2)^2}$$

Its sign is positive, but it is hard to determine the maximum. We proceed as follows. First, note that

$$\frac{\partial^2 \gamma^i(\emptyset)}{\partial (\sigma^j)^2} = \frac{-4(\sigma^i - (2 - \sigma^i)\pi^2)\sigma^i\pi^2}{(\sigma^i\sigma^j + (2 - \sigma^i)(2 - \sigma^j)\pi^2)^3}$$

whose sign is ambiguous. However,

$$\frac{\partial^3 \gamma^i(\emptyset)}{\partial (\sigma^j)^3} = \frac{12(\sigma^i - (2 - \sigma^i)\pi^2)^2\sigma^i\pi^2}{(\sigma^i\sigma^j + (2 - \sigma^i)(2 - \sigma^j)\pi^2)^4}$$

which is positive. This implies that (for any σ^i) $\frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j}$ is a convex function in σ^j which is maximized either at $\sigma^j = 0$ or at $\sigma^j = 1$. By substitution,

$$\begin{aligned} \frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=0} &= \frac{\sigma^i}{2\pi^2(2 - \sigma^i)^2} \\ \frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=1} &= \frac{2\sigma^i\pi^2}{(\sigma^i + (2 - \sigma^i)\pi^2)^2} \end{aligned}$$

Still we are left to determine the maximum possible value of $\frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j}$ because the comparison is not straightforward. But we can show that for every π , $\max_{\sigma^i} \frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=0} > \max_{\sigma^i} \frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=1}$. To prove this, first see that

$$\max_{\sigma^i} \frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=0} = \frac{1}{2\pi^2}$$

But to get $\max_{\sigma^i} \frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=1}$,

$$\frac{\partial}{\partial \sigma^i} \left(\frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=1} \right) = \frac{\partial}{\partial \sigma^i} \left(\frac{2\sigma^i\pi^2}{(\sigma^i + (2 - \sigma^i)\pi^2)^2} \right) = \frac{2\pi^2(\sigma^i + (2 - \sigma^i)\pi^2) - 4(1 - \pi^2)\sigma^i\pi^2}{(\sigma^i + (2 - \sigma^i)\pi^2)^3} \quad (2.9.A.8)$$

Note that the relevant expression in (2.9.A.8) is always positive for $\sigma^i \leq \frac{2\pi^2}{1-\pi^2}$. For a sufficiently

high π , this includes the whole range of values of σ^i . Hence, the function is maximised at $\sigma^i = 1$, and

$$\max_{\sigma^i} \frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=1} = \frac{2\pi^2}{(1+\pi^2)^2}.$$

But now it is easy to see that $\frac{1}{2\pi^2} \geq \frac{2\pi^2}{(1+\pi^2)^2}$ requires $1 + 2\pi^2 - 3\pi^4 \geq 0$, which is always true for $\pi \in (0.5, 1]$. Therefore, our claim of $\max_{\sigma^i} \frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=0} > \max_{\sigma^i} \frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=1}$ is true.

However, for low π , we have that $\arg\max_{\sigma^i} \frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=1} = \frac{2\pi^2}{1-\pi^2} \in [0, 1]$. In particular, this happens for $\pi^2 \leq \frac{1}{3}$. Even in this case, it is easy to show that $\frac{1}{2\pi^2} \geq \frac{2\pi^2 \left(\frac{2\pi^2}{1-\pi^2} \right)}{\left((1-\pi^2) \left(\frac{2\pi^2}{1-\pi^2} \right) + 2\pi^2 \right)^2}$ requires $\pi^2 \leq \frac{2}{3}$, i.e. it is always the case in the range of parameters of interest. As a consequence, we have that $\max_{\sigma^i} \frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=0} > \max_{\sigma^i} \frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} \Big|_{\sigma^j=1}$. Since we want $\frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j}$ as big as possible, we can set it as $\frac{1}{2\pi^2}$ for our sufficiency conditions.

Given this, the lowest value the numerator of $\frac{\partial \sigma^i}{\partial \sigma^j}$ from (2.9.A.6) can be found by making the relevant replacement from above to (2.9.A.7). Therefore,

$$\min_{\sigma^i, \sigma^j} \frac{\partial F(c_D^i)}{\partial \sigma^i} \geq \frac{1}{\bar{c} + \varepsilon} \frac{1}{2} \left[\frac{1}{2\pi^2} (-2\pi^2) - 1 \right] = \frac{-1}{\bar{c} + \varepsilon}.$$

To see this, note that $\min_{\sigma^i, \sigma^j} (\sigma^j - (2 - \sigma^j)\pi^2) = -2\pi^2$, $\min_{\sigma^i, \sigma^j} \gamma^i(\emptyset) \geq 0$, $\min_{\sigma^i, \sigma^j} (\gamma^i(\emptyset) - \gamma^i(I)) \geq 0$. Therefore, our first sufficient condition for the uniqueness of the equilibrium is

$$\frac{-\frac{1}{\bar{c} + \varepsilon}}{1 - \frac{1}{\bar{c} + \varepsilon}} > -1,$$

which simplifies to $\bar{c} \geq 2$, as assumed.

Looking now at the upper bound, again by replacing in (2.9.A.7) note that

$$\max_{\sigma^i, \sigma^j} \frac{\partial F(c_D^i)}{\partial \sigma^i} \leq \frac{1}{\bar{c} + \varepsilon} \frac{1}{2} \left[\frac{1}{2\pi^2} (1 - \pi^2) + 1 + \pi^2 - 1 \right] = \frac{1}{2(\bar{c} + \varepsilon)} \left[\frac{1 - \pi^2}{2\pi^2} + \pi^2 \right].$$

To see this, note that $\max_{\sigma^i, \sigma^j} (\sigma^j - (2 - \sigma^j)\pi^2) = 1 - \pi^2$, $\max_{\sigma^i, \sigma^j} \gamma^i(\emptyset) \leq 1$, $\max_{\sigma^i, \sigma^j} (\gamma^i(\emptyset) - \gamma^i(I)) \leq 1$. Therefore, our second sufficient condition for the uniqueness of the equilibrium is

$$\frac{\frac{1}{2(\bar{c} + \varepsilon)} \left[\frac{1 - \pi^2}{2\pi^2} + \pi^2 \right]}{1 - \frac{1}{\bar{c} + \varepsilon}} < 1$$

The numerator is maximised at $\pi = \frac{1}{2}$, hence the condition simplifies to $\bar{c} + \varepsilon > \frac{15}{16}$. Again, this is satisfied for $\bar{c} \geq 2$.

Step 4: Fourth, we show that a symmetric equilibrium where $\sigma^{i*} = \sigma^{j*} = \sigma^*$ always exists. Therefore, it is also unique among the set of signal-based equilibria.

Because of symmetry, the equilibrium must be the fixed point of

$$\sigma^i = \sigma^j = \sigma^* = F(c_D(\sigma^*)) \quad (2.9.A.9)$$

where from (2.1)

$$c_D(\sigma^*) = \frac{1}{2} \left[\left(\frac{(\sigma^*)^2 + (1 - \sigma^*)(2 - \sigma^*)\pi^2}{(\sigma^*)^2 + (2 - \sigma^*)^2\pi^2} - \frac{1 - \sigma^*}{2 - \sigma^*} \right) (\sigma^* - (2 - \sigma^*)\pi^2) + 1 \right] - \frac{1}{2}v$$

Looking at (2.9.A.9), note that both LHS and RHS are continuous on the $[0, 1]$ interval. Moreover, $RHS(\sigma^* = 0) \geq 0 = LHS(\sigma^*)$ and $RHS(\sigma^* = 1) < 1 = LHS(\sigma^* = 1)$. As a consequence, there exists a solution in the $[0, 1]$ interval. From the previous steps, we know that this solution is unique.

Step 5: Finally, we show that in the symmetric equilibrium it is optimal to endorse the state suggested by the most informative signal.

Assume that player j behaves as in the equilibrium described above. Now, by endorsing the wrong state in period 2 player i shifts beliefs from $\gamma^i(2) = 1$ to $\gamma^i(1)$ if it is the only one publishing in that period, and from $\gamma^i(\emptyset)$ to $\gamma^i(1)$ if both outlets publish in period 2. In both cases, sticking to the correct state is at weakly dominant.

If outlet i chooses to publish in period 1, by endorsing the least likely state outlet i is indifferent if it is the only one to publish in that period. If instead outlet j publishes in period 1 as well, the expected reputation of outlet i by endorsing the state suggested by the signal is $\pi^2\gamma^i(\emptyset) + (1 - \pi^2)\gamma^i(1)$. By endorsing the opposite state, the expected reputation is $\pi\gamma^i(1) + (1 - \pi) [\pi\gamma^i(\emptyset) + (1 - \pi)\gamma^i(1)]$. Again, the former is strictly bigger than the latter because $\gamma^i(\emptyset) \geq \gamma^i(1)$. \square

Proof of Lemma 2.2

Proof. To show this, we compare the cost threshold in monopoly and duopoly shutting down the preemption concerns, i.e. assuming $v = 0$. We want to show that in this case $c_D > c_M$. This would require

$$\frac{1}{2} [(\gamma(\emptyset) - \gamma(1)) (\sigma - (2 - \sigma)\pi^2) + 1] > (1 - \pi)(\gamma(R) - \gamma(W)) \quad (2.9.A.10)$$

Observe that $\gamma(1) = \gamma(W) = \frac{1-\sigma}{2-\sigma}$. Moreover, define $\gamma(\emptyset) - \gamma(1) := A$. We can now rearrange equation (2.9.A.10) so that it becomes

$$\frac{1}{2} [A\sigma + 1] > (1 - \pi)(\gamma_R - X) + \frac{1}{2}A(2 - \sigma)\pi^2 \quad (2.9.A.11)$$




Figure 2.A.1: Proof of Lemma 2.2: Proving $\frac{\partial^2 A}{\partial \sigma^2} > 0$. Orange plane: $\frac{\partial^2 A}{\partial \sigma^2}$, blue plane: $0 \cdot \sigma + 0 \cdot \pi$ in $\pi - \sigma$ space.

Now, after the relevant substitutions A can be simplified as $A = \frac{\sigma^2}{(2-\sigma)(\sigma^2 + (2-\sigma)^2\pi^2)}$.

As a consequence,

$$\frac{\partial A}{\partial \sigma} = \frac{2\sigma(2-\sigma)(\sigma^2 + (2-\sigma)^2\pi^2) - \sigma^2(\sigma^2 + 3(2-\sigma)^2\pi^2 - 2\sigma(2-\sigma))}{((2-\sigma)(\sigma^2 + (2-\sigma)^2\pi^2))^2} \quad (2.9.A.12)$$

Signing (2.9.A.12) is not easy in its current form. However, it is clear that $\lim_{\sigma \rightarrow 0} \frac{\partial A}{\partial \sigma} = 0$.

Moreover, we can rearrange A in a more tractable way. In particular, $A = \frac{1}{(2-\sigma)(1+\pi^2 B^2)}$ where $B = \frac{2-\sigma}{\sigma}$. Since $B > 0$ and $\frac{\partial B}{\partial \sigma} = -\frac{2}{\sigma^2} < 0$, it is now easy to see that

$$\frac{\partial A}{\partial \sigma} = \frac{1 + \pi^2 B - 2\pi^2 B \frac{\partial B}{\partial \sigma} (2-\sigma)}{((2-\sigma)(1 + \pi^2 B^2))^2} > 0.$$

The sign of $\frac{\partial^2 A}{\partial \sigma^2}$ is even more complicated, but as A is defined over just two parameters, $\sigma \in [0, 1]$ and $\pi \in (0.5, 1]$, we can prove graphically that $\frac{\partial^2 A}{\partial \sigma^2} > 0$. In particular, Figure 2.A.1 shows that $\frac{\partial^2 A}{\partial \sigma^2}$ (the orange plane) is always strictly above the zero (blue plane) for the entire set of relevant parameters.

It is now straightforward to see that in equation (2.9.A.11) $\frac{\partial \text{LHS}}{\partial \sigma} > 0$ and $\frac{\partial^2 \text{LHS}}{\partial \sigma^2} > 0$ so the LHS is strictly increasing and convex. Moreover, $\frac{\partial \text{RHS}}{\partial \sigma} > 0$.

To complete the proof, we show that $\text{LHS}(\sigma = 0) > \text{RHS}(\sigma = 1)$ for all $\pi \in (0.5, 1)$.

This requires

$$\frac{1}{2} > \frac{1-\pi}{1+\pi} + \frac{1}{2} \frac{\pi^2}{1+\pi^2}$$

which further simplifies to

$$1 - 3\pi + 2\pi^2 - 2\pi^3 < 0$$

Noticing that the LHS of the above is strictly decreasing in π , and it remains negative for both $\pi = \frac{1}{2}$ and $\pi = 1$, completes the proof. \square

Proof of Proposition 2.5

Proof. This follows directly from the strict inequality of equation (2.9.A.10) and the fact that v only reduces its LHS, without affecting the RHS. \square

Proof of Corollary 2.1

Proof. The behavior of the monopolist is unchanged with respect to Section 2.3.1. Looking at the duopoly case, by Bayes' rule

$$\gamma^i(R, \cdot) = \frac{(1 - \sigma^i)\pi + \sigma^i}{(1 - \sigma^i)\pi + \sigma^i + \pi} = \gamma^i(R)$$

$$\gamma^i(W, \cdot) = \frac{1 - \sigma^i}{2 - \sigma^i} = \gamma^i(W)$$

Therefore, the cost threshold for research is given by

$$\begin{aligned} & \frac{1}{2} \left[\sigma^j \left(\frac{v}{2} + \gamma^i(R) \right) + (1 - \sigma^j) \gamma^i(R) \right] + \frac{1}{2} \gamma^i(R) - c \geq \\ & \frac{1}{2} \left[\sigma^j \left(v + \pi \gamma^i(R) + (1 - \pi) \gamma^i(W) \right) + (1 - \sigma^j) \left(\frac{v}{2} + \pi \gamma^i(R) + (1 - \pi) \gamma^i(W) \right) \right] + \frac{1}{2} \left(\frac{v}{2} + \pi \gamma^i(R) + (1 - \pi) \gamma^i(W) \right), \end{aligned}$$

which simplifies to

$$c \leq (1 - \pi)(\gamma^i(R) - \gamma^i(W)) - \frac{1}{2}v := c'_D \quad (2.9.A.13)$$

Note that the first part of (2.9.A.13) is the same as c_M , and the only term that changes is $-\frac{1}{2}v$, making it smaller than c_M .

In terms of existence and uniqueness of the equilibrium in this set up, note that σ^{i*} and σ^{j*} are the solution of the same fixed point problem, i.e.

$$\sigma^* = F(c'_D(\sigma^*))$$

where $c'_D = c_M - \frac{1}{2}v$. The same logic of the proof of Proposition 2.2 applies here as well. Hence the equilibrium exists and it is unique and symmetric. \square

Proof of Corollary 2.2

Proof. Consider first the case of monopoly. Here, only the high quality outlet can publish in period 2, and this is observable. As a consequence,

$$\gamma(2) = 1$$

$$\gamma(1) = \frac{1 - \sigma}{2 - \sigma}$$

The monopolist chooses to investigate when $c \leq 1 - \gamma(1) := c_M''$.

In duopoly, the beliefs are updated the same way. Each outlet is considered independently and only the timing matters. The threshold is, therefore, given by

$$\frac{1}{2} \left[\sigma^j \left(\frac{v}{2} + 1 \right) + (1 - \sigma^j) \right] + \frac{1}{2} - c \geq \frac{1}{2} \left[\sigma^j (v + \gamma^i(1)) + (1 - \sigma^j) \left(\frac{v}{2} + \gamma^i(1) \right) \right] + \frac{1}{2} \left(\frac{v}{2} + \gamma^i(1) \right).$$

It follows then that $c_D'' = 1 - \gamma^i(1) - \frac{1}{2}v = c_M'' - \frac{1}{2}v < c_M''$ as claimed.

In terms of existence and uniqueness, note that σ^* is the solution of

$$\sigma^* = F(c''(\sigma^*))$$

The RHS is continuous on the $[0, 1]$ interval and, irrespective of the market structure, it is either strictly increasing and convex or flat. Moreover, $\text{RHS}(\sigma^* = 0) \geq \text{LHS}(\sigma^* = 0)$ and $\text{RHS}(\sigma^* = 1) < 1 = \text{LHS}(\sigma^* = 1)$ since $\bar{c} > 1$. \square

Proof of Proposition 2.6

Proof. We proceed in steps as outlined in Proposition 2.4. We drop the bars from σ for convenience.

Step 1: We begin by showing that in any signal-based equilibria outlets' period 1 decision on whether to research or publish is described by a threshold on c . This follows from the discussion in the text. Let σ^i and σ^j be the conjectured strategies. Then equation (2.2) defines the threshold c_D^i for outlet i .

$$c^i \leq \frac{1}{2} \left[(\gamma^i(\emptyset) - \gamma^i(1)) (\sigma^j - (2 - \sigma^j)\pi^2) + (2 - \sigma^j) (1 - \gamma^i(1)) - \sigma^j(1 - u) \right] + 1 - \frac{3}{2}u := \bar{c}_D^i \quad (2.2)$$

where $\gamma(\emptyset) = \frac{\sigma^i \sigma^j + (1 - \sigma^i)(2 - \sigma^j)\pi^2}{\sigma^i \sigma^j + (2 - \sigma^i)(2 - \sigma^j)\pi^2}$ and $\gamma(1) = \frac{1 - \sigma^i}{2 - \sigma^i}$. The problem is identical for player j .

Step 2: Next, we show that for any σ^j there is only one σ^i that solves the equilibrium fixed

point for player i .

All of the definitions from Proposition 2.4 remain unaltered.

We want to show that, for every σ^j , there is only one σ^i that solves $\sigma^i = F(\bar{c}_D^i(\sigma^i, \sigma^j))$.

1. The LHS is linear, with slope equal to 1, starting at 0 and ending at 1.
2. Now, $\bar{c}_D^i(\sigma^i = 1, \sigma^j) = c_D^i(\sigma^i = 1, \sigma^j, v = 0) + (1 - u)(1 - \frac{\sigma^j}{2})$, where each term is less than or equal to 1. But since $\bar{c} \geq 2.5$, therefore $\bar{c}_D^i(\sigma^i = 1, \sigma^j) < \bar{c}$. As a result, the RHS evaluated at $\sigma^i = 1 < 1 = \text{LHS at } \sigma^i = 1$;
3. The RHS evaluated at $\sigma^i = 0$ is greater than or equal to zero.
4. For any σ^j , both LHS and RHS are continuous in σ^i .

Hence, they cross at least once and there is at least one solution to this fixed point problem.

Further, note that \bar{c}_D^i behaves the same way as c_D^i with respect to σ^i . Therefore, the rest of the proof in this step is as before.

Step 3: Third, we show that if an equilibrium exists, it is unique for $\bar{c} \geq 2.5$.

Other than changing the relevant definitions to include σ , nothing changes in this step until we evaluate $\frac{\partial \bar{c}_D^i}{\partial \sigma^j}$

$$\frac{\partial \bar{c}_D^i}{\partial \sigma^j} = \frac{1}{2} \left[\frac{\partial \gamma^i(\emptyset)}{\partial \sigma^j} (\sigma^j - (2 - \sigma^j)\pi^2) + \gamma^i(\emptyset) + \pi^2(\gamma^i(\emptyset) - \gamma^i(1)) - (2 - u) \right]. \quad (2.9.A.14)$$

Again, the rest of the proof remains unaltered until we find the first sufficient condition. The lowest value of the numerator of $\frac{\partial \sigma^i}{\partial \sigma^j}$ from (2.9.A.6) can be found by making the relevant replacement from above to (2.9.A.14). Therefore,

$$\min_{\sigma^i, \sigma^j} \frac{\partial F(\bar{c}_D^i)}{\partial \sigma^j} \geq \frac{1}{\bar{c} + \varepsilon} \frac{1}{2} \left[\frac{1}{2\pi^2} (-2\pi^2) - (2 - u) \right] = \frac{-1}{\bar{c} + \varepsilon} \left(\frac{3 - u}{2} \right).$$

Therefore, our new first sufficient condition for the uniqueness of the equilibrium is

$$\frac{-\frac{1}{\bar{c} + \varepsilon} \left(\frac{3 - u}{2} \right)}{1 - \frac{1}{\bar{c} + \varepsilon}} > -1,$$

which simplifies to $\bar{c} \geq \frac{5 - u}{2}$. The highest value possible of $\frac{5 - u}{2}$ is 2.5 at $u = 0$, which is assumed.

Looking now at the upper bound, again by replacing in (2.9.A.14) we get

$$\max_{\sigma^i, \sigma^j} \frac{\partial F(\bar{c}_D^i)}{\partial \sigma^i} \leq \frac{1}{\bar{c} + \varepsilon} \frac{1}{2} \left[\frac{1}{2\pi^2} (1 - \pi^2) + 1 + \pi^2 - 2 + u \right] = \frac{1}{2(\bar{c} + \varepsilon)} \left[\frac{1}{2\pi^2} + \pi^2 - \frac{3}{2} + u \right].$$

Therefore, our second new sufficient condition for the uniqueness of the equilibrium is

$$\frac{\frac{1}{2(\bar{c} + \varepsilon)} \left[\frac{1}{2\pi^2} + \pi^2 - \frac{3}{2} + u \right]}{1 - \frac{1}{\bar{c} + \varepsilon}} < 1$$

The numerator is maximised at $\pi = \frac{1}{\sqrt{2}}$, hence the condition simplifies to $\bar{c} + \varepsilon > \frac{u+2}{2}$. Again, this is satisfied for $\bar{c} \geq 2.5$ since $\frac{5-u}{2} > \frac{u+2}{2}$ for $u \in [0, 1]$.

Step 4: Fourth, we show that a symmetric equilibrium where $\sigma^{i*} = \sigma^{j*} = \sigma^*$ always exists. Therefore, it is also unique among the set of signal-based equilibria.

Because of symmetry, the equilibrium must be the fixed point of

$$\sigma^i = \sigma^j = \sigma^* = F(c_D(\sigma^*)) \quad (2.9.A.15)$$

where from (2.2)

$$c_D(\sigma^*) = \frac{1}{2} \left[\left(\frac{(\sigma^*)^2 + (1 - \sigma^*)(2 - \sigma^*)\pi^2}{(\sigma^*)^2 + (2 - \sigma^*)^2\pi^2} - \frac{1 - \sigma^*}{2 - \sigma^*} \right) (\sigma^* - (2 - \sigma^*)\pi^2) - \sigma^*(1 - u) \right] - \frac{3}{2}(1 - u) \quad (2.9.A.16)$$

Looking at (2.9.A.15), note that both LHS and RHS are continuous on the $[0, 1]$ interval. Moreover, $\text{RHS}(\sigma^* = 0) \geq 0 = \text{LHS}(\sigma^*)$ and $\text{RHS}(\sigma^* = 1) < 1 = \text{LHS}(\sigma^* = 1)$. As a consequence, there exists a solution in the $[0, 1]$ interval. From the previous steps, we know that this solution is unique.

Step 5: Finally, we show that in the symmetric equilibrium it is optimal to endorse the state suggested by the most informative signal.

This is true because now there is more incentive to build a reputation. Since reputation requires matching the state, there is even less reason to not endorse the state suggested by the most informative equilibrium. \square

Proof of Proposition 2.7

Proof. We drop the bars for convenience. First, note that \bar{c}_D is decreasing in u . This is so because it can be rearranged as

$$\bar{c}_D = \frac{1}{2} [(\gamma(\emptyset) - \gamma(1)) (\sigma^* - (2 - \sigma^*)\pi^2) - \sigma^*] + \frac{3}{2} - u \left(\frac{3}{2} - \frac{\sigma^*}{2} \right)$$




Figure 2.A.2: Proof of Proposition 2.7: Proving $LHS < RHS$. Orange plane: RHS, blue plane: LHS and green plane: $0.\sigma + 0.\pi$ in the $\pi - \sigma$ space.

where $\frac{3}{2} - \frac{\sigma^*}{2} > 0$ for any $\sigma^* \in [0, 1]$. Also, c_M and σ_M^* do not change with u .

Second, consider the case when $u = 1$. We will show that $\bar{c}_D < c_M$. This requires

$$\frac{1}{2} [(\gamma(\emptyset) - \gamma(1))(\sigma - (2 - \sigma)\pi^2)] < (1 - \pi)(\gamma(R) - \gamma(W)).$$

Using the terminology introduced in Lemma 2.2, we can rewrite the above as

$$\frac{1}{2} A\sigma < (1 - \pi)(\gamma(R) - \gamma(W)) + \frac{1}{2} A(2 - \sigma)\pi^2.$$

Now, both the LHS and the RHS of the above equation are only functions of two variables, π and σ , which are defined on compact and continuous sets. Therefore, we can plot them in a graph (see Figure 2.A.2) and check that the above is true.

Third, consider the case of $u = 0$. We want to show that $\bar{c}_D > c_M$. This is equivalent to showing that

$$\frac{1}{2} [(\gamma(\emptyset) - \gamma(1))(\sigma - (2 - \sigma)\pi^2) - \sigma] + \frac{3}{2} > (1 - \pi)(\gamma(R) - \gamma(W)).$$

We showed in Lemma 2.2 that $c_D(v = 0) > c_M$. It is easy to check that $\bar{c}_D(u = 0) = c_D(v = 0) + 1 - \frac{1}{2}\sigma$ where $1 - \frac{1}{2}\sigma > 0$ for all $\sigma \in [0, 1]$. Therefore, $\bar{c}_D(u = 0) > c_D(v = 0) > c_M$.

Combining the three parts above, our result follows. \square

Proof of Proposition 2.8

Proof. We drop the bars for convenience. Reminding ourselves that

$$V = \left(\frac{(2 - \sigma^*)}{4} \right) \pi u + \frac{1}{2} \sigma^* (1 - u) + \left(\frac{1}{2} \right)^2 (1 - (1 - \sigma^*)^2) u,$$

we first take the first derivative of V with respect to π (we drop the stars and D in what follows for convenience).

$$\begin{aligned} \frac{\partial V}{\partial \pi} &= u \frac{(2 - \sigma)^2}{4} + \left[-\pi u \frac{(2 - \sigma)}{2} + \frac{1 - u}{2} + u \frac{(1 - \sigma)}{2} \right] \frac{\partial \sigma}{\partial \pi} \\ &= u \frac{(2 - \sigma)^2}{4} + \frac{1 - u(2\pi + \sigma(1 - \pi))}{2} \frac{\partial \sigma}{\partial \pi} \end{aligned} \quad (2.9.A.17)$$

Now, we need to show under what conditions $\frac{\partial \sigma}{\partial \pi} < 0$. Reminding that σ is implicitly defined by (2.9.A.15) define

$$K := \sigma - \left[\frac{1}{2} [(\gamma(\emptyset) - \gamma(1)) (\sigma - (2 - \sigma)\pi^2) - \sigma(1 - u)] + \frac{3}{2}(1 - u) \right] \frac{1}{\bar{c} + \varepsilon} - \frac{\varepsilon}{\bar{c} + \varepsilon}$$

Further, using the definitions in the proof of Lemma 2.2, we can rewrite K as

$$K = \sigma - \frac{1}{2(\bar{c} + \varepsilon)} [A(\sigma - (2 - \sigma)\pi^2) - \sigma(1 - u)] - \frac{3}{2(\bar{c} + \varepsilon)}(1 - u) - \frac{\varepsilon}{2(\bar{c} + \varepsilon)}$$

Differentiating and simplifying, we first obtain

$$\begin{aligned} \frac{\partial K}{\partial \pi} &= -\frac{1}{2(\bar{c} + \varepsilon)} \left[\frac{-2\pi B^2(\sigma - (2 - \sigma)\pi^2)}{(2 - \sigma)(1 + \pi^2 B^2)^2} - \frac{2\pi}{1 + \pi^2 B^2} \right] \\ &= \frac{1}{2(\bar{c} + \varepsilon)} \frac{1 + B}{(1 + \pi^2 B^2)^2} > 0, \end{aligned}$$

and second we obtain

$$\begin{aligned} \frac{\partial K}{\partial \sigma} &= 1 - \frac{1}{2(\bar{c} + \varepsilon)} \left[\frac{\partial A}{\partial \sigma} (\sigma - (2 - \sigma)\pi^2) + (1 + \pi^2)A - (1 - u) \right] \\ &= 1 + \frac{1}{2(\bar{c} + \varepsilon)}(1 - u) - \frac{1}{2(\bar{c} + \varepsilon)} \left[\frac{\partial A}{\partial \sigma} (\sigma - (2 - \sigma)\pi^2) + (1 + \pi^2)A \right] \\ &= 1 + \frac{1}{2(\bar{c} + \varepsilon)}(1 - u) - \frac{1}{\bar{c} + \varepsilon} \frac{\partial c_D}{\partial \sigma} \end{aligned}$$

where c_D is the cost threshold we derived in Proposition 2.4.

We can now show that $\frac{\partial c_D}{\partial \sigma} \leq 1$ in the neighborhood of the equilibrium σ . The proof for this is presented in Proposition 2.12 (Appendix C) for a generic prior q . Therefore, it is also

true in our special case of $q = \frac{1}{2}$.

Putting these two facts together and using the Implicit Function Theorem, we can now conclude that $\frac{\partial \sigma}{\partial \pi} < 0$.

Finally, we want to find the condition under which $\frac{\partial V}{\partial \pi} < 0$. From (2.9.A.17), this happens when

$$u \frac{(2 - \sigma)^2}{4} < \frac{1 - u(2\pi + \sigma(1 - \pi))}{2} \left(-\frac{\partial \sigma}{\partial \pi} \right),$$

where $\left(-\frac{\partial \sigma}{\partial \pi} \right) := \sigma_\pi > 0$. This can then be rearranged to

$$u < \left(\frac{(2 - \sigma)^2}{2} \frac{1}{\sigma_\pi} + 2\pi + \sigma(1 - \pi) \right)^{-1} := \bar{u}^V$$

Now, it is easy to see that the denominator of $\bar{u}^V > 1$ because $2\pi > 1$. Therefore, \bar{u}^V exists and lies between 0 and 1. \square

Proof of Lemma 2.3

Proof. Comparing the source's expected utility given by expressions in (2.3) and (2.4) and simplifying gives the following condition to prefer two firms:

$$\mu(1 - \pi)(2\sigma_M - \sigma_D(4 - \sigma_D)) \leq 2\sigma_M - \sigma_D^2 \quad (2.9.A.18)$$

Now we discuss different cases based on possible values of σ_M and σ_D .

Case 1: v is very high so that $\sigma_D = 0$. Substituting in 2.9.A.18 gives that the source prefers to send the story to both outlets if

$$\mu \leq \frac{1}{1 - \pi} > 1$$

Therefore, if v is very large it is possible that $\mu > 1$ (so that the source cares relatively more about matching the state) and $\sigma_D = 0$ (so that in duopoly no one does research), but still the source prefers to share information with both the outlets. This happens because $\pi > .5$ and the source still cares about getting the information out quickly.

Case 2: v is high enough so that $\sigma_D < \sigma_M$. Now, the RHS of equation (2.9.A.18) is greater than zero. But first, σ_D might not be too small so that in the LHS < 0 i.e. $2\sigma_M \leq \sigma_D(4 - \sigma_D)$. In this case, sending to both is always preferred independent of μ . Therefore, sending to both is preferred if

$$\sigma_D < \sigma_M \leq \frac{\sigma_D(4 - \sigma_D)}{2}.$$

Second, σ_D might in fact be very small so that on the LHS > 0 i.e. $2\sigma_M > \sigma_D(4 - \sigma_D)$.

In this case, sending to both is preferred only if

$$\mu \leq \frac{2\sigma_M - \sigma_D^2}{2\sigma_M - \sigma_D(4 - \sigma_D)} \frac{1}{(1 - \pi)}.$$

Case 3: v is small so that $\sigma_D > \sigma_M$. Again there are two possible situations. First, consider the case in which σ_D is not too large so that the RHS of equation (2.9.A.18) is still positive, i.e. $2\sigma_M \geq \sigma_D^2 \implies \sigma_M \geq \frac{\sigma_D^2}{2}$. Now, in this case we want to see whether the LHS can be negative i.e. if $\sigma_M < \frac{\sigma_D(4 - \sigma_D)}{2}$. But this must be true because $\sigma_D > \sigma_M$ and we know that $\frac{\sigma_D(4 - \sigma_D)}{2} > \sigma_D$. Therefore, the LHS is negative and the RHS is positive, so the condition outlined in (2.9.A.18) is satisfied. Sending to both is always preferred if

$$\frac{\sigma_D^2}{2} \leq \sigma_M < \sigma_D.$$

Second, σ_D might in fact be very large so that the RHS is negative, i.e. $\sigma_M < \frac{\sigma_D^2}{2}$. Now, it cannot be that the LHS is positive because that requires $\sigma_M > \frac{\sigma_D(4 - \sigma_D)}{2}$ which contradicts $\sigma_M < \sigma_D$. Therefore, LHS must also be negative. From condition (2.9.A.18), the source prefers both outlets only if

$$\mu \geq \frac{\sigma_D^2 - 2\sigma_M}{\sigma_D(4 - \sigma_D) - 2\sigma_M} \frac{1}{(1 - \pi)}.$$

Case 4: v is such that $\sigma_D = \sigma_M := \sigma$. When this is the case, the condition (2.9.A.18) reduces to

$$-\mu(1 - \pi)(2 - \sigma) < (2 - \sigma)$$

which is always true. Therefore, sending to both is preferred.

Our result follows from combining all the above cases. \square

Proof of Proposition 2.9

Proof. The proof is by construction. We have already constructed the equilibrium frontier and the set of all possible equilibria for a given \bar{c} and ε .

We now show what happens as $\varepsilon \rightarrow 0$. Consider σ_M first. From Proposition 2.2, observe that as $\varepsilon \rightarrow 0$ $\text{LHS}(\sigma = 0) = 0 \approx \text{RHS}(\sigma = 0) = \frac{\varepsilon}{\bar{c} + \varepsilon} \rightarrow 0$ in equation (2.9.A.3). Therefore, for any π the only fixed point equilibrium $\rightarrow 0$.

Now, consider σ_D at $v = 0$. Fix a π . We know that as $\varepsilon \rightarrow 0$, since $c_D(\sigma = 0) = \frac{1}{2}$, we have that $\text{RHS}(\sigma = 0) \rightarrow \frac{1}{2\bar{c}}$ in equation (2.9.A.9). But this is strictly greater than $\text{LHS}(\sigma = 0) = 0$. Therefore, the equilibrium fixed point $\sigma_D > 0$ and also $\frac{\sigma_D^2}{2} > 0$. Moreover, this is true for any π .

Therefore, in the $\sigma_D - \sigma_M$ space as $\varepsilon \rightarrow 0$, the equilibrium frontier lies below the $\sigma_M = \frac{\sigma_D^2}{2}$ line.

Now, let us look at what happens as $\varepsilon \rightarrow \infty$. Given that the fixed point is defined by $\sigma^* = \frac{c^* + \varepsilon}{\bar{c} + \varepsilon}$, both σ_M and σ_D approach 1 (without ever being exactly equal to 1). However, because the frontier is defined for $v = 0$ case, the frontier lies close to and to the right of the $\sigma_M = \sigma_D$ line.

Combining the two observations above with Lemma 2.3, we get our proposition. \square

B Allowing for sitting on information

In this appendix, we show that allowing outlets to “sit on information” (i.e. just refrain from publishing until period 2 without acquiring the additional signal) does not preclude the equilibrium outlined in Proposition 2.4. We prove it formally for sufficiently low π and then use mathematical simulation to argue that it holds more generally. Uniqueness of such an equilibrium (among signal-based equilibria), however, is not obvious anymore. We make only one change with respect to the model described in Section 2.2. Now $d^i \in \{res, pub, wait\}$, where $d^i = wait$ means that the outlet does not acquire the second signal but still publishes in period 2.

This addition poses some challenges in the tractability of the model because the choice is no longer just between two options and strategies are not necessarily just thresholds in c . However, even in this more complicated setup we can show a few results. First, for sufficiently low π , it is possible to find values of v such that the equilibrium described in Proposition 2.4 exists; waiting is never a best response if the other player never waits and $\sigma_D^* > \sigma_M^*$. Second, we can simulate the model showing that we can assign values to v such that, for the resulting equilibrium σ_D^* , publishing in period 1 is better than waiting and at the same time $\sigma_D^* > \sigma_M^*$.

We begin with the following lemma considering that we are interested in the (candidate) equilibrium strategies described in Proposition 2.4 where $d^i = wait$ is never played in equilibrium.

Lemma 2.4. *It is always possible to find off path beliefs such that, for sufficiently high v , $Eu^i(d^i = wait) \leq Eu^i(d^i = pub)$.*

Proof. Note that $\gamma(W_{II}, \cdot)$ is off-path in the equilibrium we are considering. For any $\gamma(\emptyset)$ and $\gamma(1)$ as defined above, the expected utility from choosing $d^i = wait$ is

$$\frac{1}{2}\sigma^j \left(\frac{v}{2} + \pi\gamma(\emptyset) + (1-\pi)\gamma(1) \right) + \left(\frac{1}{2}(1-\sigma^j) + \frac{1}{2} \right) (\pi + (1-\pi)\gamma(W_{II}, \cdot)) \quad (2.9.B.19)$$

On the other hand, the expected utility from publishing immediately is given by

$$\frac{1}{2}\sigma^j (v + \gamma(1)) + \left(\frac{1}{2}(1-\sigma^j) + \frac{1}{2} \right) \left(\frac{v}{2} + \pi^2(\gamma(\emptyset) - \gamma(1)) + \gamma(1) \right) \quad (2.9.B.20)$$

Comparing (2.9.B.19) and (2.9.B.20) and solving for v , we find that $Eu^i(d^i = wait) \leq$

$Eu^i(d^i = pub)$ when

$$v \geq \sigma^j \pi (\gamma(\emptyset) - \gamma(1)) - (2 - \sigma^j) [\pi^2 \gamma(\emptyset) + (1 - \pi^2) \gamma(1) - \pi - (1 - \pi) \gamma(W_{II}, .)] \quad (2.9.B.21)$$

Therefore, it is possible to find v and $\gamma(W_{II}, .)$ such that the above condition is satisfied. \square

This makes intuitive sense as a sufficiently high scoop value should always deter sitting on information. From now on, we set $\gamma(W_{II}, .) = 0$ and we define $\bar{v} := \sigma^j \pi (\gamma(\emptyset) - \gamma(1)) - (2 - \sigma^j) [\pi^2 \gamma(\emptyset) + (1 - \pi^2) \gamma(1) - \pi]$.

We can now move to the main proposition.

Proposition 2.10. *For sufficiently low π , it is possible to find values of v such that the equilibrium described in Proposition 2.4 exists. In such an equilibrium, waiting is never a best response if the other player follows the equilibrium strategies and $\sigma_D^* > \sigma_M^*$.*

Proof. Suppose that player j always publishes when low type and chooses just between publishing or researching when high type. Moreover, suppose that the audience conjectures that both players use the equilibrium strategies described by Proposition 2.4. For this to be an equilibrium in the new setup, it is sufficient to prove that, given the correct audience's beliefs updating, for every σ , $Eu^i(d^i = wait) \leq Eu^i(d^i = pub)$. To show this, first we prove through Figure 2.B.3 that $\frac{\partial \bar{v}}{\partial \pi} > 0$. Moreover, Figure 2.B.4 shows that there exists a range of π such that $\argmax_{\sigma} \bar{v}(\pi) = 1$. In the figure, it happens for $\pi \in [.5, .6]$. As a consequence, for every $v \geq \bar{v}(\sigma = 1, \pi \in [0.5, 0.6])$ it is true that, for every σ , $Eu^i(d^i = wait) \leq Eu^i(d^i = pub)$. In other words, if the audience conjectures an equilibrium where no types and no players choose to wait and the choice for the high type is just between publishing and researching described by a threshold strategy on c , behaving in this way is an equilibrium strategy for the outlets.

Finally, Figure 2.B.5 plots c_M and $c_D(\bar{v}(\sigma = 1))$ for sufficiently small π , proving that we can still increase v from $\bar{v}(\sigma)$ maintaining the necessary condition for $\sigma_D^* > \sigma_M^*$, i.e. $c_D \geq c_M$. \square

When $\pi > 0.6$, we can show the existence of our candidate equilibrium through mathematical simulations. Consider for example the following set of parameters: $\pi = 0.75$, $v = 0.7$, $\bar{c} = 2$, $\varepsilon = 0.1$. In this case, the equilibrium described in Proposition 2.4 (assuming it still exists) gives a solution $\sigma_D^* = 0.118219$.¹⁴ Suppose now that player i expects player j to never wait and choose to research (if it is high type) with probability 0.118219. Further,

¹⁴We simulated the model with Mathematica. The code is available upon request.




Figure 2.B.3: Proof of Proposition 2.10: Proving $\frac{\partial \bar{v}}{\partial \pi} > 0$. Orange plane: $\frac{\partial \bar{v}}{\partial \pi}$, blue plane: $0.5\sigma + 0.5\pi$ in the $\pi - \sigma$ space.




Figure 2.B.4: Proof of Proposition 2.10: Proving $\operatorname{argmax}_{\sigma} \bar{v}(\pi) = 1$. Orange plane: $\bar{v}(\pi)$, blue plane: $0.5\sigma + 0.5\pi$ in the $\pi - \sigma$ space for $\pi \in [0.5, 0.6]$.




Figure 2.B.5: Proof of Proposition 2.10: Proving $c_D(\bar{v}(\sigma = 1)) > c_M$. Orange plane: $c_D(\bar{v}(\sigma = 1))$, blue plane: c_M in the $\pi - \sigma$ space for $\pi \in [0.5, 0.6]$.

suppose the audience think that both outlets never wait and research (if they are high types) with probability 0.118219. In this case, $Eu^i(d^i = wait) = 0.754218$ and $Eu^i(d^i = pub) = 0.841236$. Hence, there is no incentive to choose waiting instead of publishing, and the meaningful choice is just between researching and publishing. The solution to this problem is the same as that described by Proposition 2.4. Finally, Figure 2.B.6 shows that, for $\pi = 0.75$ and $v = 0.7$, it is still true that $c_D \geq c_M$ for every σ . More generally, Figure 2.B.7 plots c_M and c_D (in the $\pi - \sigma$ space) by replacing v with the corresponding \bar{v} . Still, c_D is above c_M throughout the entire range of parameters of our model.




Figure 2.B.6: $c_D > c_M$ for $\pi = 0.75$ and $v = 0.7$. Orange line: c_M , blue line: c_D as a function of σ)




Figure 2.B.7: $c_D(v = \bar{v}) > c_M$ for every combination of σ and π . Orange plane: $c_D(v = \bar{v})$, blue plane: c_M in the $\pi - \sigma$ space.

C Generic prior on the type

This appendix shows that our main results are qualitatively unaffected by the assumption of prior $\Pr(\theta^i = h) = \frac{1}{2}$. In this section, we assume a generic prior $\Pr(\theta^i = h) = q \in (0, 1)$, leaving the rest of the model unchanged. We consider monopoly, duopoly and their comparison for when θ is unknown to the reader.

Monopoly

The proposition of the main result is unchanged in monopoly, as q enters only in the readers' beliefs updating.

Proposition 2.11. *If there is one media outlet and θ is not known to the audience, there exists a unique equilibrium in which the high quality outlet conducts research in $t = 1$ iff*

$$c \leq (1 - \pi)(\gamma(R) - \gamma(W)) := c_M$$

where $\gamma(R)$ and $\gamma(W)$ are the audiences' beliefs about the outlet's quality after it gets the state right and wrong respectively. As a consequence, $\sigma^* = F(c_M(q)) = \frac{c_M(q, \sigma^*) + \varepsilon}{\bar{c} + \varepsilon}$.

Proof. Suppose that a high type outlet chooses $d = res$ with probability σ . Reminding ourselves from the main text that by Bayes' rule,

$$\begin{aligned}\gamma(R) &= \frac{q(\sigma + (1 - \sigma)\pi)}{q(\sigma + (1 - \sigma)\pi) + (1 - q)\pi} \\ \gamma(W) &= \frac{q(1 - \sigma)(1 - \pi)}{q(1 - \sigma)(1 - \pi) + (1 - q)(1 - \pi)} = \frac{q(1 - \sigma)}{1 - q\sigma}.\end{aligned}$$

A high quality optimally chooses res if

$$\gamma(R) - c \geq \pi\gamma(R) + (1 - \pi)\gamma(W) \implies c \leq (1 - \pi)(\gamma(R) - \gamma(W)) := c_M(q)$$

In equilibrium the conjectured σ must be equal to the actual one, hence it must be that

$$\sigma^* = F(c_M(q, \sigma^*)). \quad (2.9.C.22)$$

We need to check if such a fixed point exists. To do so, three observations are in order. First, note that both the LHS and RHS of the above are continuous in σ^* . Second, $\text{LHS}(\sigma^* = 0) = 0 < \text{RHS}(\sigma^* = 0) = \frac{\varepsilon}{\bar{c} + \varepsilon}$ (as $c_M(q) = 0$ at $\sigma^* = 0$). Third, $\text{LHS}(\sigma^* = 1) = 1 > \text{RHS}(\sigma^* = 1) = F\left(\frac{(1 - \pi)q}{q + (1 - q)\pi}\right)$, so the equilibrium is the solution of $\sigma^* = \frac{c_M(q, \sigma^*) + \varepsilon}{\bar{c} + \varepsilon}$ and LHS and RHS must cross at least once.

Finally, we need to check for the uniqueness of the fixed point. To show this, it is sufficient to prove that the derivative of the RHS with respect to σ is smaller than 1. Note that

$$\frac{\partial \text{RHS}}{\partial \sigma^*} = \frac{1 - \pi}{\bar{c} + \varepsilon} \left[\frac{\pi(1 - \pi)q(1 - q)}{(q(\sigma + (1 - \sigma)\pi) + (1 - q)\pi)^2} + \frac{q(1 - q)}{(1 - q\sigma^*)^2} \right] > 0$$

Moreover, we can rewrite the equation as

$$\frac{\partial \text{RHS}}{\partial \sigma^*} = \frac{(1 - \pi)q(1 - q)}{\bar{c} + \varepsilon} \left[\frac{\pi(1 - \pi)}{(q(\sigma + (1 - \sigma)\pi) + (1 - q)\pi)^2} + \frac{1}{(1 - q\sigma^*)^2} \right]$$

It is easy to see that, in the range of parameters of the model, $(1 - \pi)q(1 - q) \leq \frac{1}{8}$; $\frac{\pi(1 - \pi)}{(q(\sigma + (1 - \sigma)\pi) + (1 - q)\pi)^2} \leq 1$ because $\pi(1 - \pi)$ is at most $\frac{1}{4}$ and $q(\sigma + (1 - \sigma)\pi) + (1 - q)\pi$ is at least $\frac{1}{2}$ (when $\sigma = 0$ and $\pi = \frac{1}{2}$); $\frac{1}{(1 - q\sigma^*)^2} \leq 1$. As a consequence,

$$\frac{\partial \text{RHS}}{\partial \sigma^*} < \frac{1}{8} [1 + 1] < 1$$

and this completes the proof. \square

Duopoly

For the case of duopoly, we look directly at symmetric equilibria, showing that there exists a unique symmetric equilibrium.

Proposition 2.12. *If there are two media outlets and θ is not known to the audience, there exists a unique symmetric equilibrium where $\sigma^{i*} = \sigma^{j*} := \sigma^* = F(c_D(q))$ such that*

$$c_D(q) = [(\gamma(\emptyset) - \gamma(1)) (q\sigma^* - (1 - q\sigma^*)\pi^2) + 1 - q] - \frac{1}{2}v$$

where $\gamma(\emptyset) = \frac{q((\sigma^*)^2 q + (1 - \sigma^*)(1 - q\sigma^*)\pi^2)}{(q\sigma^*)^2 + (1 - q\sigma^*)^2 \pi^2}$ and $\gamma(1) = \frac{q(1 - \sigma^*)}{1 - q\sigma^*}$.

Proof. We focus directly on symmetric equilibria where each high type outlet uses a threshold strategy on c in the decision on whether to publish or investigate. Define σ as the probability (from the point of view of the other players) that a high quality outlet chooses to do research. For the same logic as in Proposition 2.4, the threshold is given by

$$c^i \leq [(\gamma(\emptyset) - \gamma(1)) (q\sigma - (1 - q\sigma)\pi^2) + (1 - q\sigma)(1 - \gamma(1))] - \frac{1}{2}v := c_D(q) \quad (2.1)$$

where, by Bayes' rule, $\gamma(\emptyset) = \frac{q(\sigma^2 q + (1 - \sigma)(1 - q\sigma)\pi^2)}{q^2 \sigma^2 + (1 - q\sigma)^2 \pi^2}$ and $\gamma(1) = \frac{q(1 - \sigma)}{1 - q\sigma}$.

Given that cost is uniformly distributed in $[-\varepsilon, \bar{c}]$ and that in equilibrium the conjectured

probability of investigation must be equal to the actual one, the (symmetric) equilibrium level of σ , if it exists, must be the solution of

$$\sigma = F(c_D(q, \sigma)) \quad (2.9.C.23)$$

where

$$F(c_D(q, \sigma)) = \begin{cases} 0 & c_D(q, \sigma) < -\varepsilon \\ \frac{c_D(q, \sigma) + \varepsilon}{\bar{c} + \varepsilon} & -\varepsilon \leq c_D(q, \sigma) \leq \bar{c} \\ 1 & c_D(q, \sigma) > \bar{c} \end{cases}$$

and

$$f(c_D(q, \sigma)) = \begin{cases} 0 & c_D(q, \sigma) < -\varepsilon \\ \frac{1}{\bar{c} + \varepsilon} & -\varepsilon \leq c_D(q, \sigma) \leq \bar{c} \\ 0 & c_D(q, \sigma) > \bar{c} \end{cases}$$

Note that:

1. The LHS of equation (2.9.C.23) is linear, with slope equal to 1, starting at 0 and ending at 1;
2. $\text{RHS}(\sigma = 0) \geq 0 = \text{LHS}(\sigma = 0)$;
3. $\text{RHS}(\sigma = 1) < 1 = \text{LHS}(\sigma = 1)$;
4. Both LHS and RHS are continuous in σ .

Hence, they cross at least once and there is at least one solution to this fixed point problem.

To show uniqueness, we can rewrite $c_D(q)$ as

$$c_D = AE + 1 - q - \frac{1}{2}v$$

where $A := \gamma(\emptyset) - \gamma(1) = \frac{q^2(1-q)}{(1-q\sigma)[q^2 + \pi^2 B^2]}$, $B := \frac{1-q\sigma}{\sigma}$ and $E := q\sigma - (1 - q\sigma)\pi^2$.

It is easy to see that $\frac{\partial E}{\partial \sigma} \geq 0$. Moreover, it is also true that $\frac{\partial A}{\partial \sigma} \geq 0$. To see this, note that

$$\frac{\partial A}{\partial \sigma} = \frac{-q^2(1-q) [-q(q^2 + \pi^2 B^2) + 2\pi^2 B \frac{\partial B}{\partial \sigma} (1 - \sigma q)]}{((1 - q\sigma) [q^2 + \pi^2 B^2])^2} \geq 0$$

because $\frac{\partial B}{\partial \sigma} \leq 0$. However, the sign of E is ambiguous, with $E < 0$ for $\sigma < \frac{\pi^2}{q(1+\pi^2)} := \sigma^T$. We claim that the following two conditions are sufficient for uniqueness:

1. $\frac{\partial c_D(q)}{\partial \sigma} \leq 1$ for $\sigma \leq \sigma^T$;

$$2. \frac{\partial^2 c_D(q)}{\partial \sigma^2} \geq 0 \text{ for } \sigma \geq \sigma^T;$$

The argument is as follows: as $\text{RHS}(\sigma = 0) \geq 0 = \text{LHS}(\sigma = 0)$ and $\text{RHS}(\sigma = 1) < 1 = \text{LHS}(\sigma = 1)$, the fixed point is:

1. Only at $\sigma = 0$, as $\text{RHS}(\sigma = 0) = \text{RHS}(\sigma = \sigma^T)$ and below that in between. Moreover, there cannot be any additional crossing point above σ^T because the RHS would be coming from below, and, as it is convex, it cannot be that they cross and $\text{RHS}(\sigma = 1) < 1 = \text{LHS}(\sigma = 1)$.
2. If the solution is not at 0, the first time they cross it must be that the LHS comes from below. There are two sub-cases:
 - If the first crossing point is in $\sigma \leq \sigma^T$, then there cannot be others in the same interval as $\frac{\partial c_D(q)}{\partial \sigma} \leq 1$. Moreover, there cannot be any other crossing point above σ^T because the RHS would be coming from below, and, as it is convex, it cannot be that they cross and $\text{RHS}(\sigma = 1) < 1 = \text{LHS}(\sigma = 1)$.
 - If the first crossing point is above σ^T , it must be unique as a second solution would violate $\text{RHS}(\sigma = 1) < 1 = \text{LHS}(\sigma = 1)$.

We now prove that the two sufficient conditions outlined above apply to our model.

First, a sufficient condition for $\frac{\partial c_D(q)}{\partial \sigma} \leq 1$ for $\sigma \leq \sigma^T$ is $\frac{\partial E}{\partial \sigma} A \leq 1$. This implies $(1 + \pi^2)q^3(1 - q) \leq (1 - q\sigma)(q^2 + \pi^2 B^2)$. As the RHS is decreasing in σ , this condition must hold for the highest possible σ , i.e. for $\sigma = \sigma^T$. Substituting and simplifying, this requires $q(1 - q) \leq \frac{1}{\pi^2(1 + \pi^2)}$. The LHS is at most $\frac{1}{4}$ while the RHS is at least $\frac{1}{2}$, hence the condition is always satisfied.

Second, a sufficient condition for convexity of $c_D(q)$ for $\sigma \geq \sigma^T$ is $\frac{\partial^2 A}{\partial \sigma^2} \geq 0$. To show that it is always the case, note that

$$\frac{\partial^2 A}{\partial \sigma^2} = -q^2(1 - q) \frac{\frac{\partial^2 D}{\partial \sigma^2} D^2 - 2D \frac{\partial D}{\partial \sigma}^2}{D^4} \quad (2.9.C.24)$$

where $D = (1 - q\sigma)[q^2 + \pi^2 B^2]$, $\frac{\partial D}{\partial \sigma} = -q(q^2 + \pi^2 B^2) + \pi^2 2B \frac{\partial B}{\partial \sigma}(1 - \sigma q) < 0$ and $\frac{\partial^2 D}{\partial \sigma^2} = -q\pi^2 2B \frac{\partial B}{\partial \sigma} + 2\pi^2 \left[\left(\frac{\partial B}{\partial \sigma}^2 + \frac{\partial^2 B}{\partial \sigma^2} B \right) (1 - \sigma q) - qB \frac{\partial B}{\partial \sigma} \right] > 0$. A sufficient condition for (2.9.C.24) to be positive is $2 \frac{\partial D}{\partial \sigma}^2 \geq \frac{\partial^2 D}{\partial \sigma^2} D$.




Figure 2.C.8: Proof of Proposition 2.12: Proving LHS > RHS in (2.9.C.25). Orange plane: LHS–RHS, blue plane: $0 * x + 0 * \pi$ in the $\pi - x$ space.

By substitution, this implies

$$\begin{aligned}
 2 \left[-q(q^2\sigma^2 + \pi^2(1-q\sigma)^2) \frac{1}{\sigma^2} - 2\pi^2 \frac{(1-q\sigma)^2}{\sigma^3} \right]^2 &\geq \left[q\pi^2 2\sigma \frac{(1-q\sigma)}{\sigma^4} + 2\pi^2(1-q\sigma) \frac{1}{\sigma^4} + \frac{4(1-q\sigma)^2}{\sigma^4} \pi^2 + 2q\pi^2 \sigma \frac{(1-q\sigma)}{\sigma^4} \right] (1-q\sigma)(q^2 + \pi^2 B) \\
 \sigma^2 \left[-q(q^2\sigma^2 + \pi^2(1-q\sigma)^2) - 2\pi^2 \frac{(1-q\sigma)^2}{\sigma} \right]^2 &\geq 3\pi^2(1-q\sigma)^2(q^2\sigma^2 + \pi^2(1-q\sigma)^2) \\
 \sigma^2 q^2(q^2\sigma^2 + \pi^2(1-q\sigma)^2)^2 + 4\pi^4(1-q\sigma)^4 + 4\pi^2(1-q\sigma)^2 q\sigma(q^2\sigma^2 + \pi^2(1-q\sigma)^2) &\geq 3\pi^2(q^2\sigma^2 + \pi^2(1-q\sigma)^2)(1-q\sigma)^2
 \end{aligned}
 \tag{2.9.C.25}$$

where the second line follows by multiplication of both sides by σ^4 and the third by dividing both sides by 2 and working out explicitly the square on the LHS. Note that σ and q always appear together in the last line of (2.9.C.25). As a consequence, we can redefine $\sigma q := x$ and check whether the condition holds for $x \in [0, 1]$ and $\pi \in [0.5, 1]$. We prove this graphically using figure 2.C.8. It plots the difference between LHS and RHS of (2.9.C.25) for the whole range of possible values of x and π , showing that this difference is always positive. This completes the proof. \square

Monopoly-Duopoly comparison

Finally, we show that sufficient conditions for competition leading to more research than monopoly can be found in this set up as well.

Proposition 2.13. *There exists a nonempty interval of v values where $\sigma_D^*(q) > \sigma_M^*(q)$.*

Proof. A sufficient condition for the proposition to hold is that, for some values of v , $c_D(q) > c_M(q)$. Setting $v = 0$, and defining $B = \frac{1-q\sigma}{\sigma}$ note that:

$$\begin{aligned}
c_D(q) &= (\gamma(\emptyset) - \gamma(1))(q\sigma - (1 - q\sigma)\pi^2) + 1 - q \quad (2.9.C.26) \\
&= q\sigma \frac{q^2(1 - q)}{(1 - q\sigma)(q^2 + \pi^2 B^2)} - \pi^2(1 - q\sigma) \frac{q^2(1 - q)}{(1 - q\sigma)(q^2 + \pi^2 B^2)} + 1 - q \\
&= (1 - q) \left(\frac{q^3\sigma}{(1 - q\sigma)(q^2 + \pi^2 B^2)} - \frac{\pi^2 q^2}{q^2 + \pi^2 B^2} + 1 \right) \\
&= (1 - q) \left(\frac{q^2 + (1 - q\sigma)\pi^2 B^2}{(1 - q\sigma)(q^2 + \pi^2 B^2)} - \frac{\pi^2 q^2}{q^2 + \pi^2 B^2} \right) \\
&= \frac{1 - q}{1 - q\sigma} \left(\frac{q^2 + (1 - q\sigma)\pi^2 B^2}{(q^2 + \pi^2 B^2)} - \frac{\pi^2 q^2(1 - q\sigma)}{q^2 + \pi^2 B^2} \right)
\end{aligned}$$

where the first equality follows by substitution and the rest is a series of rearrangements. Note that, as $q \in (0, 1)$, neither $1 - q$ nor $1 - q\sigma$ are ever 0. Similarly, by substitution,

$$\begin{aligned}
c_M(q) &= (1 - \pi) \left(\frac{q(\sigma + (1 - \sigma)\pi)}{q\sigma + q(1 - \sigma)\pi + (1 - q)\pi} - \frac{q(1 - \sigma)}{1 - q\sigma} \right) \quad (2.9.C.27) \\
&= (1 - \pi)q \left(\frac{\sigma + (1 - \sigma)\pi}{q\sigma(1 - \pi) + \pi} - \frac{1 - \sigma}{1 - q\sigma} \right) \\
&= \frac{(1 - \pi)q\sigma(1 - q)}{(1 - q\sigma)(q\sigma(1 - \pi) + \pi)}
\end{aligned}$$

As a consequence, by comparison of (2.9.C.26) and (2.9.C.27), $c_D(q) > c_M(q)$ implies

$$\frac{q^2 + (1 - q\sigma)\pi^2(B^2 - q^2)}{(q^2 + \pi^2 B^2)} > \frac{(1 - \pi)q\sigma}{(q\sigma(1 - \pi) + \pi)} \quad (2.9.C.28)$$

Note that both LHS and RHS of (2.9.C.28) are decreasing in π . The case of RHS is straightforward. For the LHS, a sufficient condition is

$$(1 - \sigma q)2\pi(B^2 - q^2)(q^2 + \pi^2 B^2 - 2\pi B^2(q^2 + (1 - q\sigma)\pi^2(B^2 - q^2))) < 0$$

This simplifies to $-\sigma q 2\pi q^2 B^2 - (1 - \sigma q)2\pi q^4$ that is always negative.

As a consequence, a sufficient condition for $c_D(q) > c_M(q)$ is $\text{LHS}(\pi = 1) > \text{RHS}(\pi = 0.5)$. By substitution, this implies

$$\frac{q^2 + (1 - q\sigma)(B^2 - q^2)}{(q^2 + B^2)} > \frac{q\sigma}{1 + q\sigma}$$

After few simplifications and substituting back the value of B , we obtain

$$\sigma^2 q^2 \frac{2q\sigma - 1}{\sigma^2} + \frac{(1 - q\sigma)^3}{\sigma^2} > 0$$

A sufficient condition for this to hold is

$$1 - 3q\sigma + 2q^2\sigma^2 + q^3\sigma^3 > 0$$

Noticing that $q\sigma$ is bounded between 0 and 1, the condition is always satisfied and this completes the proof. □

D Monopoly with public signal

In this appendix we assume that the audience learns the actual timing with positive probability z in monopoly. This helps us establish that additional learning in our benchmark duopoly model happens not only because the timing is revealed with some probability but also because the audience uses additional information from outlets matching the state. In this set up, the outlet does not know whether the audience has learned the timing or not when taking its decision. Then, the condition for doing research is

$$z\gamma_M(2) + (1 - z)\gamma_M(\emptyset) - c \geq z\gamma_M(1) + (1 - z)(\pi\gamma_M(\emptyset) + (1 - \pi)\gamma_M(1)) \quad (2.9.D.29)$$

Note that $\gamma_M(2) = \gamma(2) = 1$ and $\gamma_M(1) = \gamma(1) = \frac{1-\sigma}{2-\sigma}$. However,

$$\gamma_M(\emptyset) = \frac{\sigma + (1 - \sigma)\pi}{\sigma + (1 - \sigma)\pi + \pi} \neq \gamma(\emptyset) = \frac{\sigma^2 + (1 - \sigma)(2 - \sigma)\pi^2}{\sigma^2 + (2 - \sigma)^2\pi^2}$$

because in duopoly the audience can learn also from the other player getting the state wrong. Hence, it is confused only if both outlets publish simultaneously and they both get the state right.

For comparison, we can write the duopoly condition for $v = 0$ a bit differently. Define χ the probability that the opponent behaves in a way that reveals the timing to the reader. Note that χ is “artificial” because it is the probability that j does not research when player i does (i.e. $\frac{1}{2}(1 - \sigma^j) + \frac{1}{2}$ on the LHS) and vice-versa (i.e. $\frac{1}{2}\sigma$ on the RHS). In such cases, the action of player j is fully revealing of the timing, irrespective of the endorsement. The duopoly condition for research is then

$$\chi\gamma(II) + (1 - \chi)\gamma(\emptyset) - c \geq \chi\gamma(I) + (1 - \chi)(\pi^2\gamma(\emptyset) + (1 - \pi^2)\gamma(I)) \quad (2.9.D.30)$$

Comparing (2.9.D.29) and (2.9.D.30) reveals that they are similar, but not identical. Even if we set $z = \chi$, the difference in $\gamma(\emptyset)$ and in the π^2 term of the RHS is still there. Hence, our result is not just due to the fact that the publication timing of the opponent reveals information about the timing of the other player. The content of the endorsements plays a role as well.

Chapter 3

Optimal Innovation Time-off

3.1 Introduction

Employee-driven innovations have become critical for the growth of big organizations in recent decades. Employees are usually more capable of identifying new ideas and problems with current workstreams and finding their solutions. First-hand experience working with the technology and specific expertise gives them a distinct advantage in the innovation process.¹ For instance, almost half of Google's products have been employee-driven innovations. Paul Buchheit, credited with developing *GMail* in Google, initially sought to refine the emailing experience after identifying that it was difficult to search within specific email inboxes. He later developed *AdSense* technology to recommend advertisements based on email searches, which now makes Google around \$10 billion each year in revenue.²

Recognizing the importance of employee-driven innovations, many firms have introduced policies of innovation time-off for their employees. For instance, Google permits its employees to take 20% time-off from their regular work to work on a project of their choice. Similarly, LinkedIn runs a yearly time-off program called the InCubator to support its employees' ideas.

However, there are apparent tradeoffs for organizations in providing employees with time off to work on their creative endeavors. On the one hand, employees would need to take time off from regular tasks to work on their projects that may lead to innovations and higher future profits for the firm. But on the other hand, providing the time off induces the agent to work on potentially unfruitful avenues, which is costly to the organizations as it diverts resources away from present tasks. In the presence of this tradeoff, we ask how a firm should

¹This is especially true for tech firms in Silicon Valley that spend millions of dollars in their recruitment processes to select the best candidates who are of the highest ability.

²Another example of successful employee-driven innovation is Flickr. It originally a side project for the team at Ludicorp whose sole product at the time was a web-based game called Game Neverending. See Tate (2012) for more examples and anecdotes.

optimally delegate authority to its employees to work on their creative projects.

(Currently) We build a simple two-period delegation model to answer this question. In our model, a principal (she) finances an agent's (he) work. The agent either works on a regular task, or he may request to work on his creative task. The regular task brings a normalized zero net benefit to either player. The creative task involves the agent working on his idea, which has the potential to generate a breakthrough and a higher payoff of 1 for both players. However, the type of idea that the agent possesses in any period is his private information and may be either high or low. A high-type idea has a higher ex-ante potential to succeed and generate a breakthrough than a low-type one. While the agent would like to gain authority to work on his idea independent of its type, the principal only finds it profitable to finance a high idea. In each period, the agent may draw a new idea if he does not already have a high idea.

The principal designs an optimal delegation policy which is contingent on past observed outcomes. One may imagine several plans that the principal may use. Fixed probability of granting authority across periods, declining probability, increasing probability, or a combination of the three depending on the previous outcome are a few examples. Our main result shows that the optimal mechanism resembles a time-based screening contract in which the agent may choose to seek authority today or tomorrow. If the agent seeks authority today, then he faces a potential future punishment for lack of performance. Specifically, the principal reduces the probability of granting authority in the next period when he does not produce a breakthrough in the first period. If the agent does not seek authority today, then the principal rewards him with more authority tomorrow *independent of his type*. Thus, in the optimal policy, the principal punishes the *persistent* good type but not the bad type. Either way, creativity only gets a limited opportunity.

There is some anecdotal evidence to support this observation. Google managers are known to clamp down on those employees that take the 20% time off but do not produce breakthroughs sufficiently quickly. In a Quora blog post³, a Google employee writes:

“Unofficially, 20% projects are no longer encouraged. They led to many problems because it took a great deal of time away from an employee's primary team (without any measurable successes).”

While at the same time, some successful employees mention how it is important to define the objectives of their projects to be able to benefit from the time-off policy. Defining objectives allows the managers to measure success and reward (or punish) accordingly in their projects.

³See <https://www.quora.com/How-does-Google-company-T1-textquoterights-Google-Innovation-Time-Of-f-20-time-policy-work-in-practice>

The intuition for the above result is as follows. The principal needs to give sufficient incentives to a low-type agent to deter seeking authority in the first period. She achieves this by using the future probability of authority in each contingency, i.e., if he does not seek authority and if he does and fails to produce a breakthrough. Therefore, she rewards the agent for continuing work on the regular task *and* also punishes him for seeking authority and not providing a breakthrough. Consequently, only the high-type agent gets penalized in the optimal incentive-compatible mechanism.

It is interesting to note that it is not possible to achieve the desired honesty of the low-type using just one of the tools. The principal not only rewards the low-type but also punishes the high-type in the optimal policy. In doing so, she faces an intertemporal tradeoff. Punishing the agent following failures not only relaxes the low-type's first-period incentive constraint but also makes the high-type's second-period constraint tighter. However, this hurts the principal as a high potential idea is persistent and taking away authority from such an agent reduces her expected payoff. The principal, therefore, tries to minimize the adverse effect on the high-type while at the same time maintaining the low-type's incentive to report honestly. How she uses the two tools optimally to balance the two effects depends on the various parameters.

We show that somewhat counterintuitively, the principal may sometimes grant authority to the agent even when ex-ante expected benefit of doing so is lower than the cost. It will never happen if we restrict the mechanism to a single period. However, in two periods, the intertemporal tradeoff kicks in. In such situations, it is more costly not to give authority to high types than to fund the low-type's creativity. Thus, while the principal is forced to punish the high-type agent for failing, she does so to the minimum possible extent. In this sense, the first-period high-type agent subsidizes the low-type in the optimal mechanism.

However, if the ex-ante cost of financing creativity is too high, then there is a discontinuous fall in the probability of getting authority for both the types in the second period. In this case, the principal would rather completely take away authority from the high-type, and risk not getting a breakthrough, than grant authority to the low-type. Again, while the principal must offer some authority to the low-type, she does so to the minimum extent. In this sense, the first-period low-type taxes the first and the second-period high-type in the optimal mechanism.

Related literature. Starting with [Holmstrom \(1982\)](#) there is a vast literature on delegation of authority in economics. [Aghion and Tirole \(1997\)](#) were the first to formalize the collaborative role played by employees in an organization. They showed that employee initiative could be increased by formal delegation and reducing the level of managerial effort. In contrast, [Rantakari \(2012\)](#) showed that employee initiative might be increased by combining formal

authority (of the manager) and limited but positive involvement of the manager. This result was achieved by combining the Aghion-Tirole model with elements of costly monitoring. While ours is a model of employee initiative, we deal with the issue of dynamic delegation of authority. A number of papers have started exploring this issue recently in different contexts; Frankel (2016), Guo (2016), Datta (2017), Guo and Horner (2017), Li, Matouschek and Powell (2017) Lipnowski and Ramos (2018) are a few.

Among these, Guo (2016) and Datta (2017) deal with the delegation in environments where the principal learns about ex-ante private information of the agent owing to experimentation. While it is common to model innovation in an experimentation-type framework, we do not explicitly need to include experimentation and learning in our model. All we require is ex-ante uncertainty over when breakthrough occurs and better information of the agent. Moreover, we are interested in the question of where ideas come from and when are they developed further. Our concern is not how a given idea is developed.

Guo and Horner (2017) is the closest in this sense to our paper. In their setting, a principal commits to an allocation policy for a perishable good (much like delegating authority) when the agent's type is persistent. While the agent would like the good either way, the principal interested in maximizing efficiency would want to grant the good only if the agent is of high type. However, there are two critical points of departure for us. First, we have a one-sided persistence of type only. In our model, only the good-type is persistent. For us, an agent who has discovered a good idea would prefer to see it through than drawing another idea, which at best is the same as the current idea. Moreover, we require that there might not be an immediate conclusion of the creative task. Second, we allow for breakthroughs to be observable. As we discuss in our analysis, these two features together produce our fundamental intertemporal tradeoff and produce our main result.

Other papers such as Li et al. (2017), Rantakari (2017) and Lipnowski and Ramos (2018) do not assume commitment by the principal and are interested in the equilibrium allocation rules that arise in similar settings using perfect public equilibrium (PPE). In Li et al. (2017), the authors show how good early choices of subordinates are rewarded with later authority. Consequently, they make more selfish decisions that end up hurting organizations in the long run. Lipnowski and Ramos (2018) build a model in which a principal sequentially delegates project choice to an agent who can assess its quality but has lower standards of acceptance. Similar to Li et al. (2017), they show that in equilibrium, the principal incentivizes the current good selection of projects by allowing future bad choices.

A common theme between this literature and our paper is the idea of linking incentives

across periods or decisions. By controlling allocation or decision rights to some other (or future) units, the principal creates value to eliciting private information today. This feature appears in several different contexts, including in papers that look at relational contracts and optimal contracts.

Jackson and Sonnenschein (2007) prove that the limitation imposed by incentive constraints on attaining social efficiency can be reduced by increasing the number of copies of the decision problem and thereby linking them. In a similar spirit, Malenko (2018) develops a model of dynamic capital budgeting in which an agent with a desire to overinvest in projects privately sees the arrival and quality of investment opportunities. A principal allocates resources to the agent to be spent on these investment opportunities. The optimal mechanism involves the principal setting a dynamic spending account that gets replenished with time at a specific rate. Additionally, the mechanism outlines a threshold limit above which the agent can pass the project to the principal for auditing. Similarly, Möbius (2001) and Hauser and Hopenhayn (2008) build models of favor trading and show how the number of favors exchanged between players may be used to determine how many new favors can be exchanged.

Two related papers are Boleslavsky and Kelly (2014) and Casella (2005). In the context of environmental regulation, using a two-period model Boleslavsky and Kelly (2014) show how the regulators may vary the strength of regulations over time when the firm privately learns its compliance costs. However, again, neither do they have one-sided state persistence nor do they have a verifiable signal of the state (since the latter is not a concern in their setting). Similarly, Casella (2005) develops a mechanism of storable votes that allows an agent to gain more influence in future democratic decision-making by giving up the right to vote today when she expects future preferences to be strong.

The rest of the the paper is organized as follows: in Section 3.2 we present what our general model will look like. However, we focus currently on the simplified version of the model presented in Section 3.3. We discuss the model and present all the relevant results within the section itself.

3.2 The general model

We present here the general model that we expect to solve after solving the simple two-period model.

Players, Tasks and Types: An agent (he) is in an employment relationship with a principal (she) through time $t = 0, 1, 2, \dots$. The two players are denoted by $i \in \{A, P\}$ and share a common discount factor δ . While in the relationship, the agent could potentially work on one of the two tasks in any given period: regular or own. The regular task corresponds to

working on assigned projects. Conducting the regular task gives a payoff of $r_i > 0$ to each player with certainty.

While working on the regular task, the agent has the ability to come up with new ideas and create own tasks. The agent costlessly comes up with an own task in each period, which could either have a low or a high potential to succeed, denoted by $\theta \in \{l, h\}$. This is the type of the own task and is agent's private information. The prior probability of drawing a high potential own task in period t is p_t . If the agent works on own task of type θ , then the probability of success or breakthrough in any given period is θ with $0 < l < h < 1$ – a high potential task is more likely to succeed in any given period. Upon achieving a breakthrough in own task, each player gets a payoff of v_i from then on in perpetuity. We will assume that success in own task is sufficiently rewarding such that $r_i < \delta v_i$. If the agent fails, then both players get 0 in that period. We denote the outcome of conducting own task by $y \in \{S, F\}$, success or failure respectively i.e. either a breakthrough is observed or not.

The higher payoff to the agent from conducting own tasks pertain to the intrinsic motivation for working on own ideas. Moreover, upon achieving success in own tasks, the principal may reward the agent with promotions and more autonomy to further conduct own task which gives the agent v_A in perpetuity. The higher payoffs to the principal from achieving success in own task pertain to the benefits of getting an innovation.

We will assume that the principal can observe the outcome y of conducting any task and also the realized payoff. This implies that the principal can perfectly distinguish between the agent conducting regular vs. own task but cannot distinguish between low potential and high potential own task.

Note that since there is nothing better than a high potential idea, the agent stops drawing new ideas after getting a high potential idea. By drawing a new idea he risks losing his current high idea. But may continue drawing ideas if she has a low potential idea, because there is only a potential benefit of doing so. Thus the high type is fully persistent.

Conflict of Preferences: While both players agree on the need for innovation, there is a conflict on which type of ideas should be pursued. We will assume

$$\begin{aligned} hv_P &> r_P \text{ and } hv_A > r_A \\ \text{but } lv_P &< r_P \text{ and } lv_A > r_A \end{aligned}$$

This implies that while the agent always prefers to conduct own tasks, the principal would rather have him conduct the regular task when she has a low potential own task. In other words, the agent's intrinsic motivation to “scratch the itch” is sufficiently high independent of the type of

the task he has drawn. However, the principal would like the agent to do so if she has a high potential own task.

Policy and Payoffs: Given the realized θ_t and the policy (defined next), the agent sends a message $m_t \in \{l, h\}$. A policy (or mechanism) is a sequence of probabilities of granting authority to the agent to conduct own task for every period t denoted by a_t as a function of previous messages, outcomes and authority decisions.

$$a_t : \{l, h\}^{t+1} \times \{S, F\}^t \times \{0, 1\}^t \rightarrow \Delta(\{0, 1\})$$

Timing: The timing of the game is as follows:

1. Initially, the principal commits to a mechanism $a_t(\cdot)$ for $t = 1, 2, \dots$
2. In $t = 1$, the agent draws his type θ_1 and reports it to the principal. The principal chooses according to $a_1(\theta_1)$. Outcome y_1 is realized if $a_1 = 1$, otherwise the game moves to $t = 2$.
3. For all $t > 1$, the agent reports his θ_t to the principal and she chooses whether to grant authority using $a_t(\cdot)$.

3.3 Two-period case: model and analysis

We present here a modified two-period version of the model. Consider the following changes. An agent and a principal are in an employment relationship for two periods $t = 1, 2$. The principal bears a cost of $0 < c < 1$ per period for financing the agent's work. If the agent works on the regular task, then the principal breaks even and gets 0. The agent also has no intrinsic motivation to do the regular task and gets 0 from conducting it. On the other hand, both the agent and the principal get a benefit of 1 from success in own task.

However, there is a preference conflict between the agent and the principal captured by the following assumption on parameters: $h > c > l$ while $1 > h > l > 0$. This implies that while the principal would like the agent to only pursue high potential ideas, the agent would like to have the authority to conduct own task independent of its potential to succeed. Thus, the agent has lower standards. The relationship ends after getting a success in own task or at the end of the second period, whichever happens first. All the other parameters of the problem are as above.

The principal designs a delegation mechanism which defines for every t the probability of granting authority to the agent to work on his own task a_t . Using the revelation principle, we restrict attention to those mechanisms in which the agent reports his type and the principal

decides whether authority must be granted.

To begin with, let us look at the principal's first-best policy. The first-best is given by the situation in which the principal can perfectly observe the type of the agent's own task in each task.

Proposition 3.1. *The principal chooses $a_t = 1$ whenever $\theta_t = h$, and $a_t = 0$ otherwise in the first-best policy.*

It is easy to see why the above is the case. As $h > c > l$, the expected payoff of the principal is always positive whenever the agent has a high-potential own task. Alternately, it is negative when the agent has a low-potential own task.⁴

Next, let us look at two extreme policies that the principal may adopt under the second-best. The first is to never grant any authority and the second is to always grant authority.

Lemma 3.1. *If the principal never grants authority, then her expected payoff is zero. If the principal always grants authority, her expected payoff is*

$$p_1 h + (1 - p_1) l - c + \delta [p_1 (1 - h)(h - c) + (1 - p_1)(1 - l)(p_2 h + (1 - p_2) l - c)].$$

Going forward, we need to determine if the principal can do better than these two mechanisms, and if so, when. Below, we will call all other non-extreme mechanisms as *interior mechanisms*.

Note first that in $t = 1$, there is no outcome to condition a_1 on. Therefore, the principal chooses a_1 as a function of m_1 which could be either h or l . With some abuse of notation, we denote this by $a_1(h)$ and $a_1(l)$ respectively. In $t = 2$, however, the principal may condition her decision on three variables – the present report, the past report and the past outcome, which could either be a failure or that no authority was granted.⁵ We denote the principal's period 2 decision by $a_2(m_2; m_1, y_1)$ where $m_t \in \{h, l\}$ and $y_1 \in \{F, \emptyset\}$ is the outcome of the previous period which could be either that the authority was granted but the agent failed to produce a breakthrough (F) or that he wasn't granted authority (\emptyset). Observe that there are many variables that one must determine in the optimal mechanism. In the following paragraphs, we show how the problem can be reduced significantly.

Lemma 3.2. *The optimal interior mechanism does not condition second period delegation decision on the reports in that period.*

⁴Note that the agent does not need to be incentivized to participate in the mechanism. He always gets 0 by being in the employment relationship with the principal. It can also be assumed to be the value of the outside option.

⁵If authority was granted in the first period and it generated a success, then there is no further need for making authority decisions as the relationship comes to an end.

This follows almost immediately from the fact that the game ends after the second period. The principal is not able to incentivize truth-telling by offering any future rewards or punishments in the second period. Incentive compatibility requires that she offers the same authority to the agent independent of her second period type. Thus, period 2 decisions can only be based on the past decisions and outcomes, and not on the present reports. This reduces our problem to determining the following probabilities in the second period: $a_2(l, F)$, $a_2(l, \emptyset)$, $a_2(h, F)$ and $a_2(h, \emptyset)$. These probabilities of granting authority reflect the principal's decision for when 1) the agent reported low but was given authority and failed, 2) the agent reported low and was not given authority, 3) the agent reported high and was given authority but failed, and 4) the agent reported high but was not given authority in period 1.

We can now write the principal's authority design problem as follows:

$$\begin{aligned}
& \underset{a_t(m, y)}{\text{maximize}} && p_1(h - c)[a_1(h) + a_1(h)(1 - h)\delta a_2(h, F) + (1 - a_1(h))\delta a_2(h, \emptyset)] \\
& && + (1 - p_1)[a_1(l)(l - c) + \delta(p_2 h + (1 - p_2)l - c)(a_1(l)(1 - l)a_2(l, F) + (1 - a_1(l))a_2(l, \emptyset))] \\
& \text{subject to} && a_1(h) + \delta[a_1(h)(1 - h)a_2(h, F) + (1 - a_1(h))a_2(h, \emptyset)] \\
& && \geq a_1(l) + \delta[a_1(l)(1 - h)a_2(l, F) + (1 - a_1(l))a_2(l, \emptyset)], && (IC)_h \\
& && a_1(l)l + \delta(p_2 h + (1 - p_2)l)[a_1(l)(1 - l)a_2(l, F) + (1 - a_1(l))a_2(l, \emptyset)] \\
& && \geq a_1(h)l + \delta(p_2 h + (1 - p_2)l)[a_1(h)(1 - l)a_2(h, F) + (1 - a_1(h))a_2(h, \emptyset)], && (IC)_l
\end{aligned}$$

where $a_t(m, y) = \{a_1(h), a_1(l), a_2(l, F), a_2(l, \emptyset), a_2(h, F), a_2(h, \emptyset)\}$ is the set of all delegation probabilities that the principal chooses.

Lemma 3.3. *The low type's incentive constraint $(IC)_l$ binds in the optimal interior mechanism.*

The intuition is straightforward. Fix any incentive compatible mechanism in which both the ICs are slack, specifically $(IC)_l$ is non-binding. Now, if the principal increases $a_1(h)$ while satisfying $(IC)_l$ then the ex-ante expected profits increase on account of a higher probability of breakthrough for the high type without inducing any deviation by the low type. This means that optimality requires $(IC)_l$ to bind.

Lemma 3.4. *Let $p_2 > l$ and $\underline{\delta} < \delta < \bar{\delta}$ where $\underline{\delta} = \frac{l}{p_2 h + (1 - p_2)l}$ and $\bar{\delta} = \frac{p_2}{p_2 h + (1 - p_2)l}$. It is always optimal for the principal to grant authority following a high report and not following a low report in period 1, i.e. $a_1(h) = 1$ and $a_1(l) = 0$.*

The above lemma shows that the first-best optimal delegation policy is the same as the second-best optimal delegation policy in the first period. In the presence of future discounting, the principal prefers not to distort incentives in the first period. The first sufficiency condition on the discount factor, $\delta < \bar{\delta}$, ensures that future is not too valuable, which might lead to pushing

all authority decisions to the second period. The second sufficiency condition $\delta > \underline{\delta}$ ensures that the future is sufficiently valuable to the agent so that the low type does not want to deviate today and inefficiently seek authority. (All proofs are presented in Appendix A)

Using Lemma 3.4, it is easy to reduce the principal's interior maximization problem to

$$\begin{aligned}
& \underset{a_2(h, F), a_2(l, \emptyset)}{\text{maximize}} && p_1(h - c)[1 + (1 - h)\delta a_2(h, F)] \\
& && + (1 - p_1)\delta(p_2 h + (1 - p_2)l - c)a_2(l, \emptyset) \\
& \text{subject to} && 1 + \delta(1 - h)a_2(h, F) > \delta a_2(l, \emptyset) \quad (\text{IC})_h \\
& && a_2(l, \emptyset) = \frac{l}{\delta(p_2 h + (1 - p_2)l)} + (1 - l)a_2(h, F) \quad (\text{IC})_l
\end{aligned}$$

where we only need to determine two variables $a_2(h, F), a_2(l, \emptyset)$. Moreover, note that $(\text{IC})_h$ is slack at the optimum. Therefore, the optimal interior mechanism is the solution to above linear program as outlined by the following proposition.

Proposition 3.2. *Let $p_2 > l$. For $\underline{\delta} < \delta < \bar{\delta}$ the optimal interior mechanism is given by the following delegation probabilities:*

- $a_1^*(h) = 1$ and $a_1^*(l) = 0$,
- if $c < \bar{c}$ where $\bar{c} = \frac{p_1 h(1-h) + (1-p_1)(1-l)(p_2 h + (1-p_2)l)}{1 - (p_1 h + (1-p_1)l)}$,

$$a_2^*(h, F) = \frac{1}{1-l} \left(1 - \frac{l}{\delta(p_2 h + (1-p_2)l)} \right) < 1 \text{ and } a_2^*(l, \emptyset) = 1,$$

- if $c = \bar{c}$, then any $a_2^*(h, F)$ and $a_2^*(l, \emptyset)$ that satisfies

$$a_2^*(l, \emptyset) = \frac{l}{\delta(p_2 h + (1-p_2)l)} + (1-l)a_2^*(h, F),$$

- if $c > \bar{c}$,

$$a_2^*(h, F) = 0 \text{ and } a_2^*(l, \emptyset) = \frac{l}{\delta(p_2 h + (1-p_2)l)} < 1, \text{ and}$$

- any $a_2^*(h, \emptyset) \in [0, 1]$ and $a_2^*(l, F) \in [0, 1]$.

The above proposition characterizes the optimal mechanism for different cost ranges. Observe how the probabilities of delegation switch for the different types between first and the second period. An agent who is of high type in the first period must suffer some punishment in the second period for failing to achieve a breakthrough. This is despite the high type being fully persistent. At the same time, a low type must be rewarded for waiting to get a high idea

when he presently doesn't have one. This implies that the first period low type must get a higher probability of delegation in the second period independent his future type. Thus, the principal buys period 1 efficiency from the low type agent by inefficiently offering rewards to the low type and punishments to the high type in period 2.⁶

We note here that the flipping of delegation probabilities for the two types is unique to our environment and arises from the two differentiating features of our model – observable outcomes and persistent high type. These two features together generate a tradeoff for the principal. On the one hand, the possibility of conditioning future delegation on the past performance relaxes the low type's constraint by making deviations less likely today. But on the other hand, since the high type is fully persistent, it hurts his incentives tomorrow. This in turn reduces the expected payoffs of the principal as she must knowingly take away authority from someone who brings in positive expected profits. Such an intertemporal tradeoff is resolved by the principal by switching delegation probabilities across periods for the two types in the optimal mechanism.

Why must the principal use both the tools? This is best understood by looking at $(IC)_l$ when $a_1(h) = 1$ and $a_1(l) = 0$;

$$a_2(l, \emptyset) = \frac{l}{\delta(p_2 h + (1 - p_2)l)} + (1 - l)a_2(h, F)$$

It is easy to see that two delegation probabilities work complementarily. An increase (decrease) in one is accompanied by an increase (decrease) in the other in order to maintain incentive compatibility of the low type. However, given the incentive compatibility of the high type, the principal would like to increase $a_2(h, F)$ to the maximum possible extent *independent* of $a_2(l, \emptyset)$. This might not always be doable or in the interest of the principal. The two cost cases depict this.

To begin with, note that given the linearity of the principal's objective function and the IC constraints, the solution must be an extreme point. Now, when the cost of delegating is low enough, i.e. lower than \bar{c} , then the principal prefers to still offer authority to the high type with a positive probability and simultaneously with probability 1 to the low type. However, when the cost is high, then financing the high type for a second period while simultaneously delegating to an ex-ante unknown type becomes too costly. In this situation, the principal would rather give no authority to the high type and restrict the delegation probability to the minimum for the low type. We explain this point further using the corollary below. For this purpose, let $p_1 = p_2 = p$.

⁶Since it never happens that a low type is offered authority and the high type is not offered authority in the first period, $a_2(h, \emptyset)$ and $a_2(l, F)$ do not matter in anyone's decision-making, and we may assign any probabilities to them.

Corollary 3.1. *Let $p_1 = p_2 = p$. For $\tilde{c} \leq c < \bar{c}$ where $\tilde{c} = ph + (1 - p)l$ is the ex-ante expected benefit of delegation, the optimal mechanism involves*

$$a_2(h, F) = \frac{1}{1 - l} \left(1 - \frac{l}{\delta(ph + (1 - p)l)} \right) < 1 \text{ and } a_2(l, \emptyset) = 1.$$

The above corollary shows that even if the ex-ante expected benefit of granting authority is lower than the cost of doing so, the principal optimally delegates authority with probability 1 to such an agent some times. Observe that this would not be the case if this was a single period mechanism. In a single period, the principal would never offer authority to an agent who is too costly to fund. The difference arises because of how the principal resolves the aforementioned intertemporal tradeoff. See Figure 3.1 in relation to the explanation below.

When $c < \tilde{c}$, i.e. the cost of delegating is lower than the the expected benefit of delegating, then the preferences of the principal and period 1 low type agent are more aligned. On an average, she expects to gain by granting authority to period 1 low type in period 2. At the same time, she would like to continue granting authority in period 2 to period 1 high type. This is reflected in her isoprofit lines. The slope of the isoprofit line in $(a_2(h, F), a_2(l, \emptyset))$ space is negative. To remain at a given profit level, the principal must increase $a_2(l, \emptyset)$ following a decrease in $a_2(h, F)$. This leads to an easy resolution of the tradeoff in the optimal mechanism; the principal happily gives authority to the low type while holding the punishment to the high type to a minimum.

But when $c > \tilde{c}$, i.e the cost of delegating authority to period 1 low type in period 2 is higher than the expected benefit, then the preferences are less aligned. This means that the principal would like to take away authority from the low type i.e. reduce $a_2(l, \emptyset)$. At the same time, the desire to not punish the high type agent in period 2 remain. Thus, to maintain a given profit level a reduction in $a_2(h, F)$ must be accompanied by a reduction in $a_2(l, \emptyset)$. The intertemporal tradeoff now has a bite - the principal must decide between using $a_2(h, F)$ more intensely or less intensely.⁷ By choosing a lower $a_2(h, F)$, she can minimize the punishment to the high type but it comes at the cost of inefficiently giving authority to the low type in period 2. Alternately, she can reduce the inefficiency by increasing the high type's punishment tomorrow, i.e. a higher $a_2(h, F)$. How this tradeoff is resolved now depends on a second cost threshold \bar{c} .

\bar{c} reflects the cost threshold at which the principal is indifferent between how the tradeoff is resolved. This implies that for $\tilde{c} \leq c < \bar{c}$ the principal still prefers to grant authority to period 1 low type in period 2. This is so because in this range of cost parameters the principal's cost of

⁷Note that not using $a_2(h, F)$, i.e. setting $a_2(h, F) = 1$ is not an option for the principal because $(IC)_l$ must be respected at all times. Setting $a_2(h, F) = 1$ breaks $(IC)_l$ as it requires an $a_2(l, \emptyset) > 1$ to match it.

taking away authority from a high type is higher than the cost of inefficiently granting authority to the low type. Thus, the high type agent ‘subsidizes’ the low type in the optimal mechanism. But when $c > \bar{c}$, the cost of granting authority to the low type is higher than the cost of taking away authority from the high type. Thus, the principal prefers to set $a_2(h, F) = 0$ and minimize $a_2(l, \emptyset)$ in the process. The low type, in this case, ‘taxes’ the high type.

Finally, the optimal mechanism is evaluated by comparing the principal’s payoffs from the optimal interior mechanism with the extreme mechanisms outlined in Lemma 3.1.

Proposition 3.3. *Let $l < p_2 < \frac{l}{1-(h-l)}$. For $\underline{\delta} < \delta < \bar{\delta}$ the optimal mechanism is given by cost thresholds \underline{c} , \bar{c} and $\bar{\bar{c}}$ where $\underline{c} < \bar{c} < \bar{\bar{c}}$ such that*

- *if $c < \underline{c}$, the principal always grants authority in the optimal mechanism;*
- *if $\underline{c} \leq c < \bar{c}$, the principal grants authority as follows in the optimal mechanism*

$$a_1^*(h) = 1, a_1^*(l) = 0, a_2^*(h, F) = \frac{1}{1-l} \left(1 - \frac{l}{\delta(p_2 h + (1-p_2)l)} \right) < 1 \text{ and } a_2^*(l, \emptyset) = 1,$$

and any $a_2^(h, \emptyset) \in [0, 1]$ and $a_2^*(l, F) \in [0, 1]$;*

- *if $\bar{c} \leq c < \bar{\bar{c}}$, the principal grants authority as follows in the optimal mechanism*

$$a_1^*(h) = 1, a_1^*(l) = 0, a_2^*(h, F) = 0 \text{ and } a_2^*(l, \emptyset) = \frac{l}{\delta(p_2 h + (1-p_2)l)} < 1,$$

and any $a_2^(h, \emptyset) \in [0, 1]$ and $a_2^*(l, F) \in [0, 1]$;*

- *if $c \geq \bar{\bar{c}}$, the principal never grants authority.*

The above proposition outlines the overall optimal mechanism for different cost ranges. As expected, we show that when the cost of granting authority is too low or too high, the principal prefers to implement the extreme mechanism instead of the interior one. When the cost is low, the principal always grants authority; when the cost is too large, the principal never grants authority. In between, the principal prefers the optimal interior mechanism.

3.4 Conclusion

We showed using a simple two-period model how organizations may limit the creativity of their employees. Organizations end up being inefficiently harsh on employees who are capable of achieving breakthroughs in own ideas to limit the behavior of those employees who do not have good ideas. Such mechanisms are, therefore, not likely to be successful in promoting creativity among the employees. Our model suggests that organizations must look for alternative ways of fostering creativity.

A few points are in order about the model that we built in this paper. First, there are some obvious issues in extending our model to more than two periods. Since we need that the agent's outcome is observable, when we expand our model to include many periods, the problem blows up immediately. Second, this is not a model of experimentation even though we are attempting to model innovation. By assuming that θ is perfectly known to the agent, we are essentially killing any learning, a standard of experimentation models. One may introduce the tradeoff between the regular task and own task as one between exploitation and exploration à la [Manso \(2011\)](#) in our current two-period framework. However, it is not clear what new insights we obtain from doing so. Third, there are no monetary transfers in our model. This is so because we are interested in employee-driven innovation where the employer-employee are already in a relationship and the employee wants to conduct innovation driven by intrinsic motivation. Finally, we can extend our model in meaningful ways to include moral hazard or multiple types. Our model, therefore, shows further scope of research.

3.5 Figures




Figure 3.1: Optimal delegation mechanism for different c ranges

3.6 Appendix

A Proofs from the main text

Proof of Lemma 3.4

Proof. We know from Lemma 3.3 that $(IC)_l$ binds in the optimal mechanism. Rearranging $(IC)_l$ and substituting in the principal's expected profit function we get:

$$\begin{aligned} & p_1(h - c) [a_1(h) + a_1(h)(1 - h)\delta a_2(h, F) + (1 - a_1(h))\delta a_2(h, \emptyset)] \\ & + (1 - p_1) [a_1(h)l + \delta(p_2h + (1 - p_2)l)(a_1(h)(1 - l)a_2(h, F) + (1 - a_1(h))a_2(h, \emptyset))] \\ & - (1 - p)c [a_1(l) + a_1(l)(1 - l)\delta a_2(l, F) + (1 - a_1(l))\delta a_2(l, \emptyset)] \end{aligned}$$

It is now easy to verify that

$$\frac{\partial \pi_P}{\partial a_1(l)} = -(1 - p_1)c [1 + \delta(1 - l)a_2(l, F) - \delta a_2(l, \emptyset)] < 0$$

since $1 > \delta a_2(l, \emptyset)$. Now, given the linearity of the profit function in $a_1(l)$, it is immediate that the principal should set $a_1^*(l) = 0$ in the optimal mechanism.

Now, substitute $a_1^*(l) = 0$ and $(1 - a_1(h))a_2(h, \emptyset) = a_2(l, \emptyset) - a_1(h)\frac{l}{\delta(p_2h + (1 - p_2)l)} - a_1(h)(1 - l)a_2(h, F)$ from $(IC)_l$ in the original profit function and maximize with respect to $a_2(h)$:

$$\frac{\partial \pi_P}{\partial a_1(h)} = p_1(h - c) \left[1 - \delta(h - l)a_2(h, F) - \frac{l}{p_2h + (1 - p_2)l} \right]$$

which is always positive for $\delta < \bar{\delta} = \frac{p_2}{p_2h + (1 - p_2)l}$. Thus, $a_1^*(h) = 1$.

Moreover, since $a_2(h, F), a_2(l, \emptyset)$ are numbers between 0 and 1, from $(IC)_l$ we get the second sufficiency condition

$$\frac{l}{\delta(p_2h + (1 - p_2)l)} < 1 \implies \delta > \underline{\delta} = \frac{l}{p_2h + (1 - p_2)l}.$$

□

Proof of Proposition 3.2

Proof. Substitute $a_1^*(h) = 1$ and $a_1^*(l) = 0$ in the IC conditions and the principal's profit function. This reduces the principal's optimization problem to

$$\begin{aligned}
 & \underset{a_2(h, F), a_2(l, \emptyset)}{\text{maximize}} && p_1(h - c)[1 + (1 - h)\delta a_2(h, F)] \\
 & && + (1 - p_1)\delta(p_2h + (1 - p_2)l - c)a_2(l, \emptyset) \\
 \text{subject to} &&& 1 + \delta(1 - h)a_2(h, F) > \delta a_2(l, \emptyset) && (\text{IC})_h \\
 &&& a_2(l, \emptyset) = \frac{l}{\delta(p_2h + (1 - p_2)l)} + (1 - l)a_2(h, F) && (\text{IC})_l
 \end{aligned}$$

First, note that $(\text{IC})_h$ must be slack at the optimum. Second, by the fact that we have already assumed $\delta < \bar{\delta}$, it is easy to verify that $a_2(h, F) < 1$.

Now, given the linearity of the profit function and the IC constraints it is easy to derive the optimal mechanism by comparing the slopes of the isoprofit lines and $(\text{IC})_l$. In the $(a_2(h, F), a_2(l, \emptyset))$ space, the slope of the isoprofit line is given by

$$\frac{-p_1(h - c)(1 - h)}{(1 - p_1)(p_2h + (1 - p_2)l - c)}$$

and the slope of $(\text{IC})_l$ is $1 - l$.

- If $p_2h + (1 - p_2)l > c$, then the slope of isoprofit line is negative. Moreover, a higher profit is a shift of the isoprofit line to the right. This implies that in the optimal mechanism $a_2^*(l, \emptyset) = 1$ and $a_2^*(h, F) = \frac{1}{1-l} \left(1 - \frac{l}{\delta(p_2h + (1-p_2)l)}\right)$.
- If $p_2h + (1 - p_2)l < c$, then the slope of the isoprofit line is positive and a higher profit is a shift of the line to the right and down. Two cases are possible depending on the comparison of slopes

$$\begin{aligned}
 & \frac{p_1(h - c)(1 - h)}{(1 - p_1)(c - p_2h + (1 - p_2)l)} \leq 1 - l \\
 \implies c & \leq \bar{c} = \frac{p_1h(1 - h) + (1 - p_1)(1 - l)(p_2h + (1 - p_2)l)}{1 - (p_1h + (1 - p_1)l)}.
 \end{aligned}$$

- If $c < \bar{c}$, then the isoprofit line is steeper than $(\text{IC})_l$. The optimal mechanism as in the case above, i.e. $a_2^*(l, \emptyset) = 1$ and $a_2^*(h, F) = \frac{1}{1-l} \left(1 - \frac{l}{\delta(p_2h + (1-p_2)l)}\right)$.
- If $c > \bar{c}$, then the isoprofit line is flatter than $(\text{IC})_l$. The optimal mechanism now is $a_2^*(l, \emptyset) = \frac{l}{\delta(p_2h + (1-p_2)l)}$ and $a_2^*(h, F) = 0$.

This completes the proof of Proposition 3.2 and Corollary 3.1. \square

Proof of Proposition 3.3

Proof. Begin by rewriting the value to the principal of always granting authority to the agent. From Lemma 3.1 we know that it is equal to

$$p_1 h + (1 - p_1)l - c + \delta[p_1(1 - h)(h - c) + (1 - p_1)(1 - l)(p_2 h + (1 - p_2)l - c)].$$

Using the fact that $\bar{c} = \frac{p_1 h(1 - h) + (1 - p_1)(1 - l)(p_2 h + (1 - p_2)l)}{1 - (p_1 h + (1 - p_1)l)}$, we can simplify the above as

$$p_1 h + (1 - p_1)l - c + \delta[1 - (p_1 h + (1 - p_1)l)](\bar{c} - c). \quad (3.6.A.1)$$

Consider first the situation of $c = \bar{c}$. From Proposition 3.2, we know that

$$a_1^*(h) = 1, a_1^*(l) = 0, a_2^*(h, F) = 0 \text{ and } a_2^*(l, \emptyset) = \frac{l}{\delta(p_2 h + (1 - p_2)l)} < 1,$$

and any $a_2^*(h, \emptyset) \in [0, 1]$ and $a_2^*(l, F) \in [0, 1]$ is an optimal interior mechanism. At this optimal mechanism, the value to the principal is

$$p_1(h - c) + (1 - p_1)l \left(1 - \frac{c}{p_2 h + (1 - p_2)l}\right). \quad (3.6.A.2)$$

Comparing the value to the principal from (3.6.A.1) with (3.6.A.2), it is easy to see that the latter is better. Therefore, at $c = \bar{c}$, the principal prefers the optimal interior mechanism.

It also now immediately follows, that for the case $c > \bar{c}$, either the interior mechanism is optimal or the extreme where no one is granted authority is optimal. The extreme mechanism where the agent always gets authority can not be optimal anymore since such a mechanism performs worse than the extreme one outlined above. To check whether it is optimal to grant no authority we need to check if

$$p_1(h - c) + (1 - p_1)l \left(1 - \frac{c}{p_2 h + (1 - p_2)l}\right) > 0.$$

Simplifying the above gives the condition,

$$c < \frac{(p_1 h + (1 - p_1)l)(p_2 h + (1 - p_2)l)}{p_1 p_2 h + (1 - p_1 p_2)l} := \bar{\bar{c}}.$$

Finally, we now need to determine what happens when $c < \bar{c}$. Observe that at $c = 0$, the extreme mechanism where the agent always gets authority is better than the interior mechanism.

Reminding ourselves that the optimal mechanism from Proposition 3.2 is given by

$$a_1^*(h) = 1, a_1^*(l) = 0, a_2^*(h, F) = \frac{1}{1-l} \left(1 - \frac{l}{\delta(p_2h + (1-p_2)l)} \right) < 1 \text{ and } a_2^*(l, \emptyset) = 1,$$

and any $a_2^*(h, \emptyset) \in [0, 1]$ and $a_2^*(l, F) \in [0, 1]$. The value to the principal is

$$p_1h \left(1 - \frac{1-h}{1-l} \frac{l}{\tilde{c}_2} \right) + \delta \left(p_1h \frac{1-h}{1-l} + (1-p_1)\tilde{c}_2 \right) \quad (3.6.A.3)$$

where $\tilde{c}_2 = p_2h + (1-p_2)l$. Also, at $c = 0$, the value to the principal from the extreme mechanism is given by

$$p_1h + (1-p_1)l + \delta(p_1h(1-h) + (1-p_1)\tilde{c}_2(1-l)). \quad (3.6.A.4)$$

Comparing (3.6.A.4) with (3.6.A.3) we get that the extreme mechanism gives a higher value if

$$1 - p_1 + p_1h \frac{1-h}{1-l} \frac{1}{\tilde{c}_2} > \delta \left((1-p_1)\tilde{c}_2 + p_1h \frac{1-h}{1-l} \right),$$

which is always verified because $\tilde{c}_2 < 1$ and under the assumption $p_2 < \frac{l}{1-(h-l)}$, we have that $\delta < 1$.

Combining the above with the fact that the interior optimal mechanism is better for the principal at $c = \bar{c}$ and that the value is a linear (decreasing) function of c , it must be that there exists a $\underline{c} \in (0, \bar{c})$ where

- for $c < \underline{c}$, the principal prefers the extreme mechanism where the agent always gets the authority, and
- for $c \geq \underline{c}$, the principal prefers the optimal interior mechanism.

This completes the proof. □

Bibliography

- Aghamolla, Cyrus**, “Essay on Analyst Herding.” PhD dissertation, Columbia University 2016.
- Aghion, Philippe and Jean Tirole**, “Formal and real authority in organizations,” *Journal of political economy*, 1997, 105 (1), 1–29.
- Ali, Omer**, “Disclosure in Multistage Projects,” *Working paper*, 2017.
- Allcott, Hunt and Matthew Gentzkow**, “Social media and fake news in the 2016 election,” *Journal of Economic Perspectives*, 2017, 31 (2), 211–36.
- Andreottola, Giovanni and Antoni Italo De Moragas**, “Fake News, Media Competition and Political Accountability,” *Working paper* 2017.
- Angelucci, Charles and Julia Cagé**, “Newspapers in times of low advertising revenues,” *American Economic Journal: Microeconomics*, 2019, 11 (3), 319–64.
- Armstrong, Mark**, “Public service broadcasting,” *Fiscal Studies*, 2005, 26 (3), 281–299.
- Aumann, Robert J and Sergiu Hart**, “Long cheap talk,” *Econometrica*, November 2003, 71 (6), 1619–1660.
- Bandura, Albert**, “Self-efficacy: toward a unifying theory of behavioral change,” *Psychological review*, 1977, 84 (2), 191.
- Barrera, Oscar, Sergei Guriev, Emeric Henry, and Ekaterina Zhuravskaya**, *Facts, alternative facts and fact checking in times of post-truth politics*, Vol. 11, Centre for Economic Policy Research, 2017.
- Bénabou, Roland and Jean Tirole**, “Self-confidence and personal motivation,” *The Quarterly Journal of Economics*, 2002, 117 (3), 871–915.
- Besley, Timothy and Andrea Prat**, “Handcuffs for the grabbing hand? Media capture and government accountability,” *American Economic Review*, 2006, 96 (3), 720–736.
- Bizzotto, Jacopo, Jesper Rüdiger, and Adrien Vigier**, “Dynamic Persuasion with Outside Information,” *Working paper*, 2018.
- Bobtcheff, Catherine, Jérôme Bolte, and Thomas Mariotti**, “Researcher’s dilemma,” *The Review of Economic Studies*, 2016, 84 (3), 969–1014.
- Boleslavsky, Raphael and David L Kelly**, “Dynamic regulation design without payments: The importance of timing,” *Journal of Public Economics*, 2014, 120, 169–180.
- **and Tracy R Lewis**, “Evolving influence: Mitigating extreme conflicts of interest in advisory relationships,” *Games and Economic Behavior*, 2016, 98, 110–134.
- Cairncross, Frances**, “The Cairncross Review. A Sustainable Future for Journalism,” Technical Report, Department for Digital, Culture, Media & Sport, HM Government of UK February 2019.
- Carson, Andrea**, *Investigative Journalism, Democracy and the Digital Age*, Routledge, 2019.
- Casella, Alessandra**, “Storable votes,” *Games and Economic Behavior*, 2005, 51 (2), 391–419.
- Chan, Ying**, “Journalism and Digital Times: Between Wider Reach and Sloppy Reporting,” in

- M. Dragomir and M. Thompson, eds., *Digital Journalism. Making News, Breaking News*, Open Society Foundations, 2014.
- Che, Yeon-Koo and Johannes Horner**, “Optimal design for social learning,” *Working Paper*, 2015.
- Chen, Heng and Wing Suen**, “Competition for attention in the news media market,” Working paper 2016.
- Datta, Bikramaditya**, “Delegation and Learning,” *Working paper*, 2017.
- Ellman, Matthew and Fabrizio Germano**, “What do the papers sell? A model of advertising and media bias,” *The Economic Journal*, 2009, 119 (537), 680–704.
- Ely, Jeffrey C**, “Beeps,” *The American Economic Review*, 2017, 107 (1), 31–53.
- and **Martin Szydlowski**, “Moving the Goalposts,” *Working paper*, 2017.
- Forges, Françoise and Frédéric Koessler**, “Long persuasion games,” *Journal of Economic Theory*, 2008, 143 (1), 1–35.
- Frankel, Alexander**, “Discounted quotas,” *Journal of Economic Theory*, 2016, 166, 396–444.
- Fuchs, William**, “Contracting with repeated moral hazard and private evaluations,” *The American Economic Review*, 2007, 97 (4), 1432–1448.
- Fudenberg, Drew and Jean Tirole**, “Preemption and rent equalization in the adoption of new technology,” *The Review of Economic Studies*, 1985, 52 (3), 383–401.
- , **Richard Gilbert, Joseph Stiglitz, and Jean Tirole**, “Preemption, leapfrogging and competition in patent races,” *European Economic Review*, 1983, 22 (1), 3–31.
- Gentzkow, Matthew**, “Trading dollars for dollars: The price of attention online and offline,” *American Economic Review: Papers & Proceedings*, 2014, 104 (5), 481–88.
- and **Jesse M Shapiro**, “Media bias and reputation,” *Journal of Political Economy*, 2006, 114 (2), 280–316.
- and —, “Competition and Truth in the Market for News,” *Journal of Economic Perspectives*, 2008, 22 (2), 133–154.
- Golosov, Mikhail, Vasiliki Skreta, Aleh Tsyvinski, and Andrea Wilson**, “Dynamic strategic information transmission,” *Journal of Economic Theory*, 2014, 151, 304–341.
- Goltsman, Maria, Johannes Hörner, Gregory Pavlov, and Francesco Squintani**, “Mediation, arbitration and negotiation,” *Journal of Economic Theory*, 2009, 144 (4), 1397–1420.
- Gratton, Gabriele, Richard Holden, and Anton Kolotilin**, “When to drop a bombshell,” *The Review of Economic Studies*, 2017, 85 (4), 2139–2172.
- Guo, Yingni**, “Dynamic delegation of experimentation,” *American Economic Review*, 2016, 106 (8), 1969–2008.
- and **Johannes Horner**, “Dynamic mechanisms without money,” 2017.
- Guttman, Ilan**, “The timing of analysts’ earnings forecasts,” *The Accounting Review*, 2010, 85 (2), 513–545.
- , **Ilan Kremer, and Andrzej Skrzypacz**, “Not only what but also when: A theory of dynamic voluntary disclosure,” *American Economic Review*, 2014, 104 (8), 2400–2420.
- Hafer, Catherine, Dimitri Landa, and Patrick Le Bihan**, “Who Gets the Credit? News-Gathering Competition and Political Accountability,” Working paper 2018.
- , —, and —, “Breaking News: Media Competition and Political Polarization,” Working paper 2019.

- Hauser, Christine and Hugo Hopenhayn**, “Trading favors: Optimal exchange and forgiveness,” *Working paper*, 2008.
- Henry, Emeric and Marco Ottaviani**, “Research and the Approval Process: The Organization of Persuasion,” *American Economic Review*, March 2019, 109 (3), 911–55.
- Hirshleifer, David, Angie Low, and Siew Hong Teoh**, “Are overconfident CEOs better innovators?,” *The Journal of Finance*, 2012, 67 (4), 1457–1498.
- Holmstrom, Bengt**, *On the theory of delegation*, Northwestern University, 1982.
- Honryo, Takakazu**, “Dynamic persuasion,” *Journal of Economic Theory*, 2018, 178, 36 – 58.
- Hopenhayn, Hugo A and Francesco Squintani**, “Preemption games with private information,” *The Review of Economic Studies*, 2011, 78 (2), 667–692.
- and — , “Patent rights and innovation disclosure,” *The Review of Economic Studies*, 2015, 83 (1), 199–230.
- Hörner, Johannes and Andrzej Skrzypacz**, “Learning, experimentation and information design,” *Working Paper*, 2016.
- Jackson, Matthew O. and Hugo F. Sonnenschein**, “Overcoming Incentive Constraints by Linking Decisions,” *Econometrica*, 2007, 75 (1), 241–257.
- Kamenica, Emir and Matthew Gentzkow**, “Bayesian persuasion,” *The American Economic Review*, 2011, 101 (6), 2590–2615.
- Knobel, Beth**, *The Watchdog Still Barks: How Accountability Reporting Evolved for the Digital Age*, Fordham Univ Press, 2018.
- Koellinger, Philipp, Maria Minniti, and Christian Schade**, ““I think I can, I think I can”: Overconfidence and entrepreneurial behavior,” *Journal of economic psychology*, 2007, 28 (4), 502–527.
- Kremer, Ilan, Yishay Mansour, and Motty Perry**, “Implementing the “Wisdom of the Crowd”,” *Journal of Political Economy*, 2014, 122 (5), 988–1012.
- Krishna, Vijay and John Morgan**, “The art of conversation: eliciting information from experts through multi-stage communication,” *Journal of Economic theory*, 2004, 117 (2), 147–179.
- Lee, Angela Min-Chia**, “How fast is too fast? Examining the impact of speed-driven journalism on news production and audience reception.” PhD dissertation, University of Texas at Austin 2014.
- Li, Jin, Niko Matouschek, and Michael Powell**, “Power dynamics in organizations,” *American Economic Journal: Microeconomics*, 2017, 9 (1), 217–41.
- Lipnowski, Elliot and Joao Ramos**, “Repeated delegation,” *Working paper*, 2018.
- Lizzeri, Alessandro, Margaret A Meyer, and Nicola Persico**, *The incentive effects of interim performance evaluations*, University of Pennsylvania, Center for Analytic Research in Economics and the Social Sciences, 2002.
- Malenko, Andrey**, “Optimal dynamic capital budgeting,” *Working paper*, 2018.
- Manso, Gustavo**, “Motivating innovation,” *The Journal of Finance*, 2011, 66 (5), 1823–1860.
- Möbius, Markus**, “Trading favors,” *Working paper*, 2001.
- Morris, Stephen**, “Political Correctness,” *Journal of Political Economy*, 2001, 109 (2), 231–265.
- Orlov, Dmitry**, “Optimal Design of Internal Disclosure,” *Working paper*, 2013.
- , **Andrzej Skrzypacz, and Pavel Zryumov**, “Persuading the principal to wait,” *Working paper*, 2018.

- Perego, Yacopo and Sevgi Yuksel**, “Media Competition and Social Disagreement,” Working paper 2018.
- Prat, Andrea and David Strömberg**, “The political economy of mass media,” *Advances in economics and econometrics*, 2013, 2, 135.
- Rantakari, Heikki**, “Employee initiative and managerial control,” *American Economic Journal: Microeconomics*, 2012, 4 (3), 171–211.
- , “Relational Influence,” *Working paper*, 2017.
- Renault, Jérôme, Eilon Solan, and Nicolas Vieille**, “Dynamic sender–receiver games,” *Journal of Economic Theory*, 2013, 148 (2), 502–534.
- , —, and —, “Optimal dynamic information provision,” *Games and Economic Behavior*, 2017, 104, 329–349.
- Rochet, Jean-Charles and Jean Tirole**, “Platform competition in two-sided markets,” *Journal of the European Economic Association*, 2003, 1 (4), 990–1029.
- and —, “Two-sided markets: a progress report,” *The RAND Journal of Economics*, 2006, 37 (3), 645–667.
- Rosenberg, Howard and Charles S Feldman**, *No time to think: The menace of media speed and the 24-hour news cycle*, A&C Black, 2008.
- Shapiro, Jesse M.**, “Special interests and the media: Theory and an application to climate change,” *Journal of Public Economics*, 2016, 144, 91–108.
- Smolin, Alex**, “Dynamic Evaluation Design,” *Working Paper*, 2017.
- Sobbrio, Francesco**, “Citizen-editors’ endogenous information acquisition and news accuracy,” *Journal of Public Economics*, 2014, 113, 43–53.
- Starbird, Kate, Dharma Dailey, Owla Mohamed, Gina Lee, and Emma S Spiro**, “Engage early, correct more: How journalists participate in false rumors online during crisis events,” in “Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems” ACM 2018, p. 105.
- Strömberg, David**, “Media and politics,” *Annual Review of Economics*, 2015, 7 (1), 173–205.
- Tate, Ryan**, *The 20% Doctrine: How tinkering, goofing off, and breaking the rules at work drive success in business*, HarperCollins, 2012.
- Vaccari, Federico**, “Direct Competition in the News Market,” Working paper 2018.