Essays in Microeconomic Theory

by

Erica Meixiazi Liu

Department of Economics
Duke University

Date: _____

Approved:

_____
Philipp Sadowski, Co-Supervisor

_____
Curtis Taylor, Co-Supervisor

_____
James Anton

_____
David McAdams

Dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Economics
in the Graduate School of
Duke University
2022

ABSTRACT

Essays in Microeconomic Theory

by

Erica Meixiazi Liu

Department of Economics
Duke University

Date: _____
Approved:

_____
Philipp Sadowski, Co-Supervisor

_____
Curtis Taylor, Co-Supervisor

_____
James Anton

_____
David McAdams

An abstract of a dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Economics
in the Graduate School of
Duke University

2022

Abstract

This dissertation consists of three main chapters - chapter 2, 3, and 4. These are different research problems studying the use of information. In chapter 2, we study a mechanism design problem where the Principal hires two agents to inspect a product, the quality of which is uncertain. The main research question we ask is the comparison of the inspection protocols. We found the optimality of sequential inspection among the protocols considered. We then extend the analysis to allow agents to differ and try to understand which is the better to order the agents when using the sequential protocol. In chapter 3, we study an information design problem where a present biased agent could commit to an information choice to help herself to save more that she would have. We provide a full characterization of the optimal information choice for a risk averse present biased agent. After that, as an effort to further understand the interaction between risk aversion and present bias, we introduce the EZKP framework and find a counterexample where risk aversion itself has no impact on the form of the information choice in a special case where the elasticity of the intertemporal substitution is fixed. In chapter 4, we study an information design in games problem where a designer chooses information for the agents to induce joint effort. We provide two examples illustrating the relative strength of two prominent constructions of the optimal information structure in the literature.

Dedication

To Kim-Sau Chung, my first advisor, for showing me the joy and beauty of intellectual pursuits.

Contents

List of Figures

Acknowledgements

The submission of this dissertation marks the end of my journey in the academic world. All the excitements and the frustrations, the intriguing and the disappointing, and everything I have learned along the way will be treasured memories together with the small discoveries documented here. I thank everyone that has been part of this journey, my advisors, family, friends, and people I came across and got to know.

Among all the people who have helped me getting to this destination point, I am most indebted to my advisors, Curtis Taylor and Philipp Sadowski. They were the anchor of my academic life and kept me going with their continuous support when I was most devastated by the disruptions. This dissertation would not look the same if not for the extensive guidance and feedback that I have received from them. Besides my main advisors, I am also grateful for having David McAdams and James Anton on my dissertation committee and Adam Rosen as our DGS.

In the darkest of times, it was quotes, poems, songs, life stories of fighters from every walk of life, and my parents' unconditional faith in me that gave me strength to stay in the arena and finish the fight. Today is the day I say bye to academia, and I am no longer one of the "most privileged people on earth" as Ariel Rubinstein said in his article for "lost" graduate students. But I still "owe something in return", not for those privileges, but for the pure joys that I have experienced while wandering in the academic world. I would be grateful if I could have a chance to return this favor to our society and make my small contribution someday.

Chapter 1

Introduction

This dissertation is a collection of results studying the use of information in three different scenarios. The first project, which is a joint work with Curt, is a study on the contracting of information acquisition where decision making and information acquisition are conducted by different parties. This captures many real-life situations where an individual seeks expert options on the unknown quality of a product. The focus of this study is on the comparison of different protocols. When the principal could hire multiple experts to inspect the product, it is important to know whether she should approach the experts at the same time or one at a time. To simplify the analysis, we first consider agents to be symmetric. Once we've established the optimality of sequential inspection as a protocol, a next question to ask is: what if the agents are not the same? They could differ in how skilled they are as well as their specialization. Given the same effort choice, some expert is better at detecting the defect than others. In other cases, for example, we could have human resource managers who are better at telling whether an applicant could be a good fit for the firm and others better at telling whether they could be a misfit. Once the experts become different, which one to approach first becomes an interesting question to ask. While contracting for information acquisition is well studied, the specific question on the optimal ordering the agents in a sequence has not received much attention. We find examples where it is better to first consult the more skilled expert and the expert who specializes in looking for evidence contrary to our prior belief.

The second project is a study on the possible use of informational tools to help alleviate the dynamic inconsistency problem of a present biased agent. A present biased agent discounts the near future too much and finds only in the future that she has not saved enough. This is an issue many of us could relate to when we end up in a situation where the money is not enough for life after retirement. Traditional intervention from the government includes programs like 401(k) where working people are encouraged to save more. While these "hard" tools could be effective, we are living in an environment where there is an abundance of information. If there is some way that we could help encourage more saving through the guidance of what news people should read, then this could be a better approach as news is freely available one click away. We consider a behavioral agent with present bias as well as risk aversion which is also the common risk attitude and try to see if we could solve this problem through the costless choice of information. As we will see in our analysis, it turns out not knowing exactly what the economic condition will be tomorrow sometimes could help with our savings. After finding that optimal information choice, we are going to dig deeper and see how this present bias and the risk attitude interact in our decision on what information to acquire.

The third project is a study on the constructions of the optimal information structure for information design in incomplete information coordination games. This class of problems feature an information designer choosing the informational environment for a set of strategic players to induce a desirable outcome. The general properties of the optimal information structure have been limited in their guidance for solving this type of problem in applied settings where constructions are often ad hoc. As part of the effort towards finding an easy-to-follow con-

struction procedure for applied words, we take two prominent constructions from the literature, the *pooling* and *Hoshino (2017)*, and highlight their relative strength in two examples. We find that the prior distribution over states plays a role in determining their performance.

Chapter 2

Quality Control

2.1 Introduction

Consulting experts is an important way to gather information for better decision making in many situations. A car manufacturer hires inspectors to check if there is significant safe issue before introducing a new model to the market to avoid future recalls. A patient visits several doctors for different diagnosis in case early signs of a serious disease might be overlooked. An employer puts an applicant through several rounds of interviews to see if he's a good fit hoping not to waste time and effort training and onboarding someone who would not stay. A journal editor invites multiple referees to help evaluate a paper so that it's less likely for them to publish questionable results that would tarnish the reputation of the journal.

Before reaching out to the experts with a contract, our decision maker needs to decide first what is the best way to approach these experts. Should we approach the experts one by one or simultaneously? Having the experts inspect a product sequentially helps avoid the case where more than one of them obtaining the same conclusive message on the existence of some defect, but also makes inducing the later inspectors more costly as they become more pessimistic about discovering a defect since earlier people have failed. The comparison of the different protocols is therefore unclear.

Instead of solving for the optimal mechanism, in this paper, we take an indirect approach and compare a selection of mechanisms, or protocols. We model the

inspection as costly information acquisition where agents choose a costly effort to acquire a signal about the underlying state. Specifically, the information acquisition technology is such that the probability of discovering the defect is determined by the effort. We start out with the benchmark model where the two agents share the same information acquisition technology and show that sequential inspection dominates simultaneous inspection.

In general, the difficulty in obtaining analytical solution for principal-expert problems is well recognized in the literature as stated in Carroll (2015). While we managed to get an analytical result for convex effort costs in the benchmark, thanks to the simplicity of our information acquisition technology, this becomes impossible when we extend to settings with asymmetric agents. However, the superiority of the sequential inspection protocol for symmetric agents leads us to a following question about how we should order the agents once their technology starts to differ. We compromised the generality of the analysis and imposed functional form restrictions among others and found some insights that could potentially help guide future studies seeking a complete understanding of the problem. In particular, we found examples where it is better to ask the more skilled inspector first if they both look for evidence of the defect and that it is better to ask the expert who specialize in looking evidence confirming the state that is deemed to be unlikely when they differ in which state they could confirm.

Compared to information acquisition, contracting for experimentation has received more attention as studies by Bergemann and Hege (1998, 2005), Manso (2011), Hörner and Samuelson (2013), and Halac et al. (2016), among others. The closest paper to ours is Gerardi and Maestri (2012) where they also study contracting for information acquisition. However, there is some essential differences from our model as their effort choice is binary, the cost of effort is constant, and

the realization of signal is only observable to the agent which implies that truthful revelation is no longer guaranteed.

In the following, we first introduce the model and then analyze the benchmark case with symmetric agents by comparing three different inspection protocols: simultaneous, sequential, and random sequencing. After that, we extend the analysis to allow the agents to differ first in their probability getting the confirming message with the same effort and then in the direction of their information acquisition.

2.2 Model

The Principal has a product of unknown quality $\omega \in \Omega = \{G, B\}$ and needs to choose an action $\alpha \in \mathcal{A} = \{t, nt\}$. Her preference over the actions depends on the underlying state in the following way:

$$v(\alpha, \omega) = \begin{cases} -L, & \text{if } \alpha = nt, \omega = B \\ 0, & \text{otherwise} \end{cases}. \tag{2.1}$$

Given the utility function, the Principal is indifferent between the actions *treat* and *not treat* in the *good* state while strictly prefers *treat* to avoid loss in the bad state.

Before taking an action, the Principal hires two agents A_1 and A_2 to inspect the product. The agents share the same information acquisition technology exerting effort $a_i \in [0, 1]$ with increasing and strictly convex cost function $c : [0, 1] \to \mathbb{R}_+$ that is sufficiently smooth with $c \in C^3$ to look for an informative signal $s \in S = \{B, \varnothing\}$ to learn about the bad state. We first consider the case where agents can only look for bad news with the conditional probabilities of getting each signal depending

on their effort choices only in the bad state:

$$Prob(s = B|\omega = B) = \theta_i a_i, \tag{2.2}$$

$$Prob(s = \varnothing|\omega = B) = 1 - \theta_i a_i, \tag{2.3}$$

$$Prob(s = \varnothing|\omega = G) = 1, \tag{2.4}$$

with $\theta_i \in (0, 1)$ which is an index for the agent's capability. The agents' payoffs depend on the transfer $w \in \mathbb{R}_+$ they receive from the Principal as well as the cost of effort incurred in the following way:

$$u(w_i, a_i) = w_i - c(a_i). \tag{2.5}$$

The common prior belief on the state being bad is denoted by p. Both the Principal and the agents are risk neutral, and the agents are protected by limited liability. In the following analysis, we consider two cases where the agents are symmetric ($\theta_1 = \theta_2 = 1$) and asymmetric ($\theta_2 < \theta_1 = 1$). We'll impose an additional assumption on the cost function throughout which we'll see later is sufficient[1] for the existence of interior solutions:

Assumption 2.2.1. *The marginal cost of effort is zero but increases significantly when there is zero effort, and it increases at a strictly increasing speed everywhere. Specifically,*

$$c'(0) = 0, c''(0) > pL, and\ c'''(\cdot) > 0\ \forall a \in [0, 1]. \tag{2.6}$$

[1] A weaker set of sufficient conditions that would work as well is: $c'(0) < pL$, and $2c''(a) + ac'''(a) > pL,\ \forall a$.

2.3 Benchmark: Symmetric Agents

The Principal asks the agents to inspect following one of the three protocols: (1) both agents inspect simultaneously; (2) A_1 inspects first and if the signal is $\varnothing$ A_2 inspects; (3) A_1 inspects first with probability $\pi \in [0, 1]$ and A_2 inspects first with probability $1 - \pi$. In this section, we analyze the three scenarios and compare the efficiency in terms of the Agents' induced efforts and Principal's welfare when possible.

Simultaneous Inspection: Given the information acquisition technology, there are a total of four possible outcomes:

$$(s_1, s_2) \in \Big\{(B, B), (B, \varnothing), (\varnothing, B), (\varnothing, \varnothing)\Big\}. \tag{2.7}$$

Since the agents can get an empty message with zero cost and the Principal uses transfers to induce effort, it is optimal for the Principal to give zero transfers for empty messages. For the remaining cases, let w_i be the payment to the agent A_i who discovers the defect when the other agent does not, and $\bar{w}$ be the payment to each agent when both discover the defect. The Principal's problem is:

$$\min_{a_1, a_2, w_1, w_2, \bar{w}} p[L(1 - a_1)(1 - a_2) + a_1(1 - a_2)w_1 + a_2(1 - a_1)w_2 + a_1 a_2 2\bar{w}] \tag{2.8}$$

$$s.t.\ \ a_i = \arg\max p[a_i(1 - a_j)w_i + a_i a_j \bar{w}] - c(a_i), \forall i. \tag{2.9}$$

By solving the Agent's problem and using the first order condition to replace the incentive constraints:

$$c'(a_i) = p[(1 - a_j)w_i + a_j \bar{w}], \tag{2.10}$$

the Principal's problem can then be rewritten as:

$$\min_{a_1,a_2} K_{SM} = pL(1-a_1)(1-a_2) + a_1c'(a_1) + a_2c'(a_2), \tag{2.11}$$

where

$$a_ic'(a_i) = p[a_i(1-a_j)w_i + a_ia_j\bar{w}]. \tag{2.12}$$

Comparing with the first-best problem allows us to see how the second-best efforts compare to the first-best:

Lemma 2.3.1. $a^* < a^{FB}$.

Proof. The optimal efforts are pinned down by the first order condition of the Principal's problem above. Given the symmetry of the agents, the optimal effort in this second-best case is characterized by:

$$pL(1-a^*) = c'(a^*) + a^*c''(a^*). \tag{2.13}$$

An interior solution $a^* \in (0,1)$ exists since the second order condition of this transformed unconstrained optimization problem is satisfied as the Hessian $D^2K_{SM}(a_1^*, a_2^*)$ is a positive definite symmetric matrix:

$$\left.\frac{\partial^2}{\partial a_i^2}K_{SM}\right|_{a_i=a^*} = 2c''(a^*) + a^*c'''(a^*) > 0, \tag{2.14}$$

$$\begin{vmatrix} 2c''(a^*) + a_1c'''(a^*) & pL \\ pL & 2c''(a_2) + a^*c'''(a^*) \end{vmatrix} > c''(0)^2 - (pL)^2 > 0,$$

$$\text{where } \left.\frac{\partial^2}{\partial a_ia_j}K_{SM}\right|_{a_i=a_j=a^*} = pL, \tag{2.15}$$

since $c''(0) > pL$ and $c'''(\cdot) > 0$ by Assumption 2.2.1, and neither boundary $a = 0$ or 1 is optimal:

$$\frac{\partial}{\partial a_i} K_{SM}\bigg|_{a_i=0} = -pL + c'(0) < 0, \text{ by Assumption 2.2.1,}$$
$$\text{and } \frac{\partial}{\partial a_i} K_{SM}\bigg|_{a_i=1} = c'(1) + c''(1) > 0. \quad (2.16)$$

For the first-best solution, the Principal's problem is:

$$\min_{a_1,a_2} \; pL(1-a_1)(1-a_2) + c(a_1) + c(a_2). \quad (2.17)$$

Similarly, the optimal effort in the first-base case is characterized by[2]:

$$pL(1-a^{FB}) = c'(a^{FB}). \quad (2.18)$$

By subtracting these two characterizations, we get:

$$pL(a^{FB} - a^*) = c'(a^*) + a^* c''(a^*) - c'(a^{FB}). \quad (2.19)$$

Suppose $a^* \geqslant a^{FB}$, strict convexity of the cost function implies that the right hand side is strictly positive while the left hand side is weakly negative which is impossible. Therefore, we have $a^* < a^{FB}$. ■

Sequential Inspection: There are three possible outcomes in this case:

$$\Big\{(B), (\varnothing, B), (\varnothing, \varnothing)\Big\}. \quad (2.20)$$

[2]The proof of the existence of interior solution is similar to the second-best case and therefore omitted.

Similar to our analysis above, it is optimal to set the transfers equal to zero for empty messages. The Principal's problem is:

$$\min_{a_1,a_2,w_1,w_2,\bar{w}} p[L(1-a_1)(1-a_2) + a_1 w_1 + a_2(1-a_1)w_2] \tag{2.21}$$

$$s.t.\ \ a_2 = \arg\max \frac{p(1-a_1)}{1-pa_1} a_2 w_2 - c(a_2), \tag{2.22}$$

$$a_1 = \arg\max pa_1 w_1 - c(a_1), \tag{2.23}$$

where $\frac{p(1-a_1)}{1-pa_1}$ is A_2's updated belief on the state being bad since he moves second and therefore it must be that A_1 has received an empty message given the protocol. Using the characterizations of optimal efforts solved from the agents' problems:

$$w_2 = \frac{(1-pa_1)}{p(1-a_1)} c'(a_2), \text{ and } w_1 = \frac{1}{p} c'(a_1), \tag{2.24}$$

we can rewrite the Principal's problem as:

$$\min_{a_1,a_2} K_{SQ} = pL(1-a_1)(1-a_2) + a_1 c'(a_1) + (1-pa_1)a_2 c'(a_2). \tag{2.25}$$

A natural question to ask here is whether it's optimal for the Principal to induce a higher effort from A_1 or A_2, which cannot be easily answered by checking the marginal effects of the effort choices from the objective function above. As we can see from this objective function, the Principal's expected loss comes from two sources: the expected loss when the defect has not been detected which is $pL(1-a_1)(1-a_2)$, and the expected transfer to A_1 and A_2 which is $a_1 c'(a_1) + (1-pa_1)a_2 c'(a_2)$. If we have equal optimal efforts, a marginal increase in a_1 and a_2 will induce the same amount of reduction in expected loss but different amount of increase in the expected transfers. Different from the analysis of the simultaneous inspection

case where the marginal increase in the expected transfer takes the same form $\frac{d}{da_i}a_i c'(a_i)$ for both $i = 1, 2$, here the dynamic nature of this protocol implies that A_1's effort choice affects A_2's updated belief and hence effort choice. This means that, compared to the simultaneous case, the marginal increase in expected transfer due to a marginal increase in A_1's effort, $\frac{d}{da_1}a_1 c'(a_1) - pa_2 c'(a_2)$, is smaller by a constant; while contribution from marginally increasing A_2's effort, $(1 - pa_1)\frac{d}{da_2}a_2 c'(a_2)$, is smaller by a fraction. The reason why we said that the comparison of a_1^{**} and a_2^{**} cannot be easily answered comes from the difficulty in telling which has a more significant impact at the optimum.

It turns out that Assumption 2.2.1 is a sufficient condition for inducing a lower effort from the first inspector A_1 to be optimal.

Proposition 2.3.2. $a_1^{**} < a_2^{**}$.

Proof. The optimal efforts are characterized by the first order conditions of the Principal's problem:

$$-pL(1 - a_2^{**}) + c'(a_1^{**}) + a_1^{**}c''(a_1^{**}) - pa_2^{**}c'(a_2^{**}) = 0, \tag{2.26}$$

$$-pL(1 - a_1^{**}) + (1 - pa_1^{**})(c'(a_2^{**}) + a_2^{**}c''(a_2^{**})) = 0. \tag{2.27}$$

Subtracting and rearranging these two equations gives us:

$$pL(a_1^{**} - a_2^{**}) + a_2^{**}c''(a_2^{**}) - a_1^{**}c''(a_1^{**}) = pc'(a_2^{**})(a_1^{**} - a_2^{**}) + pa_1^{**}a_2^{**}c''(a_2^{**}) + c'(a_1^{**}) - c'(a_2^{**}), \tag{2.28}$$

where the left hand side can be rewritten as:

$$\int_{a_2^{**}}^{a_1^{**}} [pL - (c''(a) + ac'''(a))]\, da. \tag{2.29}$$

Suppose $a_1^{**} \geqslant a_2^{**}$. Note first that the cost function is strictly increasing and strictly convex implies that the right hand side is strictly positive. Assumption 2.2.1 implies that:

$$pL - (c''(a) + ac'''(a)) < pL - c''(0) < 0, \text{ since } c'''(\cdot) > 0 \text{ and } c''(0) > pL. \tag{2.30}$$

This implies that the left hand side is strictly negative, which is impossible. Therefore, $a_1^{**} < a_2^{**}$. ■

Compare the Principal's objective function to that in the simultaneous case, we can see that the only difference is the extra term $-pa_1a_2c'(a_2)$. Suppose that the agents are exerting the same level of effort with $a_1 = a_2 = a$. To see why this could not be optimal, let us consider local deviations by marginally increasing a_1 and a_2. The convexity of the cost function implies that it is better for the Principal to marginally increasing a_2 instead of a_1 as:

$$\left.\frac{\partial K_{SQ}}{\partial a_1}\right|_{a_1=a_2=a} > \left.\frac{\partial K_{SQ}}{\partial a_2}\right|_{a_1=a_2=a}, \tag{2.31}$$

where

$$\left.\frac{\partial K_{SQ}}{\partial a_1}\right|_{a_1=a_2=a} = -pL(1-a) + c'(a) + ac''(a) - pac'(a), \tag{2.32}$$

$$\left.\frac{\partial K_{SQ}}{\partial a_2}\right|_{a_1=a_2=a} = -pL(1-a) + c'(a) + ac''(a) - pac'(a) - pa^2c''(a). \tag{2.33}$$

This local argument explains the direction of adjustments for effort choices a_1 and a_2 which are substitutes when the Principal adopts the sequential inspection protocol.

Lemma 2.3.3. *Sequential inspection is a better choice of protocol for the Principal than simultaneous inspection.*

Proof. Compare the Principals' expected payoff functions:

$$K_{SM} - K_{SQ} = pa_1a_2c'(a_2) > 0, \tag{2.34}$$

since the cost function is strictly increasing. K_{SQ} everywhere below K_{SM} then implies that $\min K_{SQ} < \min K_{SM}$. ■

Essentially, this is a result of the dynamic nature of this protocol which allows the Principal to save on the transfer to A_2 when A_1 has already learned that the state is bad.

Random Sequencing: Given that the order to inspect becomes random, there are six possible outcomes in this case:

$$\Big\{(s_1 = B), (s_1 = \varnothing, s_2 = B), (s_1 = \varnothing, s_2 = \varnothing), (s_2 = B), \\ (s_2 = \varnothing, s_1 = B), (s_2 = \varnothing, s_1 = \varnothing)\Big\}. \tag{2.35}$$

Since the agent's identity is irrelevant for the value of the signal to the Principal and it is similarly optimal to give zero transfers for empty messages, the indices of the transfers w_1 and w_2 rewarding successful detection of the defect now correspond to the realized order in the sequence, instead of the identity index, of the agent. This randomization distorts the incentives of all parties by changing the Principal's probability of getting the bad news as well as the agents' probability of getting w_1 or w_2 for a successful detection for any pair of effort choice (a_1, a_2). For either agent, when they're asked to inspect, it could be either of the two cases: (1) they are the first one in line; (2) they are the second inspector but the first one couldn't confirm if there is a defect. The randomization probability is common knowledge, which

implies that the agents' beliefs on them being the first one asked to inspect are respectively:

$$\sigma_1 = \frac{1}{1+\lambda(1-pa_2)}, \text{ and } \sigma_2 = \frac{1}{1+\frac{1}{\lambda}(1-pa_1)}, \text{ where } \lambda = \frac{1-\pi}{\pi} \in [0,\infty]. \quad (2.36)$$

The order matters since it determines their belief on the underlying state. Now we're ready to proceed with the analysis by first understanding how good this random sequencing protocol is for any randomization probability π, and then figuring out if there is a best randomization probability if we allow it to be a choice variable of the Principal.

For any given π, the Principal's problem is:

$$\begin{aligned} \min_{a_1,a_2,w_1,w_2} p\Big\{ & L(1-a_1)(1-a_2) + w_1[\pi a_1 + (1-\pi)a_2] + w_2[\pi(1-a_1)a_2 \\ & + (1-\pi)(1-a_2)a_1]\Big\} \end{aligned} \quad (2.37)$$

$$s.t.\ a_i = \arg\max\ [\sigma_i p w_1 + (1-\sigma_i)\frac{p(1-a_j)}{1-pa_j}w_2]a_i - c(a_i),\ \forall i = 1,2. \quad (2.38)$$

Before going into the detailed analysis, we can see from the objective function that it is a weighted average of the Principal's expected payoffs using the sequential inspection protocol with sequences (A_1, A_2) and (A_2, A_1) where the agents have to use the same efforts (a_1, a_2) in both cases. Recall from our previous analysis that the optimal efforts are no longer symmetric when using the sequential inspection. This implies that, whatever efforts the agents choose, the Principal's minimal expected payoff using sequential inspection cannot be achieved simultaneously for both realizations of sequences (A_1, A_2) and (A_2, A_1). This rough intuition tells us the Principal cannot be better off with random sequencing than sequential inspection. While we have been examining only the objective function in this constrained

optimization problem, let us now formalize this intuition in the following result:

Proposition 2.3.4. *Sequential inspection is a better choice of protocol for the Principal than random sequencing.*

Proof. We first characterize the optimal efforts as functions of transfers by solving the agent's problem. Plug in σ_1 and σ_2 from above and solve the transfers from the characterizations of a_1^{***} and a_2^{***}:

$$w_1 = \frac{-\frac{1}{\lambda}(1-a_1)[1+\frac{1}{\lambda}(1-pa_1)]c'(a_1) + \lambda(1-a_2)[1+\frac{1}{\lambda}(1-pa_1)]c'(a_2)}{\lambda p(1-a_2) - \frac{1}{\lambda}p(1-a_1)}, \quad (2.39)$$

$$w_2 = \frac{[1+\frac{1}{\lambda}(1-pa_1)]c'(a_1) - [1+\frac{1}{\lambda}(1-pa_1)]c'(a_2)}{\lambda p(1-a_2) - \frac{1}{\lambda}p(1-a_1)}. \quad (2.40)$$

Plug in the transfers and rewrite the Principal's problem as an unconstrained optimization problem:

$$\min_{a_1,a_2} K_R = \pi\Big\{pL(1-a_1)(1-a_2) + a_1c'(a_1) + a_2c'(a_2)(1-pa_1)\Big\} + (1-\pi)\Big\{pL(1-a_1)(1-a_2) + a_2c'(a_2) + a_1c'(a_1)(1-pa_2)\Big\}. \quad (2.41)$$

Comparing this objective function with K_{SQ}, this confirms our earlier conjecture that random sequencing is a convex combination of the sequential inspections. Given this, the strategy to show that the Principal is worse off with random sequencing is to show that, given the current optimal choices of efforts (a_1^{***}, a_2^{***}), the Principal's expected payoff from any of the two realizations of sequence is strictly higher than that when using sequential inspection. Consider the case[3] where the realized

[3]The analysis of the other case where the realized sequence is (A_2, A_1) is the same given the

sequence is (A_1, A_2), this is true so long as $(a_1^{***}, a_2^{***}) \neq (a_1^{**}, a_2^{**})$.

Rearrange the first order conditions to prepare for comparison with those characterizing the sequential solutions:

$$\begin{aligned} -pL(1-a_2^{***}) + c'(a_1^{***}) + a_1^{***}c''(a_1^{***}) =& \pi p a_2^{***} c'(a_2^{***}) \\ &+ (1-\pi) p a_2^{***}[c'(a_1^{***}) + a_1^{***}c''(a_1^{***})], \end{aligned} \tag{2.42}$$

$$\begin{aligned} -pL(1-a_1^{***}) + c'(a_2^{***}) + a_2^{***}c''(a_2^{***}) =& (1-\pi) p a_1^{***} c'(a_1^{***}) \\ &+ \pi p a_1^{***}[c'(a_2^{***}) + a_2^{***}c''(a_2^{***})]. \end{aligned} \tag{2.43}$$

We can similarly rearrange the first order conditions in the sequential analysis where A_1 is the first to inspect:

$$-pL(1-a_2^{**}) + c'(a_1^{**}) + a_1^{**}c''(a_1^{**}) = p a_2^{**} c'(a_2^{**}), \tag{2.44}$$

$$-pL(1-a_1^{**}) + c'(a_2^{**}) + a_2^{**}c''(a_2^{**}) = p a_1^{**}[c'(a_2^{**}) + a_2^{**}c''(a_2^{**})]. \tag{2.45}$$

Suppose $(a_1^{***}, a_2^{***}) = (a_1^{**}, a_2^{**}) = (a_1, a_2)$, the left hand side in both sets of equations are equal, which implies that the right hand side of the corresponding equations must be equal as well:

$$c'(a_2) - c'(a_1) = a_1 c''(a_1), \tag{2.46}$$

$$c'(a_1) - c'(a_2) = a_2 c''(a_2), \tag{2.47}$$

which cannot hold at the same time since the cost function is strictly convex and we have shown before that $(a_1^{**}, a_2^{**}) \neq (0, 0)$. Since π is arbitrary, this implies that

symmetry of the agents.

$(a_1^{***}(\pi), a_2^{***}(\pi)) \neq (a_1^{**}, a_2^{**})$, $\forall \pi \in (0,1)$. Given this, the minimal value of

$$\min_{a_i, a_j} \; pL(1-a_i)(1-a_j) + a_i c'(a_i) + (1-pa_i)a_j c'(a_j) \tag{2.48}$$

cannot be achieved, and therefore, the Principal cannot be better off with random sequencing. ■

Given that sequential inspection is the boundary case of random sequencing with $\pi = 0$ or 1 and the symmetry of the problem, a natural guess would be $\pi = \frac{1}{2}$ is the worst for the Principal. We first present a lemma as a preparation for proving this corollary.

Lemma 2.3.5. *$a_1^{***}(\pi) = a_2^{***}(\pi)$ if and only if $\pi = \frac{1}{2}$.*

Proof. If $\pi = \frac{1}{2}$, note from equations (2.42) and (2.43) which characterizes the optimal efforts, these two equations become symmetric once we plug in $\pi = \frac{1}{2}$ and therefore it has to be $a_1^{***} = a_2^{***}$.

For the only if part, suppose[4] $a_1^{***} = a_2^{***} = a \neq 0$, equations (2.42) and (2.43) are satisfied at the same time which implies that:

$$(1-\pi)pa^2c''(a) = \pi pa^2c''(a) \Longrightarrow \pi = \frac{1}{2}, \text{ since } a \neq 0. \tag{2.49}$$

■

Corollary 2.3.6. *Among all random sequencing protocols, the Principal is worst off to randomize the sequences (A_1, A_2) and (A_2, A_1) with equal probabilities $\frac{1}{2}$.*

[4]It's easy to check $a_1^{***} = a_2^{***} = 0$ is never optimal.

Proof. Apply the Envelope Theorem:

$$\frac{\partial}{\partial \pi} K_R(a_1^{***}(\pi), a_2^{***}(\pi), \pi) = p a_1^{***}(\pi) a_2^{***}(\pi)(c'(a_1^{***}(\pi)) - c'(a_2^{***}(\pi))) = 0,$$
$$\text{if and only if } \pi = \frac{1}{2} \text{ by Lemma (2.3.5).} \quad (2.50)$$

Since $c \in C^3$ which guarantees enough smoothness such that no interior point $\pi \in (0, \frac{1}{2}) \bigcup (\frac{1}{2}, 1)$ can be a local maximizer or minimizer. It remains to compare $\pi = \frac{1}{2}$ with the boundary cases $\pi = 0$ and 1 which are the sequential inspections. Proposition 2.3.4 implies that:

$$K_R(a_1^{***}, a_2^{***})\big|_{\pi=\frac{1}{2}} > K_R(a_1^{***}, a_2^{***})\big|_{\pi=0} = K_R(a_1^{***}, a_2^{***})\big|_{\pi=1}. \quad (2.51)$$

■

Let us delay the discussion of these two results until after we compare this protocol with simultaneous inspection:

Lemma 2.3.7. *Random sequencing is a better choice of protocol for the Principal than simultaneous inspection.*

Proof. Rearrange the objective function in equation (2.41) from the previous proof:

$$\min_{a_1, a_2} K_R = pL(1-a_1)(1-a_2) + a_1 c'(a_1) + a_2 c'(a_2) - p a_1 a_2 [\pi c'(a_2) + (1-\pi) c'(a_1)] \quad (2.52)$$

Compare the Principal's expected payoff functions:

$$K_{SM} - K_R = p a_1 a_2 [\pi c'(a_2) + (1-\pi) c'(a_1)] > 0, \quad (2.53)$$

since the cost function is strictly increasing. K_R everywhere below K_{SM} then implies

that $\min K_R < \min K_{SM}$. ■

The intuition for why random sequencing dominates simultaneous inspection is essentially the same as before when we compared sequential with simultaneous: the Principal avoids receiving a duplicate message confirming the bad state (i.e. the outcome (B, B)). Extending this intuition also helps shed light on why sequential inspection dominates random sequencing. Since fully informing the second inspector about what was learned from the first inspector (i.e. it has to be that the first inspector has received an empty message) allows the Principal to be more efficient in using a transfer to induce effort, randomizing the order of the inspectors damages the value of learning and therefore makes it more costly for the Principal. Along this line of reasoning, it would not be hard to imagine that randomizing with equal probabilities turns out to be the worst among all possible randomizations.[5]

Besides the boundary case $\pi = 0$ and 1 which is the sequential inspection, the other potentially interesting case is randomizing with equal probabilities. While we have proved that it is the worst choice, a related question to ask is if we could find a class of contracts where randomizing with equal probabilities can be optimal. We have discussed the value of the ordering of information to the agent and the Principal which explains why randomizing with equal probabilities is worst as it maximally reduces this information among all choices of π. Given this, if we restrict the class of the contracts and make the Principal less effective at inducing efforts using transfers, the advantage of sequential inspections might vanish. It turns out that randomizing with equal probabilities is indeed optimal if the Principal can no longer customize the rewards to successful detection on the order of the inspector:

[5] To what extent this conjecture depends on our model setup in terms of assumptions on functional forms and properties is beyond the scope of this paper.

Proposition 2.3.8. *Consider the class of contracts where the agent receives a transfer w for a successful detection regardless of his order in the sequence. The Principal is best off randomizing the sequences (A_1, A_2) and (A_2, A_1) with equal probabilities $\frac{1}{2}$.*

Proof. The approach we have taken before is to substitute out the transfers using the first order conditions of the agents' problem and transform the Principal's constrained optimization problem to an unconstrained optimization problem with effort choices. The restriction on the class of contracts means that we have one transfer variable but two agents. This implies that we can no longer freely adjust the efforts when solving the Principal's problem and the transformed problem now has a constraint which comes from the first order conditions in the agents' problem:

$$c'(a_1) = pw\frac{1 + \lambda(1 - a_2)}{1 + \lambda(1 - pa_2)}, \tag{2.54}$$

$$c'(a_2) = pw\frac{\lambda + 1 - a_1}{\lambda + 1 - pa_1}. \tag{2.55}$$

Divide both sides and rearrange the equation to get the following equality constraint on a_1, a_2:

$$\begin{aligned} g(a_1, a_2, \pi) = {} & c'(a_1)\Big\{(1-\pi)^2(1-pa_2) + \pi(1-\pi)[1 + (1-pa_2)(1-a_1)] + \pi^2(1-a_1)\Big\} \\ & - c'(a_2)\Big\{(1-\pi)^2(1-a_2) + \pi(1-\pi)[1 + (1-pa_1)(1-a_2)] + \pi^2(1-pa_1)\Big\} = 0. \end{aligned} \tag{2.56}$$

The Principal's problem with $w_1 = w_2 = w$ becomes:

$$\min_{a_1,a_2} K_R^w = pL(1-a_1)(1-a_2) + a_1c'(a_1) + a_2c'(a_2) \\ - pa_1a_2[\pi c'(a_2) + (1-\pi)c'(a_1)] \tag{2.57}$$

$$s.t.\ g(a_1, a_2, 0) = 0. \tag{2.58}$$

Note that at $\pi = \frac{1}{2}$, both K_R^w and $g(a_1, a_2, \frac{1}{2})$ are symmetric with respect to a_1 and a_2, which implies that the solutions should be equal with $a_1^{***}(\frac{1}{2}) = a_1^{***}(\frac{1}{2})$. This then implies the constraint always holds and can therefore be discarded by setting the Lagrangian multiplier $\lambda_L = 0$ without loss when $\pi = \frac{1}{2}$, since:

$$g(a, a, \frac{1}{2}) = 0, \forall a. \tag{2.59}$$

Therefore, we can apply the Envelope Theorem for the unconstrained optimization case when $\pi = \frac{1}{2}$:

$$\frac{\partial}{\partial \pi} K_R^w(a_1^{***}(\pi), a_2^{***}(\pi), \pi)\Big|_{\pi=\frac{1}{2}} = pa_1^{***}(\frac{1}{2})a_2^{***}(\frac{1}{2})(c'(a_1^{***}(\frac{1}{2})) - c'(a_2^{***}(\frac{1}{2}))) = 0. \tag{2.60}$$

■

2.4 Extension: Sequential Inspection with Asymmetric Agents

In the benchmark analysis above, we have shown the superiority of the sequential inspection protocol among all three protocols considered when agents are symmetric. If we continue to use this sequential protocol but relax the assumption

on agents to allow them to differ in their information acquisition technologies, the order of the agents in the sequence becomes relevant to the Principal's decision making as it affects the agents' belief updating and hence the effort choice. In this section, we consider a variation of the model where the two agents are asymmetric in (a) capabilities, i.e. the probabilities of discovering the defect with the same effort choice are different; (2) direction, i.e. one agent specializes in looking for good news that confirms the good state while the other agent specializes in bad news that confirms the bad state, and compare the possible orderings of the agents within the class of sequential inspections.

2.4.1 Asymmetric in Capabilities

What if agents are no longer equally skilled at detecting defects? In particular, if we choose to ask inspectors one by one, should we approach the better expert first or second? In this section, we extend the analysis of sequential inspection by assuming one agent being fully capable as before while the other one being less capable $(\theta < 1)$, and compare the two possible sequences: (1) ask the more skilled inspector first $(\theta_1 = 1, \theta_2 = \theta)$; (2) ask the less skilled inspector first: $(\theta_1 = \theta, \theta_2 = 1)$.

Recall from our model setup, each agent's probability of successful detection when there is indeed a defect is $\theta_i a_i$, and their capability θ_i is common knowledge. As we have analyzed in the previous section, there are three possible outcomes for signal realizations, $\{(B), (\varnothing, B), (\varnothing, \varnothing)\}$, when the Principal chooses the sequential inspection protocol. Using the same notations for the transfers w_1 and w_2 rewarding

each agent's successful detection, we can write out the Principal's problem:

$$\min_{a_1,a_2,w_1,w_2} p[L(1-\theta_1 a_1)(1-\theta_2 a_2) + \theta_1 a_1 w_1 + \theta_2 a_2(1-\theta_1 a_1)w_2] \tag{2.61}$$

$$s.t.\ a_2 = \arg\max \frac{p(1-\theta_1 a_1)}{1-p\theta_1 a_1}\theta_2 a_2 w_2 - c(a_2), \tag{2.62}$$

$$a_1 = \arg\max\ p\theta_1 a_1 w_1 - c(a_1), \tag{2.63}$$

where $\frac{p(1-\theta_1 a_1)}{1-p\theta_1 a_1}$ is A_2's updated belief on the state being bad given his knowledge of his order in the sequence and A_1's capability. Transform the Principal's problem in the same way by substituting out the optimal transfers solved from the agents' problems, we have:

$$\min_{a_1,a_2} pL(1-\theta_1 a_1)(1-\theta_2 a_2) + a_1 c'(a_1) + (1-p\theta_1 a_1)a_2 c'(a_2). \tag{2.64}$$

We had a result on effort comparison earlier when agents are symmetric that the second inspector puts in more effort than the first inspector. Since the efforts are substitutes to the Principal, we could tell from this result that the second inspector's effort has to be the cheaper one between the two agents. What this implies is that the agents' optimal efforts should be closer to each other when the second inspector becomes less capable with $\theta_2 = \theta$, since the second inspector's effort becomes less cheap and hence the two agents' efforts are less substitutable; and, similarly, the agents' optimal efforts should be further apart from each other when the first inspector becomes less capable with $\theta_1 = \theta$ for the opposite reason. We formalize this intuition in the following lemma:

Lemma 2.4.1. *Compared to the case when the second inspector is more capable* $(\theta_1 = \theta, \theta_2 = 1)$*, the first (second) inspector puts in more (less) effort when the first inspector is more capable* $(\theta_1 = 1, \theta_2 = \theta)$*.*

Proof. Case 1: $(\theta_1 = 1, \theta_2 = \theta)$. FOCs are:

$$c'(a_1) + a_1 c''(a_1) = pa_2 c'(a_2) + pL(1 - \theta a_2), \forall a_2, \tag{2.65}$$

$$c'(a_2) + a_2 c''(a_2) = \frac{\theta pL(1 - a_1)}{1 - pa_1}, \forall a_1. \tag{2.66}$$

Case 2: $(\theta_1 = \theta, \theta_2 = 1)$. FOCs are:

$$c'(a_1) + a_1 c''(a_1) = \theta pa_2 c'(a_2) + \theta pL(1 - a_2), \forall a_2, \tag{2.67}$$

$$c'(a_2) + a_2 c''(a_2) = \frac{pL(1 - \theta a_1)}{1 - \theta pa_1}, \forall a_1. \tag{2.68}$$

Compare the right-hand side of FOC for a_1 in both cases:

$$pa_2 c'(a_2) + pL(1 - \theta a_2) > \theta pa_2 c'(a_2) + \theta pL(1 - a_2), \forall a_2. \tag{2.69}$$

Given our assumptions on the cost function, there exists a continuously differentiable and monotonically decreasing function $a_1(a_2)$ by the Implicit Function Theorem, and the right-hand side comparison implies that $a_1(a_2)$ from case 1 is everywhere above that function from case 2. Similarly, we have a continuously differentiable and monotonically decreasing function $a_2(a_1)$ and that the right-hand side comparison of FOCs for a_2 implies that $a_2(a_1)$ from case 1 is everywhere below that function from case 2. Combine the comparisons of the reaction functions, we could compare the intersection points (a_1^1, a_2^1) for case 1 and (a_1^2, a_2^2) for case 2. ■

However, this effort comparison would not help with comparing which sequence is better for the Principal given the difference in the Principal's objective functions. Given this difficulty in obtaining an analytical solution, we take a specific quadratic cost function to illustrate the potential comparison in an example.

Example 2.4.2. *Plug in cost function $c(a) = a^2$ and solve for the optimal efforts:*

Case 1 ($\theta_1 = 1, \theta_2 = \theta$):

From the FOCs we have:

$$a_2 = \frac{\theta p L(1 - a_1)}{4(1 - pa_1)}, \tag{2.70}$$

Use this to further simplify the Principal's problem:

$$\min_{a_1} \; K^1 = 2a_1^2 + \frac{\theta^2 p^2 L^2 (1 - a_1)^2}{8(1 - pa_1)} + pL(1 - a_1)\Big[1 - \frac{(1 - a_1)\theta^2 pL}{4(1 - pa_1)}\Big], \tag{2.71}$$

and a_1 is then solved implicitly from the following equation:

$$32p^2a_1^3 + (-64p - 8p^3L + \theta^2p^3L^2)a_1^2 + (32 + 16p^2L - 2\theta^2p^2L^2)a_1 = 8pL - 2\theta^2p^2L^2 + \theta^2p^3L^2. \tag{2.72}$$

Plug in the implicit solution to have $K^1(p, \theta, L)$.

Case 2 ($\theta_1 = \theta, \theta_2 = 1$):

From the FOCs we have:

$$a_2 = \frac{pL(1 - \theta a_1)}{4(1 - \theta pa_1)}, \tag{2.73}$$

Use this to further simplify the Principal's problem:

$$\min_{a_1} \; K^2 = \frac{1}{-8(1 - \theta pa_1)}\big[16\theta pa_1^3 + (-16 - 8\theta^2p^2L + \theta^2p^2L^2)a_1^2 + 2\theta pL(4(1 + p) - pL)a_1 + (-8 + pL)pL\big], \tag{2.74}$$

and a_1 is then solved implicitly from the following equation:

$$32p^2\theta^2a_1^3 + (-64p\theta - 8p^3\theta^3L + \theta^3p^3L^2)a_1^2 + (32 + 16p^2\theta^2L - 2\theta^2p^2L^2)a_1 = 8\theta pL - 2\theta p^2L^2 + \theta p^3L^2. \quad (2.75)$$

Plug in the implicit solution to have $K^2(p, \theta, L)$.

Mathematica gives the result $K^1 < K^2$ for parameter values $p \in (0.1, 0.9)$, $\theta \in (0.1, 0.9)$, and $L \in (0.1, 9.9)$ which satisfy Assumption 2.2.1.

This example suggests it is better for the Principal to start with the more skilled inspector at least in some cases. While we do not have a complete intuition for this comparison due to the lack of analytical solutions which would help us identify the factors driving the final result, following our local arguement for the effort comparison under sequential inspection in Proposition 2.3.2 where we showed that inducing effort of the first inspector is more costly than the second inspector, we might guess that convexity of the cost function implies that the Principal should ask the less costly inspector, which is the more skilled inspector, to inspect first.

2.4.2 Asymmetric in Direction

So far the Principal's payoff depends only on the bad state. To set the stage for the analysis with agents being able to acquire both good and bad news, we consider a

variation of the Principal's preference where both states matter:

$$v(a,\omega) = \begin{cases} -L, & \text{if } a = n, \omega = B \\ 0, & \text{if } a = n, \omega = G \\ -T, & \forall \omega, \text{ if } a = t \end{cases} \cdot \tag{2.76}$$

To allow for the Principal's preferred action being state-dependent, we need an assumption on the comparison of payoffs L and T such that the action t (treatment) is preferred when the state is bad; and the action n (no treatment) is preferred when the state is good:

Assumption 2.4.3. $L > T$.

Now that both states matter to the Principal, we can have agents learn about both the good and the bad states. In particular, each agent specializes in looking for either good or bad news, and they could either receive a conclusive signal confirming the true state or an empty message. Their technologies are otherwise symmetric. To accommodate the introduction of the new parameters and the more complicated signal structures in the setup, instead of general convex cost functions, we work with quadratic costs of the following functional form to make the analysis tractable:

$$c(a) = \frac{a^2}{2}. \tag{2.77}$$

Since the agents are no longer symmetric in the current setting, we switch the notations for their efforts from a_1 and a_2 to a for the agent who specializes in looking for good news and b for bad news.

As we've seen from the previous section, the analysis for comparison of different sequences with asymmetric agents is complicated due to the difference in the

Principal's objective functions, which means all parameter values could matter. To simplify the analysis, we first allow the Principal to commit to an action after getting two empty messages in a row. This means that the first inspector's effort affects his own probability of getting a conclusive message as well as his follower's belief, and the second inspector's effort affects his signals in the same way, but their efforts would no longer influence the Principal's posterior belief and hence final action. Moreover, the agents' information acquisition technology also implies that only the agent who could get a conclusive signal leading to the action different from the committed action is useful to the Principal. We will see that this simplification turns out to be helpful for the algebra as well as yielding a neat and intuitive result.

Committed Action: We first formalize the observation that only one inspector will be asked in this case.

Observation 2.4.4. *The Principal will only ask the good-news (bad-news) agent to inspect if he has committed to treat (not treat) when receiving only empty messages.*

Proof. Suppose the Principal commits to treat the patient if all the signals he has collected from the agents are empty messages, i.e. the outcome is $(\varnothing, \varnothing)$. Consider the bad-news agents who could get either a signal confirming the bad state or an empty message as the second inspector. While his effort is costly, his information acquisition is of no value to the Principal as he is going to take the same action, treat, in either case. Therefore, it would be optimal for the Principal to set the effort of the bad-news agent at zero. ■

Now that we know only one inspector will be asked, we only need to find out which action the Principal should commit to which then determines which inspector he should ask:

Proposition 2.4.5. *When the default action at the prior belief is to not treat (treat), it is optimal to ask the bad-news (good-news) agent.*

Proof. We first solve the Principal's problem with both committed actions. (1) Commit to Treat:

The Principal's problem is:

$$\min_a V^t = [(1-p)(1-a)+p]T + \frac{a^2}{2}. \tag{2.78}$$

Solve a from FOC, we have:

$$a^* = \begin{cases} (1-p)T, & \text{if } p \in [\max\{1-\frac{1}{T}, 0\}, 1] \\ 1, & o.w. \end{cases}, \tag{2.79}$$

and

$$V(a^*) = \begin{cases} T - \frac{(1-p)^2T^2}{2}, & \text{if } p \in [\max\{1-\frac{1}{T}, 0\}, 1] \\ pT + \frac{1}{2}, & o.w. \end{cases}. \tag{2.80}$$

(2) Commit to Not Treat:

The Principal's problem is:

$$\min_b V^{nt} = pbT + p(1-b)L + \frac{b^2}{2}. \tag{2.81}$$

Solve b from FOC, we have:

$$b^* = \begin{cases} p(L-T), & \text{if } p \in [0, \min\{1, \frac{1}{L-T}\}] \\ 1, & o.w. \end{cases}, \tag{2.82}$$

and

$$V(b^*) = \begin{cases} pL - \frac{p^2(L-T)^2}{2}, & \text{if } p \in [0, \min\{1, \frac{1}{L-T}\}] \\ pT + \frac{1}{2}, & o.w. \end{cases} . \tag{2.83}$$

Next, we consider all possible cases depending on the values of parameters p, L, T.

1. $T > 1, L - T < 1$

   (a) $p \in [0, 1 - \frac{1}{T}]$

$$V(a^*) - V(b^*) \tag{2.84}$$

$$= pT + \frac{1}{2} - (pL - \frac{p^2(L-T)^2}{2}) \tag{2.85}$$

$$= \frac{1}{2}[p(L-T) - 1]^2 > 0 \tag{2.86}$$

   Optimal to use the bad-news agent. Note in this case, the default is not to treat.

   (b) $p \in [1 - \frac{1}{T}, 1]$

$$V(a^*) - V(b^*) \tag{2.87}$$

$$= T - \frac{(1-p)^2T^2}{2} - (pL - \frac{p^2(L-T)^2}{2}) \tag{2.88}$$

$$= (T - pL)\{1 - \frac{1}{2}[(L-2T)p + T]\}, \tag{2.89}$$

   which is strictly positive when $T > pL$ (i.e. $p \in [1 - \frac{1}{T}, \frac{T}{L})$) and strictly negative when $T < pL$ (i.e. $p \in (\frac{T}{L}, 1]$), since $L < 2T$. Optimal to use the bad-news agent when the default at the prior belief is not to treat and good-news agent when the default is to treat.

2. $T < 1, L - T < 1$

For any p:

$$V(a^*) - V(b^*) \tag{2.90}$$
$$= T - \frac{(1-p)^2T^2}{2} - (pL - \frac{p^2(L-T)^2}{2}) \tag{2.91}$$
$$= (T - pL)\{1 - \frac{1}{2}[(L-2T)p + T]\}, \tag{2.92}$$

which is strictly positive when $T > pL$ and strictly negative when $T < pL$, since $(L-2T)p + T \in [\min\{L-T, T\}, \max\{L-T, T\}]$ and $\max\{L-T, T\} < 1$. Optimal to use the bad-news agent when the default at the prior belief is not to treat and good-news agent when the default is to treat.

3. $T < 1, L - T > 1$

   (a) $p \in [0, \frac{1}{L-T}]$

$$V(a^*) - V(b^*) \tag{2.93}$$
$$= T - \frac{(1-p)^2T^2}{2} - (pL - \frac{p^2(L-T)^2}{2}) \tag{2.94}$$
$$= (T - pL)\{1 - \frac{1}{2}[(L-2T)p + T]\}, \tag{2.95}$$

which is strictly positive when $T > pL$ (i.e. $p \in [0, \frac{T}{L})$) and strictly negative when $T < pL$ (i.e. $p \in (\frac{T}{L}, \frac{1}{L-T}]$), since $(L-2T)p+T \in [T, \frac{L-2T}{L-T}+T]$ and $\frac{L-2T}{L-T} + T < 2$ as $T(L-T) < L$. Optimal to use the bad-news agent when the default at the prior belief is not to treat and good-news agent when the default is to treat.

(b) $p \in [\frac{1}{L-T}, 1]$

$$V(a^*) - V(b^*) \tag{2.96}$$

$$= T - \frac{(1-p)^2 T^2}{2} - (pT + \frac{1}{2}) \tag{2.97}$$

$$= -\frac{1}{2}[(1-p)T - 1]^2 < 0 \tag{2.98}$$

Optimal to use the good-news agent. Note in this case, the default is to treat.

4. $T > 1, L - T > 1$

(a) $p \in [0, \min\{1 - \frac{1}{T}, \frac{1}{L-T}\}]$

$$V(a^*) - V(b^*) \tag{2.99}$$

$$= pT + \frac{1}{2} - (pL - \frac{p^2(L-T)^2}{2}) \tag{2.100}$$

$$= \frac{1}{2}[p(L-T) - 1]^2 > 0 \tag{2.101}$$

Optimal to use the bad-news agent. Note in this case, the default is not to treat since:

$$pL - T \leqslant \min\{L - T - \frac{L}{T}, \frac{L - LT + T^2}{L - T}\} \tag{2.102}$$

$$\propto \min\{(L-T)(LT - T^2 - L), T(L - LT + T^2)\} < 0 \tag{2.103}$$

(b) $p \in [\max\{1 - \frac{1}{T}, \frac{1}{L-T}\}, 1]$

$$V(a^*) - V(b^*) \tag{2.104}$$
$$= T - \frac{(1-p)^2 T^2}{2} - (pT + \frac{1}{2}) \tag{2.105}$$
$$= -\frac{1}{2}[(1-p)T - 1]^2 < 0 \tag{2.106}$$

Optimal to use the good-news agent. Note in this case, the default is to treat since:

$$pL - T \geqslant \max\{L - T - \frac{L}{T}, \frac{L - LT + T^2}{L - T}\} \tag{2.107}$$
$$\propto \max\{(L-T)(LT - T^2 - L), T(L - LT + T^2)\} > 0 \tag{2.108}$$

(c) $p \in [\min\{1 - \frac{1}{T}, \frac{1}{L-T}\}, \max\{1 - \frac{1}{T}, \frac{1}{L-T}\}]$

i. $L - T > \frac{L}{T}$

$$V(a^*) - V(b^*) \tag{2.109}$$
$$= pT + \frac{1}{2} - (pT + \frac{1}{2}) = 0 \tag{2.110}$$

Optimal to use either one.

ii. $L - T < \frac{L}{T}$

$$V(a^*) - V(b^*) \tag{2.111}$$
$$= T - \frac{(1-p)^2 T^2}{2} - (pL - \frac{p^2 (L-T)^2}{2}) \tag{2.112}$$
$$= (T - pL)\{1 - \frac{1}{2}[(L - 2T)p + T]\}, \tag{2.113}$$

which is strictly positive when $T > pL$ (i.e. $p \in [1 - \frac{1}{T}, \frac{T}{L})$) and

strictly negative when $T < pL$ (i.e. $p \in (\frac{T}{L}, \frac{1}{L-T}]$), since $(L-2T)p+T \in [\min\{(L-2T)(1-\frac{1}{T})+T, \frac{L-2T}{L-T}+T\}, \max\{(L-2T)(1-\frac{1}{T})+T, \frac{L-2T}{L-T}+T\}]$ and $\max\{(L-2T)(1-\frac{1}{T})+T, \frac{L-2T}{L-T}+T\} < 2$. Optimal to use the bad-news agent when the default at the prior belief is not to treat and good-news agent when the default is to treat.

■

This result says that the Principal should commit to his default action to save money on information acquisition. Suppose the default action at the prior belief is to not treat but the Principal chooses to commit to treat after getting empty messages. As our earlier Observation says, the Principal should then only ask the good-news agent for the signals to be valuable. However, the default action is to not treat implies that it's very likely for the agent to get the conclusive signal confirming the good state. Inducing effort to generate a conclusive signal is costly, and therefore, this higher likelihood implies that the Principal is paying more by asking the good-news agent. Essentially, if we're ever going to pay for information, the money is used more efficiently if we spend it on looking for evidence for the unlikely case contradicting the prior.

Chapter 3

Self Persuasion of a Present Biased Agent

3.1 Introduction

How people discount the future is an important component in the analysis of their consumption saving decisions. A major cognitive bias related to discounting is that people may discount over the same amount of time differently, depending on how far into the future this time period is. This change in discounting over time then gives rise to dynamic inconsistency, which captures the phenomenon that the actual consumption saving decision made at a future time is inconsistent with the optimal decision for that time viewed from the present. To alleviate the dynamic inconsistency problem, there has been a long tradition in looking for an optimally constrained future budget set with the availability of certain commitment device (Strotz, 1955; Laibson, 1997; Amador, Werning, and Angeletos, 2006). Parallel to this traditional approach, this paper explores the possibility of using informational tools when the return to savings is uncertain.

To illustrate the informational tool that we have in mind, let us consider the following scenario: knowing that some money will be deposited to our 401(k) account at the next pay day, suppose that we want to figure out how much of it should be invested in some ETF, say Vanguard Total Stock Market. We could wait and simply pick a number at the time when that money arrives, based on some understanding of the investment holdings and the past performance of this ETF. Alternatively, while waiting, we could spend the time reading through newspa-

pers and financial analysts' reports, which are freely available online, to be more informed about how likely it is going to be a bull market. If our belief in the market condition is influenced by the reading, we might end up choosing a different number to invest.

If an agent is not present biased, acquiring information to be fully informed of the uncertain state is always optimal, as it allows the agent to make a state-contingent action plan (Blackwell, 1953). Once the agent becomes present biased, he is prone to over-consuming. In this case, there might be a state where the action that he prefers to take today differs from that tomorrow, which renders acquiring full information possibly suboptimal. This gives the agent an incentive to sometimes acquire less information as a way to induce different future actions.

The usefulness of the informational tool to a present biased agent has been previously studied in Carrillo and Mariotti (2000) and Bénabou and Tirole (2002). Despite the similarities in the model setup, we ask a different question here: what is the agent's ex-ante optimal information choice? Our question is essentially an information design question, which we answer by applying the concavification approach (Aumann and Maschler, 1995; Kamenica and Gentzkow, 2011) to solve the intra-personal game between the incarnations of the agent at different times.

Our main finding concerns an agnet with CRRA preference, and proves that full information turns out to be optimal for risk averse agents so long as they are not too present biased relative to their risk attitude, so the Blackwell result is locally robust as we start to accommodate the cognitive limitation of a behavioral agent. Once the agent becomes sufficiently present biased relative to their risk attitude, restricting access to information might be beneficial and the states matter. We find a threshold for the uncertain state as a function of only the risk aversion

and the present bias parameters for the risk averse expected-utility-maximizing agent. Any state that is above this threshold is good enough for the agent and their senstivity to further improvements of future prospects decreases. This creates the room for partial information to dominate full information as the current self could make the future self much more likely to save a decent amount of their income instead of saving a little bit more but way less likely.

We can see from our discussion above that risk aversion plays an important role in the determination of the optimal information choice to the extent that it could completely cancel out the effect of the cognitive limitation and restore the optimality of full information in certain cases. We understand the mechanism behind this effect is complicated by the double-role of the risk aversion parameter in the CRRA utility function, which ties together the agent's attitude towards tradeoffs both across time and across states. Present bias affects discounting and hence intertemporal tradeoffs. We might guess that its implication for the elasticity of intertemporal substitution allows the risk aversion parameter to cancel out the effect of the present bias. This conjecture is supported as we show that for a version of Epstein-Zin preference which separates risk aversion from elasticity of intertemporal substitution, the prediction that the behavioral agent behaves in exactly the same way as a standard agent disappears when the consumption decision is very sensitive to changes in the return on savings.

We first discuss the related literature in the remainder of this section before introducing our model and discussing the main assumptions. After that, we apply the concavification approach and explain main intuitions in a benchmark model, followed by our main result. The extension to EZ-preference is presented in the last.

3.1.1 Related Literature

Carrillo and Mariotti (2000) and Bénabou and Tirole (2002) are among the first attempts to investigate the value of having an information choice in mitigating the sub-optimality of a present biased agent's future decision. Carrillo and Mariotti (2000) model the information choice as the choice of when to stop an experimentation process, and prove the existence of an equilibrium where the agent stops before he learns completely by reaching the correct expected value of the uncertain state. In contrast, we consider an information design problem. Bénabou and Tirole (2002) study optimal information either with or without the commitment assumption.[1] In their Section I where they assume commitment, they consider a very restrictive set of information choices containing only no and full information. As a comparison, we impose no restrictions other than what is implied by the agent's Bayesian rationality. Later in their Section IV, they consider the ex-ante welfare of self-0 where information obtained could be remembered and passed on to self-1 at a zero cost, which is equivalent to our problem under the commitment assumption. However, we model the information choice differently. In their paper, the information tool is modeled as memory where self-0 could choose what to tell self-1 after observing the true state herself. Essentially, there is no common prior and this is a problem of information disclosure. We model the information choice as the design of information, or the choice of a statistical experiment, where both agents have the same information at the time when self-0 makes the decision.

Our paper fits broadly into the recent literature on information design as surveyed in Bergemann and Morris (2019). In particular, we follow the concavification approach (Kamenica and Gentzkow, 2011; Aumann and Maschler, 1995) to solve

[1] Their demand side analysis in Section I considers the optimal information with respect to the ex-ante indirect utility function; their supply side analysis in Section II considers the optimal signal (or memory) after the realization of the low state.

the one-sender-one-receiver Bayesian persuasion problem.[2] There are some recent works that similarly consider a Bayesian persuasion problem for a behavioral agent. Lipnowski and Mathevet (2018) considers a large class of behavioral agents with intrinsic preference for information, including a Gul and Pesendorfer (2001) agent with self control problems, and discuss the optimal information provision with respect to different relationships of the psychological and instrumental value of information. Our focus is on the characterization of the optimal information for a present biased agent. Mariotti, Schweizer, Szech, and von Wangenheim (2018) is a recent paper that extends Carrillo and Mariotti (2000) to study the optimal information also for a present biased agent. While their main contribution is on the optimal design of information when agents are heterogenous in their present bias, we focus on the single agent scenario but explore how the two different roles of the parameter in the CRRA utility function, risk aversion and elasticity of intertemporal substitution, affect the information choice.

Within the information design literature, there are also papers extending the analysis to incorporate non-standard preferences. Beauchêne, Li, and Li (2019) extends the standard Bayesian persuasion problem by allowing both the sender and the receiver to be ambiguity averse and characterize the optimal (ambiguous) information. We have a section on accommodating Epstein-Zin preference, but the objective function can be transformed into an expected utility form, and therefore the optimal information is a smaller departure from that in a standard Bayesian persuasion problem.

Adopting the Epstein-Zin formulation for a consumption saving problem under uncertainty has been a common practice in the asset pricing literature as a response

[2]Other approaches include the linear programming approach (Kolotilin, 2018; Lipnowski and Mathevet, 2017), the information design approach (Bergemann and Morris, 2019), and the price-theoretic approach (Dworczak and Martini, ming).

to the equity premium puzzle, for example in the long run risk literature started by Bansal and Yaron (2004). Our exercise might be helpful for further explorations of optimal information disclosure in asset pricing problems.

3.2 Model

Consider a quasi-hyperbolic agent with a three-period time horizon: $t = 0, 1$, and 2. In $t = 1$, an endowment $y = 1$ arrives, and the agent makes a consumption saving decision choosing $c \in \mathcal{A} = [0, 1]$.[3] The gross rate of return on savings, R, is uncertain. This payoff-relevant uncertain state $R \in \Omega = \{\underline{R}, \overline{R}\}$ is realized but unobservable in $t = 1$, which determines the consumption in $t = 2$. Before the endowment arrives, in $t = 0$, the agent who is endowed with the correct prior belief on the high state $p^* = Prob(R = \overline{R})$ makes a costless information acquisition decision to learn about the state by choosing $\mu \in \Delta(\Delta\Omega)$ for $t = 1$ in a *Bayes plausible* fashion.[4] Specifically, since the state space Ω being binary implies that the belief p is one-dimensional and therefore the space of posterior beliefs $\Delta\Omega$ is simply $[0, 1]$, the set of all *Bayes plausible* distributions is:

$$\mathcal{M} = \{\mu \in \Delta([0, 1]) | \int_0^1 p\mu(dp) = p^*\}.$$

Consumption in $t = 1$ is then chosen optimally for the random belief $p \in supp(\mu)$ realized in that period. The agent's preference over the uncertain consumptions is represented by the von Neumann Morgenstern (vNM) constant relative risk

[3] Following Kamenica and Gentzkow (2011), we consider the *effective* action space, that is, $\forall c \in \mathcal{A}$, there exists p such that c is an optimal choice. Our action space $\mathcal{A}$ is necessarily continuous in the general setting.

[4] As defined in Kamenica and Gentzkow (2011), "a distribution of posteriors is *Bayes plausible* if the expected posterior probability equals the prior", which translates to $\int_{\Delta(\Omega)} p\mu(dp) = p^*$ using our notations.

aversion (CRRA) utility function with a coefficient of RRA $1-\rho$ for $\rho \in (-\infty, 1]$.

The agent has two choices to make: the information μ in $t = 0$ and the consumption c in $t = 1$. These two choices are interdependent. The optimality of a consumption choice in $t = 1$ depends on the posterior belief that is realized in that period. At the same time, the optimality of a belief distribution in $t = 0$ depends on how the consumption choice to be made in the following period would respond to it. It has been standard practice to formalize this interdependence of the choices in different periods as a strategic interaction between two separate agents that are indexed respectively by the period when they need to make a choice. Following the intra-personal game literature (Carrillo and Mariotti, 2000; Bénabou and Tirole, 2002), we call the agent in $t = 0$ choosing the information *self-0* (*she*) and that in $t = 1$ choosing the consumption *self-1* (*he*). To solve this intra-personal game, we use the Sender-preferred subgame perfect equilibrium, which is defined in Kamenica and Gentzkow (2011), as our solution concept.[5]

3.2.1 Modeling the information choice

Here we have chosen to model the agent's information choice in $t = 0$ as the choice of a *Bayes plausible* distribution of posterior beliefs μ. Alternatively, we could model it as the choice of a set of signals, $\{\pi(\cdot|R) : \Omega \to \Delta\mathcal{S}\}$, which are conditional distributions over some signal realization space $\mathcal{S}$.[6] The equivalence between these two approaches is a result of the Bayesian rationality, as proved in Proposition 1 of Kamenica and Gentzkow (2011) and, among others, discussed in

[5]Compared to the standard perfect Bayesian equilibrium, by requiring the receiver to choose the sender-preferred action whenever indifferent, solution concept ensures the upper semicontinuity of the agent's indirect utility function and hence the selection of the equilibrium where the sender obtains the highest possible payoff.

[6]This set of signals is essentially a Blackwell experiment as defined in Blackwell (1953) and most recently Jakobsen (2018).

Shmaya and Yariv (2009).[7] We illustrate the equivalence for the following three types of information choices:

1. No information: If $supp(\mu)$ contains only one belief p, then the *Bayes plausibility* constraint implies that $p = p^*$ and $\mu(p^*) = 1$.[8] The agent's belief is not updated in either state, which is only possible if the signal is completely uninformative:

$$\pi(s|\overline{R}) = p^*, \forall s \in \mathcal{S}.$$

   We denote this μ as μ_n and interpret it as *no information*.

2. Full information: If $supp(\mu)$ contains two beliefs $p_1 = 0$ and $p_2 = 1$, then the *Bayes plausibility* constraint implies that $\mu(p_2; p^*) = p^*$. There is no uncertainty at either posterior belief, which is only possible if the signal is completely informative:

$$\exists s_1 \neq s_2 \in \mathcal{S}, s.t.\ \pi(s_1|\overline{R}) = 1, \text{and } \pi(s_2|\underline{R}) = 1.$$

   We denote this μ as μ_f and interpret it as *full information*.

3. Partial information: For all the other μ, the signal has to be designed in such a way that the agent is better but not completely informed in either state:

$$\exists R \in \Omega, s.t.\ \forall s \in \mathcal{S}, \pi(s|R) < 1.$$

[7] The key idea behind this equivalence is that each belief $p \in supp(\mu) \subseteq \Delta\Omega$ is essentially a posterior belief that is Bayesian updated from the prior belief p^* after some realization of the signal $s \in \mathcal{S}$. This observation then facilitates the construction of a one-to-one mapping between the signal realizations and the posterior beliefs using Bayes rule.

[8] Strictly speaking, the Bayes plausibility condition implies the dependence of the distribution function μ on the prior belief p^*. However, we use $\mu(p)$ instead of $\mu(p; p^*)$ throughout the paper for easy notation.

We collectively refer to all of them as *partial information*.

We impose two major assumptions on the informational environment. The first assumption is about the information symmetry between self-0 and self-1. Following our explanations above, whether we interpret the role of self-0 as a person or a benevolent government anticipating over-consumption behavior to arise in the future, we assume that self-1 has no access to additional sources of information when compared to self-0 and hence has no private types. In that sense, self-0 and self-1 are essentially the same person and self-0 is indeed choosing the information for self-1. The informational asymmetry between self-0 and self-1 also justifies the assumption of a common prior. This assumption implies that our model is not a signaling or cheap talk game by nature.

The second assumption is about the commitment power of self-0. Justifying this assumption is more of a challenge when we interpret self-0 as the agent himself rather than some benevolent government. We want to think of self-0's information choice as a costless but time-consuming process. Information is not be readily available but needs to be collected, for instance, from reading news articles during period $t = 0$. By the time self-1 needs to make the consumption saving decision, he could no longer change his information because he could neither afford the time to collect new information or forget what he has learned so far about the state.

There are two ways to interpret the role of self-0. She could be a person who is fully aware of her present bias and therefore correctly anticipates over-consumption behavior to arise in the future. Given this awareness, she would be interested in taking some proactive measure to alleviate the problem of over-consumption, but is limited in the tools she has available. Let us assume for now that signing up for a commitment device to save an arbitrary predefined amount

is not free and therefore she cannot afford it since her endowment does not arrive in $t = 0$.[9] In fact, in $t = 0$, her only affordable option is to customize the channels in the News app on her laptop, which we assume is all her access to the news. Some channels can be more informative on the future rate of return to savings than others. For example, The Wall Street Journal should be considered as more informative on this matter than The Washington Post. We want to consider a hypothetical list of infinitely many channels such that she could always find her ideal channel on the list.

Alternatively, she could be a benevolent planner who is aware of its subjects' present bias and considers setting up a news channel to help them out. For example, some universities provide faculty and staff free consultation with financial advisors[10]. Indeed, we could interpret any of self-0's optimal information choices as the suggestion of some financial advisor[11], together with the understanding that financial advisors differ in the forecasting models that they use. Again, we want to consider a hypothetical list of infinitely many financial advisors, each of whom uses his unique forecasting model, such that the university could always find a perfect advisor from the list for an employee.

We also have a few considerations regarding the payoffs of the game. We would like to introduce risk aversion and model it with the CRRA utility function, which is standard practice in the asset pricing literature . However, CRRA might cause a difficulty in comprehending how self-0's preference over distributions of posterior

[9] We might be able to justify this assumption for a present biased agent for the following two reasons: (1) the 401(k) contribution cannot be personalized; (2) it is nearly costless to rearrange any saving plan with a bank.

[10] See, for example, the arrangement at Duke from https://hr.duke.edu/benefits/retirement/investment-carriers, retrieved on January 14, 2019.

[11] We ignore any conflict of interests between financial advisors and clients.

beliefs is induced in our model. A key feature of the CRRA utility function is that the relative risk aversion also determines the elasticity of the intertemporal substitution with $EIS = \frac{1}{RRA}$. The relative risk aversion naturally plays a role in determining which information would be preferred as the environment is uncertain. The elasticity of intertemporal substitution also plays a role here, since the conflict of interests between self-0 and self-1 comes from the difference in discounting. Therefore, after establishing the benchmark linear case and extending the analysis to the standard von Neumann-Morgenstern expected utility formulation of CRRA, we will separate EIS from RRA using Epstein-Zin preferences. A caveat is that under those preferences the agent would demonstrate an intrinsic preference for the time of resolution of uncertainty, which may or may not interfere with self-0's induced preference over distributions of posterior beliefs as we will see in detail below.

3.3 Applying the Concavification Approach

Our model has a simple state space but a continuous action space, which differs from most models in the literature applying the concavification approach. In this section, we show how to simplify the representation of the concave closure in our model.

Anticipating self-1 choosing consumption optimally, self-0's optimization problem is:

$$\mu^* \in argmax_{\mu \in \mathcal{M}} \int_0^1 \beta\delta\{\frac{c_1^*(p)^\rho}{\rho} + \delta E_p[\frac{\big(R(1-c_1^*(p))\big)^\rho}{\rho}]\}\mu(dp), \tag{3.1}$$

$$s.t. \quad c_1^*(p) \in argmax_{c \in \mathcal{A}} \frac{c^\rho}{\rho} + \beta\delta E_p[\frac{\big(R(1-c)\big)^\rho}{\rho}], \tag{3.2}$$

where the integrand in equation 3.1 is her indirect utility function of p, denoted by $u_0(p)$, multiplied by the discount factor $\beta\delta$.

Simplifying $u_0(p)$ For $\rho \in (-\infty, 1)$, equation 3.2 has a unique interior solution, $c_1^*(p) \in (0,1)$, for any p.[12] Self-1's (gross) optimal consumption growth rate is then determined by:

$$t_1^*(p) \equiv \frac{1 - c_1^*(p)}{c_1^*(p)} = \left(\beta\delta E_p[R^\rho]\right)^{\frac{1}{1-\rho}}. \tag{3.3}$$

That $t_1^*(p)$ takes a simpler form is a result of our functional form assumption, as CRRA is a special case of constant elastic substitution (CES) functions. Given the equivalence between $t_1^*(p)$ and $\frac{1-c_1^*(p)}{c_1^*(p)}$ in equation 3.3, we proceed in the analysis with the choice of consumption growth rate, $t \in \overline{\mathbb{R}}_+$, as self-1's action. The action space which contains only actions that are attainable with some $p \in [0,1]$ becomes:

$$\widetilde{A} = [min\{(\beta\delta\underline{R}^\rho)^{\frac{1}{1-\rho}}, (\beta\delta\overline{R}^\rho)^{\frac{1}{1-\rho}}\}, max\{(\beta\delta\underline{R}^\rho)^{\frac{1}{1-\rho}}, (\beta\delta\overline{R}^\rho)^{\frac{1}{1-\rho}}\}]. \tag{3.4}$$

Substituting out $c_1^*(p)$ with $t_1^*(p)$, self-0's indirect utility function can then be simplified to:

$$u_0(t_1^*(p)) = \frac{(1 + t_1^*(p))^{-\rho}}{\rho}\left(1 + \beta^{-1}t_1^*(p)\right). \tag{3.5}$$

Simplifying $\mathcal{M}$ Many papers in the Bayesian persuasion literature have taken the *recommended action* interpretation of the signal realization and used the number of possible actions as a restriction on the number of posterior beliefs in the support of the distribution function μ, which is not helpful for our continuous action space. Instead, denote by $U_0(\cdot)$ the concave closure of the set:

$$\{u_0(t_1^*(p)) | p \in [0,1]\}. \tag{3.6}$$

[12] We discuss the special case of $\rho = 1$ in the following.

Since this set is one dimensional with $u_0(\cdot) \in \mathbb{R}$, Caratheodory's theorem implies that any point that lies in the convex hull of this set can be written as a convex combination of at most two points in the set. Specifically, any point on $U_0(\cdot)$ is either a point in $\{u_0(t_1^*(p))\}$ or a convex combination of two points in $\{u_0(t_1^*(p))\}$. This observation allows us to consider a subset of $\mathcal{M}$ containing only distributions over at most two posterior beliefs:

$$\mathcal{M} = \{\mu \in \Delta(\{p_1, p_2\}) | p_1 < p^* < p_2, \mu(p_1) = \frac{p_2 - p^*}{p_2 - p_1}\} \cup \{\delta_{p^*}\}. \tag{3.7}$$

3.4 Special Case: $\rho = 1$

In this section, we consider a benchmark model with linear utility function. For any consumption choice $c \in \mathcal{A}$ and any belief p on the high state $\overline{R}$, self-1's lifetime expected utility is:

$$U_1(c; p) = c + \beta\delta\{p\overline{R}(1-c) + (1-p)\underline{R}(1-c)\}.$$

Writing $E_p[R] \equiv p\overline{R} + (1-p)\underline{R}$ for the expected rate of return under belief p, we can simplify self-1's objective function to:

$$U_1(c; p) = c + \beta\delta E_p[R](1-c).$$

Similarly, self-0's lifetime expected utility evaluated at $t = 1$[13] is:

$$U_0(c; p) = c + \delta E_p[R](1-c).$$

[13] We have abused the notation slightly by using $U_0(c; p)$ to denote the expect lifetime utility, which is different from $U_0(\cdot)$ in the previous section that denotes the concave closure of $u_0(\cdot)$, the utility of a consumption in one period.

The linearity of the utility function captures both risk neutrality and perfect elasticity of intertemporal substitution. Given that the agent is risk neutral, information has only *instrumental* value in the sense that there is no value of information to self-0, unless the information that she would make available to self-1 would induce him to choose a different optimal action. Indeed, once we hold c fixed in self-0's objective function, the linearity of $E_p[R]$ in p implies that any provision of information in the form of a *Bayes plausible* distribution of posterior beliefs would not change the magnitude of $E_p[R]$ and hence would not affect self-0's expected utility. Therefore, the key implication of the risk neutrality assumption under von Neumann-Morgenstern expected utility preferences is that self-0's induced preference over information becomes completely driven by how useful that information is in aligning self-1's preferred actions with self-0's. This implication allows us to simply focus on the alignment of the preferred actions to figure out self-0's optimal information choice.

Given perfect EIS, each of self-0 and self-1 would prefer to either consume all of the endowment at $t = 1$ or save everything for $t = 2$. Denote self-0's and self-1's optimal consumptions for a belief p in $t = 1$ as $c_0^*(p)$ and $c_1^*(p)$ respectively. Since the marginal utility of saving is the discounted expected rate of return, which is $\delta E_p[R]$ for self-0 and $\beta\delta E_p[R]$ for self-1, and the marginal utility of consumption is simply 1, they would prefer to save only when the expected rate of return is high enough:

$$c_0^*(p) = \begin{cases} 1, & \text{if } E_p[R] < \frac{1}{\delta} \\ 0, & \text{if } E_p[R] \geqslant \frac{1}{\delta} \end{cases} \quad , \quad c_1^*(p) = \begin{cases} 1, & \text{if } E_p[R] < \frac{1}{\beta\delta} \\ 0, & \text{if } E_p[R] \geqslant \frac{1}{\beta\delta} \end{cases} \quad .$$

Therefore, perfect EIS allows us to focus without loss of generality on the binary

action space $\mathcal{A}_b = \{0, 1\}$, which is a subspace of the original continuous action space $\mathcal{A} = [0, 1]$. With both features of linear utility and binary action space, the analysis of this benchmark model is readily comparable with existing studies in the literature, including the prosecutor-judge example in Kamenica Gentzkow (2011), the confidence-maintenance section in Benabou Tirole (2002), and the three-period example in Carrilo Marriotti (2000).

Combining the implications of both the risk neutrality and the perfect EIS assumptions, we understand that our analysis could be separated into cases depending on self-0's and self-1's preferred actions, which depend on the parameters $\beta, \delta, \overline{R}$, and $\underline{R}$. Not all cases are interesting. In particular, if self-1's preferred action is the same across states, then his preferred action is independent of his belief p, which renders self-0's persuasion useless. The equilibrium is then trivial where self-0 prefers *no information* and self-1 plays a pure strategy, his state-independent preferred action. We summarize in the following lemma:

Lemma 3.4.1. *For any posterior belief $p \in [0, 1]$,*

1. *If $\overline{p}^1 \leqslant 0$, then self-1 is optimally saving everything for $t = 2$;*
2. *If $\overline{p}^1 > 1$, then self-1 is optimally consuming everything in $t = 1$.*

In either case, self-0 cannot influence self-1's action with the informational tool and therefore (weakly) prefers no information.

Therefore, for our analysis of self-0's information choice to be nontrivial, we need the parameters to be such that self-1 is *persuadable* with his preferred action being different across states. Denote the belief threshold where self-1 switches from consuming to saving as $\overline{p}^1$, which can be determined from the threshold for

$E_p[R]$ in the $c_1^*(p)$ function:

$$\overline{p}^1 \equiv \frac{\frac{1}{\beta\delta} - \underline{R}}{\overline{R} - \underline{R}}.$$

Without further calculations, we know that the present bias $\beta \in (0,1)$ implies that $\beta\delta E_p[R] < \delta E_p[R]$ for any belief p. Gence self-0's threshold for saving to be optimal is strictly smaller than self-1's, i.e. $\overline{p}^0 < \overline{p}^1$, where

$$\overline{p}^0 \equiv \frac{\frac{1}{\delta} - \underline{R}}{\overline{R} - \underline{R}}.$$

$\overline{p}^1 \in (0,1]$ then rules out the case where consuming everything is preferred by self-0 across states. Restricting the parameters to make the problem non-trivial, we are left with only two cases depending on whether self-0's preferred action is state independent or not.

Case 1: $\overline{p}^0 \leqslant 0 < \overline{p}^1 \leqslant 1$. The first case is when self-0 has a state-independent preferred action, which is saving everything. However, self-1's preferred action is state dependent, which is saving in the high state and consuming in the low state. This is the case where self-0's and self-1's preferred actions are partially aligned. If the common prior belief p^* is high enough such that $p^* \geqslant \overline{p}^1$, then self-1 would already be saving absent of self-0's information choice. It would be optimal for self-0 not to persuade self-1 by choosing *no information* μ_n, given that their preferred actions are already aligned.

If, instead, the prior belief p^* is not high enough, then self-1 would be consuming. In this scenario, the instrumental value of an information choice μ to self-0 comes from its use in persuading self-1 to save as often[14] as possible. For self-1

[14]This is a frequentist interpretation of the probability in the distribution of self-1's choice, not the probability of playing saving in self-1's strategy. Indeed, self-1 is always playing some pure

to save, he needs to be confident enough in the state being high, which requires a posterior belief p at least as high as $\overline{p}^1$. But because the prior belief p^* is too low, *Bayes plausibility* implies that it would be impossible for self-0 to persuade self-1 into saving all the time, since $p^* < \overline{p}^1$. In fact, *Bayes plausibility* requires that self-0 also needs to randomize over at least one posterior belief $p_1 < p^*$ to "balance out", which leads to self-1 consuming. If self-0 would like to consider randomizing over an additional posterior belief $\hat{p}$, this belief would have to induce self-1 to either consume or save. Since it cannot induce a different action, as we've discussed earlier in this section, the risk neutrality assumption and the von Neumann-Morgenstern formulation inherent in this linear case jointly imply that self-0 would be equally well off if we pool this belief $\hat{p}$ with either one of the two posterior beliefs depending on which action it induces.[15] Therefore, we can safely constrain ourselves to considering two and only two posterior beliefs in the support of self-0's optimal *Bayes plausible* distribution of beliefs μ.

That it is sufficient to consider μ that is only supported on at most two points for self-0's optimal information choice is an important step that we shouldn't take lightly. Formally, this means that we could without loss of generality search for self-0's optimal information on a subset $\mathcal{M}_2 \subseteq \mathcal{M}$, the set of all *Bayes plausible* belief distributions that we have introduced earlier:

$$\mathcal{M}_2 \equiv \{\mu \in \Delta(\Delta\Omega) | \int p d\mu(p; p^*) = p^*, \text{and } |supp(\mu)| \leqslant 2\},$$

strategy as we follow Kamenica Gentzkow (2011) in selecting the "sender optimal equilibrium".

[15] If, say, $\hat{p}$ induces self-1 to consume, then self-0 could alternatively consider randomizing over a belief $\tilde{p}$ which is a convex combination of $\hat{p}$ and p_1.

which set could be equivalently represented by a set of supports:

$$\tilde{\mathcal{M}}_2 \equiv \{\{p_1, p_2\} | p_1 \in [0, p^*], \text{and } p_2 \in [p^*, 1]\},$$

since the μ function can be uniquely pinned down by the *Bayes plausible* constraint with

$$\mu(p_1; p^*, p_2) = \frac{p_2 - p^*}{p_2 - p_1}, \quad \mu(p_2; p^*, p_1) = \frac{p^* - p_1}{p_2 - p_1}$$

when $|supp(\mu)| = 2$ and $\mu(p^*; p^*) = 1$ when $|supp(\mu)| = 1$. Self-0's optimization problem is reduced from infinite dimensional to two dimensional and becomes solvable with a simple linear programming approach.

This linear programming approach for Bayesian persuasion is formally introduced in Kolotilin (2018) and also discussed in Lipnowski Mathevet (2017) and Lipnowski and Mathevet (2018), who point out that we could without loss of generality search for the optimal information by focusing on the set of *outer point* posterior beliefs.[16] Further constraining the set of μ to measures supported by at most two points could also be justified following the geometric approach to solving Bayesian persuasion problems in Kamenica and Gentzkow (2011), as any point on the concave envelope of a graph, by definition, can be reached by the convex combination of two points on the graph.

Following this simplification, our next step is to determine which pair of $\{p_1, p_2\}$ would be optimal to self-0. If we pick p_2 such that $p_2 < \overline{p}^1$, then self-1 would still be consuming all the time even after self-0's information choice. Therefore, for

[16] Essentially, given the finiteness of the action space, this simplification should be no surprise after learning the revelation principle result for information design problems in Bergemann and Morris (2016), who consider the action space as the signal realization space. This "recommended action" interpretation with *straightforward* signals is also mentioned in Kamenica and Gentzkow (2011). As we've discussed earlier in this section, the finiteness of $\mathcal{S}$ when we look for the optimal information structure (S, π) should be equivalent of the finiteness of $supp(\mu)$ when we look for the optimal distribution of posterior beliefs μ.

self-0 to benefit from her information choice, we need $p_2 \in [\overline{p}^1, 1]$ to induce saving. While any p_2 in this range works, self-0 could only afford to randomize over higher p_2 with a smaller probability, since she needs to have the pair $\{p_1, p_2\}$ satisfy the *Bayes plausiblility* constraint. Her best option for p_2 then has to be the lowest one in that range, i.e. $p_2 = \overline{p}^1$. Similarly, any choice of p_1 would work the same in inducing self-1 to consume, since p_1 needs to fall on the left hand side of p^*, which is already too low for self-1 to switch to saving. Self-1 consuming is the action that is misaligned with self-0's preferred action, saving, and therefore self-0 would like to minimize its frequency by choosing the furthest possible p_1 from p^*, i.e. $p_1 = 0$.

Case 2: $0 < \overline{p}^0 < \overline{p}^1 \leqslant 1$. The second case is when self-0's preferred action is state dependent, where she prefers to consume in the low state and to save in the high state. Since self-1's preferred action is also state dependent, this is the case where self-0's and self-1's preferred actions are perfectly aligned. While the preferred actions are aligned, absent self-0's information choice, self-1 cannot have a state-contingent action plan. He is either saving or consuming all the time, depending on how the prior p^* compares with his threshold $\overline{p}^1$. Therefore, in this case, the instrumental value of information to self-0 comes from making self-1's action contingent on the state, and *full information* is therefore optimal. In particular, note that even if we get rid of the present bias by setting $\beta = 1$, information is still valuable to the agent for exactly the same reason, which is to have the agent choose the right action in either state.

We summarize the main insights from above in the following proposition:

Proposition 3.4.2. *Suppose* $0 < \overline{p}^1 \leqslant 1$. *Let* p^* *denote an arbitrary prior belief and* p *an arbitrary posterior belief. For* $\beta \in (0, 1)$, *in any equilibrium, self-1 chooses an optimal*

consumption $c_1^*(p)$, which is given by

$$c_1^*(p) = \begin{cases} 1, & \text{if } p < \overline{p}^1 \\ 0, & \text{if } p \geqslant \overline{p}^1 \end{cases} .$$

1. *If $\overline{p}^0 < 0$, then*

   (a) *when $p^* < \overline{p}^1$, the unique equilibrium is that self-0 chooses a* ***partial information*** *μ_p such that $supp(\mu_p) = \{0, \overline{p}^1\}$,*

   $$\mu_p(0; p^*, \overline{p}^1) = \frac{\overline{p}^1 - p^*}{\overline{p}^1}, \text{and } \mu_p(\overline{p}^1; p^*, 0) = \frac{p^*}{\overline{p}^1}.$$

   (b) *when $p^* \geqslant \overline{p}^1$, there are infinitely many equilibria where self-0 chooses a Bayes plausible belief distribution μ such that $supp(\mu) = \{p_1, p_2\} \subseteq [\overline{p}^1, 1]$ and the probabilities $\mu(p_1; p^*, p_2)$ and $\mu(p_2; p^*, p_1)$ are as given above. In particular, one equilibrium is that self-0 chooses the* ***no information*** *μ_n which is degenerate at p^*.*

2. *If $\overline{p}^0 > 0$, then the unique equilibrium is that self-0 chooses the* ***full information*** *μ_f such that $supp(\mu_f) = \{0, 1\}$,*

   $$\mu_f(0; p^*, 1) = 1 - p^*, \text{and } \mu_f(1; p^*, 0) = p^*.$$

3. *If $\overline{p}^0 = 0$, then any equilibrium described in either case above remains an equilibrium.*

The detailed proof is in the Appendix. We follow the concavification approach from Kamenica Gentzkow (2011) to solve for self-0's optimal information. Essentially, this approach allows us to only look at the concave envelop of self-0's indirect

utility function, which gives the highest attainable utility from any of her feasible information choices. At any prior belief p^*, this highest attainable utility has to be reachable from the convex combination of her utilities at some pair of posterior beliefs, and this pair becomes the support of self-0's optimal information choice μ.

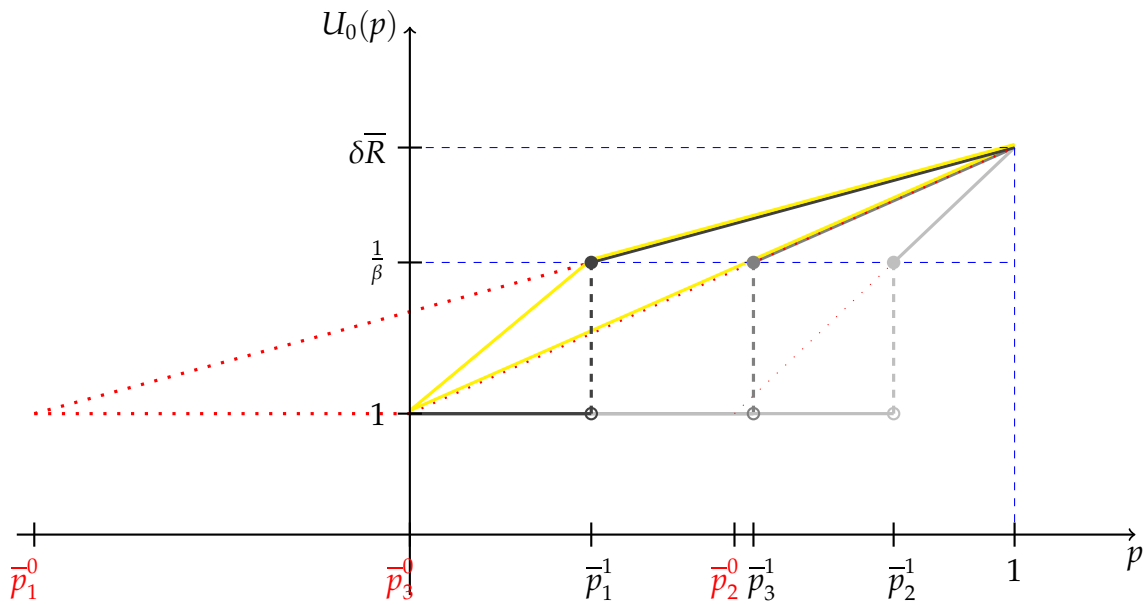

Figure 3.1: Self-0's indirect utility changes as $\overline{p}^0$ increases. The corresponding concave envelopes are colored yellow.

Following this approach, an important step is to determine the shape of the concave envelope. Recall that self-1's optimal consumption choice $c_1^*(p)$ is binary, and he switches from consuming to saving at the threshold $\overline{p}^1$, which implies that self-0's indirect utility function is constant on $[0, \overline{p}^1)$ and linear with an upward slope on $[\overline{p}^1, 1]$ with a discrete jump at $\overline{p}^1$. Therefore, the concave envelope of self-0's indirect utility function is either piecewise linear with a kink at $\overline{p}^1$ or simply linear everywhere[17], which corresponds to self-0 preferring either *no information* or *full information* at $\overline{p}^1$.

[17]In the language of Lipnowski Mathevet (2017) and Lipnowski Mathevet (2018), we have three *outer points* here: $0, \overline{p}^1, 1$.

When the prior belief is $\overline{p}^1$, we know from $c_1^*(p)$ that self-1 would be saving absent of self-0's information choice. This saving action is perfectly aligned with self-0's preferred action when $\overline{p}^0 < 0$ as in Case 1, and we know from our previous discussions that *no information* is optimal here. In this case, the concave envelope of self-0's indirect utility function should display a kink point at $\overline{p}^1$, which implies that self-0 should choose μ_p to randomize over the pair of posterior beliefs $\{0, \overline{p}^1\}$ when the prior belief $p^* \in (0, \overline{p}^1)$ and choose μ_n when $p^* \in (\overline{p}^1, 1)$. Similarly, we can tell that *full information* becomes optimal for its value in making self-1's action state contingent when $0 < \overline{p}^0 < 1$ as in CASE 2. This implies that the concave envelope should be linear everywhere and that self-0 should choose μ_f randomizing over $\{0, 1\}$ for any p^*.

Given the characterization of self-0's optimal choices of μ, we can then construct the optimal signals $\{\pi(\cdot|R) : \Omega \to \Delta\mathcal{S}\}$. The equilibrium can be alternatively characterized using the signals:

Corollary 3.4.3. *Suppose $\underline{R} < \frac{1}{\beta\delta} < \overline{R}$ and therefore $0 < \overline{p}^1 \leqslant 1$. Also suppose $\beta \in (0, 1)$. Let p^* denote an arbitrary prior belief. Consider the signal realization space $\mathcal{S} = \{\underline{r}, \overline{r}\}$.*

1. *If $\frac{1}{\delta} < \underline{R}$, then*

   (a) *when $p^* < \overline{p}^1$, the unique equilibrium is that self-0 chooses a* ***partial information*** *set of signals $\{\pi_p(\cdot|R) : \Omega \to \Delta\mathcal{S}\}$ such that:*

$$\pi_p(\overline{r}|\overline{R}) = 1, \textit{and } \pi_p(\overline{r}|\underline{R}) = \frac{p^*}{1-p^*}\frac{1-\overline{p}^1}{\overline{p}^1},$$

and self-1 chooses an optimal consumption $C(r)$ for $r \in \mathcal{S}$ such that:

$$C(r) = \begin{cases} 0, & \text{if } r = \overline{r} \\ 1, & \text{if } r = \underline{r} \end{cases}.$$

(b) *when $p^* \geqslant \overline{p}^1$, there are infinitely many equilibria where self-0 chooses a set of signals $\{\pi(\cdot|R)\}$ such that:*

$$\frac{1 - \pi(\overline{r}|\underline{R})}{1 - \pi(\overline{r}|\overline{R})} \leqslant \frac{p^*}{1 - p^*} \frac{1 - \overline{p}^1}{\overline{p}^1},$$

and self-1 consumes nothing (saves):

$$C_s(r) = 0, \forall r \in \mathcal{S}.$$

In particular, one equilibrium is that self-0 chooses the ***no information*** *signals $\{\pi_n(\cdot|R)\}$ such that:*

$$\pi_n(\overline{r}|\overline{R}) = \pi_n(\overline{r}|\underline{R}) = p^*.$$

2. *If $\frac{1}{\delta} > \underline{R}$, then the unique equilibrium is that self-0 chooses the* ***full information*** *signals $\{\pi_f(\cdot|R)\}$ such that:*

$$\pi_f(\overline{r}|\overline{R}) = \pi_f(\underline{r}|\underline{R}) = 1,$$

and self-1 chooses the same optimal consumption $C(r)$ as that given in $1(a)$.

3. *If $\frac{1}{\delta} = \underline{R}$, then any equilibrium described in either case above remains an equilibrium.*

The condition $\frac{1}{\delta} \lesseqgtr \underline{R}$ that separates the first two cases in Corollary 3.4.3 corresponds to the condition $\overline{p}^0 \lesseqgtr 0$ in Proposition 3.4.2 and is derived similarly from

the tradeoff at $\overline{p}^1$. We can think about this tradeoff in terms of the signals by considering the special case where $p^* = \overline{p}^1$. We can represent any arbitrary signal structure via $\pi(\overline{r}|\overline{R}) = \pi_1$ and $\pi(\overline{r}|\underline{R}) = \pi_2$. Further, $p^* = \overline{p}^1$ implies that we can without loss of generality restrict our attention to $\pi_1 \in [\overline{p}^1, 1]$ and $\pi_2 \in [0, \overline{p}^1]$. Any more informative signal structure would induce self-1 to be more confident about the state being high than his prior belief and hence to save after observing a high signal realization $\overline{r}$. Similarly, any more informative signal structure would also induce self-1 to be less confident about the state being high than his prior belief and hence to consume after observing a low signal realization $\underline{r}$.

If self-0 marginally increases π_1, she is reallocating probability for $\underline{r}$ to $\overline{r}$ in the high state, which causes a net change of $\delta\overline{R} - 1$ in her payoff from the high state. By our assumption, this net change is always positive, and hence she would optimally set $\pi_1 = 1$ and always truthfully reveal the high state. Similarly, if she marginally increases π_2, then there is a net change of $\delta\underline{R} - 1$ in her payoff from the low state. She would optimally set $\pi_2 = 0$ if this change is negative, i.e. $\frac{1}{\delta} > \underline{R}$, which leads to the fully informative signal structure being optimal. She would set π_2 as high as she can if the change is instead positive, $\pi_2 = \overline{p}^1$. But then she could set π_1 to $\overline{p}^1$ without incurring any loss from the high state while increasing her payoff from the low state, since the signal structure would become completely uninformative with $\pi_1 = \pi_2 = \overline{p}^1$ and self-1 would be saving even after observing the low sate signal realization $\underline{r}$.

We now offer some intuitive understanding of the *partial information* result in the Case 1a. As we've discussed above, this is the case where self-0's and self-1's preferred actions are partially aligned. The goal for self-0 is to persuade self-1 to save as often as possible, which is equivalent to inducing self-1 to have a belief

$p \geqslant \overline{p}^1$ as often as possible. To achieve this goal, when the true state is the high state $\overline{R}$, self-0 would certainly like to reveal it truthfully to self-1; when the true state is $\underline{R}$, self-0 would also like to send a high state message to self-1 as often as she can. Specifically, the more often she is sending the false high state message in the low state, the more often self-1 is saving. But self-0 couldn't afford sending the false message too often, since increasing the frequency of sending this false message also reduces self-1's confidence in believing that the state is indeed high when observing a high state message. Taking this to the extreme, if self-0 always sends a false high state message when the true state is low, self-1 would know that the signal is useless and his confidence in the high state would be as low as his prior belief p^*, and he would be optimally choosing to consume. Therefore, she can send the false high state message in the low state at most as often as $\frac{p^*}{1-p^*}\frac{1-\overline{p}^1}{\overline{p}^1}$, which induces self-1 to be just confident enough to save after observing the high state signal $\overline{r}$.

3.5 Main Result

Without further restrictions, the distribution of posterior beliefs μ is supported on $[0,1]$. In our model, any p induces a different action $c_1^*(p)$, and our action space is necessarily continuous. In the limit case of $\rho = 1$, this optimal consumption ratio $t_1^*(p)$ is either 0 or ∞, corresponding to $c_1^*(p)$ being either 1 or 0, depending on whether the belief p is above or below self-1's belief threshold defined for the benchmark case in the previous section, which is $\overline{p}^1$. For any $\rho \in (-\infty, 1)$, the optimal consumption $c_1^*(p)$ is an interior solution within $(0,1)$ and hence $t_1^*(p) \in (0,\infty)$. Instead of having self-0's lifetime expected utility function defined as a

function of c, we could rearrange it as a function of the consumption ratio t:

$$U_0(t;p) = \frac{(1+t)^{-\rho}}{\rho}\{1 + \delta E_p[R^\rho]t^\rho\},$$

from which we can further derive self-0's indirect utility from any posterior belief p:

$$U_0(t_1^*(p);p) = \frac{(1+t_1^*(p))^{-\rho}}{\rho}\{1 + \beta^{-1}t_1^*(p)\}.$$

To apply the concavification approach, we need to have an understanding of the curvature of this indirect utility function. It turns out that its curvature critically depends on the comparison of ρ and β. Specifically, for any β, if $\rho \leqslant \beta$, then this indirect utility function is globally convex independent of all the other parameters, $\delta, \overline{R}$, and $\underline{R}$. In the other case where $\rho > \beta$, these other parameters start to matter in the sense that, if $\overline{R}$ is large enough and $\underline{R}$ small enough such that there exists a unique inflection point p where $t_1^*(p) = \frac{1-\beta\rho}{\rho-\beta}$, then the function starts out being convex from $p = 0$ until it hits this inflection point and becomes concave all the way until $p = 1$. Otherwise, if either the high state $\overline{R}$ is too low or the low state $\underline{R}$ is too high, there is no inflection point and the function is either globally convex or concave.

At the inflection point, if it exists, we have:

$$[\beta\delta E_p[R^\rho]]^{\frac{1}{1-\rho}} = \frac{1-\beta\rho}{\rho-\beta}.$$

This equation holds for all $\rho < 1$ but might break down in the limit as we take $\rho \to 1$. In fact, it holds in the limit if and only if $\beta\delta E_p[R] = 1$. Connecting back to our analysis of the benchmark, we realize this is exactly the equation that gives us self-1's belief threshold, $\overline{p}^1$, where he switches from consuming to saving. Indeed, this

inflection point corresponds to the discontinuity point in the linear case. We've learnt from Lemma 1 that the only interesting case is when there exists such a discontinuity point for some $p \in (0,1)$. Moreover, we also learnt from Corollary 1 that, in this interesting case, whether full information or partial information[18] is optimal critically depends on the low state $\underline{R}$. Specifically, we characterized a threshold, $\frac{1}{\delta}$, such that full (partial) information is optimal if $\underline{R}$ is below (above) this threshold. Not surprisingly, our result on the optimal information for the current case parallels that from the benchmark linear case:

Proposition 3.5.1. *Consider $\underline{R}, \overline{R} \in \mathbb{R}_+$ such that $\underline{R} < \overline{R}$. If $\rho \leqslant \beta$, full information is always optimal. If $\rho > \beta$, for any $\overline{R}$ large enough such that $\overline{R} > R_0$, where*

$$R_0 \equiv [\frac{1}{\beta\delta}(\frac{1-\beta\rho}{\rho-\beta})^{1-\rho}]^{\frac{1}{\rho}},$$

there exists a unique $\underline{R}^*$ such that:

1. *if $\underline{R} \leqslant \underline{R}^*$, then full information is optimal;*
2. *if $\underline{R}^* < \underline{R} < R_0$, then there exists $\hat{p} \in (0,1)$ such that, when the prior belief $p^* \in (0, \hat{p})$, partial information supported on $\{0, \hat{p}\}$ is optimal;*
3. *if $\underline{R} \geqslant R_0$, or $\underline{R}^* < \underline{R} < R_0$ with a prior belief $p^* \in [\hat{p}, 1]$, then no information is optimal;*

otherwise, if $\overline{R} \leqslant R_0$, full information is again optimal.

This proposition tells us that, as we take ρ away from 1, we get to maintain the characterization of the optimal information, but only within a small enough

[18] Here, by partial information, we are referring to the case where partial information is optimal for low prior belief and no information is optimal for high prior belief.

neighborhood bounded by β. Once ρ becomes too different from 1, full information becomes optimal regardless of the values of all the other parameters. The natural question to ask is: why does it matter how ρ compares to β? We will soon realize that there is no easy answer to this simple question, because ρ plays two roles. As we decrease ρ from 1, we increase the agent's aversion to risk as well as decreasing the agent's aversion to consumption smoothing. We might come to a conjecture that, given that we are comparing ρ to β which itself is about across-time tradeoff, the change in EIS might be the dominant driving force. But because both across-state and across-time tradeoffs are affected at the same time, there is no way to understand the mechanism behind the sharp change in result at $\rho = \beta$ unless we can separate the two roles of ρ.

3.6 Extension: EZKP preferences

Our motivation for disentangling the elasticity of intertemporal substitution (EIS) and risk aversion (RRA) arises from the fact that the effect of self-0's informational tool in influencing self-1's action through these two channels are otherwise intertwined. This lack of separation has been acknowledged as a key critique of the expected utility formulation with CRRA utility. It is particularly problematic in our context. EIS captures the agent's preference across time, which is necessarily confounded with β which is a discount factor and hence also operates across time. RRA, in contrast, captures the agent's preference across states. Since uncertainty only applies to $t = 2$, β does not interact with RRA directly. Adopting the Epstein-Zin separation of preferences will give us enough flexibility to better comprehend the roles of EIS and RRA in the optimality of self-1's information choice.

We follow the standard approach, the Epstein-Zin formulation developed in Epstein Zin (1989), to disentangle EIS and RRA. The parameter $\alpha \in (-\infty, 1)$ is used to capture RRA with $RRA = 1 - \alpha$, while $\rho \in (-\infty, 1)$ is used to capture EIS with $EIS = \frac{1}{1-\rho}$. The essential idea behind the Epstein-Zin formulation is that we take care of the uncertainty in each period by considering the certainty equivalent consumption $CE_\alpha[c;p]$ and then evaluate the path of certainty equivalent consumptions using a different utility function that is parametrized by ρ. While the general formulation is recursive, we need to adapt it to our setting which is essentially static due to the quasi-hyperbolic agent making a consumption saving decision only in $t = 1$.

We want to derive self-0's and self-1's objective functions by first considering their certainty equivalent consumptions in $t = 2$. Since the uncertain R is multiplicative in our model, the certainty equivalent consumption is:

$$CE_\alpha[c;p] = (E_p[R^\alpha])^{\frac{1}{\alpha}}(1-c) = CE_\alpha[R;p](1-c),$$

where $CE_\alpha[R;p]$ denotes the certainty equivalent rate of return.[19] Evaluating the certainty equivalent consumption from each period using $u(c) \equiv \frac{c^\rho}{\rho}$, we get self-1's lifetime utility under the Epstein-Zin formulation:

$$U_1(c;p) = \frac{c^\rho}{\rho} + \beta\delta\frac{[CE_\alpha[R;p](1-c)]^\rho}{\rho},$$

[19] For any consumption choice c and belief p, the expected utility from the uncertain consumption $t = 2$ is:

$$E_p[v(c)] = p\frac{[\overline{R}(1-c)]^\alpha}{\alpha} + (1-p)\frac{[\underline{R}(1-c)]^\alpha}{\alpha},$$

where $v(c) \equiv \frac{c^\alpha}{\alpha}$ captures the risk aversion.

which can be rearranged as:

$$U_1(c;p) = \frac{c^\rho}{\rho} + \beta\delta CE_\alpha[R;p]^\rho \frac{(1-c)^\rho}{\rho}.$$

Similarly, self-0's lifetime utility evaluated at $t = 1$ becomes:

$$U_0(c;p) = \frac{c^\rho}{\rho} + \delta CE_\alpha[R;p]^\rho \frac{(1-c)^\rho}{\rho}.$$

If we set $\rho = \alpha$, then we are back to standard CRRA expected utility case with $CE_\alpha[R;p]^\alpha = E_p[R^\alpha]$, which we have analyzed in the previous section; if we additionally set them equal to 1, then we are back to our benchmark model featuring linear utilities as $CE_1[R;p]^1 = E_p[R]$.

We can solve for self-0's and self-1's optimal actions from these $t = 1$ utility functions. To better understand how the misalignment in preference induces the misalignment in preferred actions, instead of comparing optimal consumptions, we compare optimal consumption ratios:

$$t_i^*(p) \equiv \frac{1 - c_i^*(p)}{c_i^*(p)}.$$

One way to measure the misalignment in actions is to consider the ratio of their respective $t_i^*(p)$, from which we can see that only the elasticity of intertemporal substitution matters:

$$\frac{t_1^*(p)}{t_0^*(p)} = \beta^{\frac{1}{1-\rho}}.$$

To better use this separation of EIS and RRA to bridge the gap in intuition between the benchmark linear case and the standard CRRA expected utility case, we first shut down the EIS channel by setting $\rho = 1$ and allow RRA to be completely

flexible. What we have seen in the above equation is that EIS plays a key role in determining the misalignment which is essential for the relevance of informational tools in alleviating insufficient saving caused by present bias. As we will see below, for the special case where $\rho = 1$, it turns out that risk aversion has no impact on information choice.

One caveat in using the Epstein-Zin formulation is that the agent would necessarily demonstrate some intrinsic preference for the time of resolution of uncertainty, as proved in Kreps Porteus (1978) and re-stated in Epstein Zin (1989). Specifically, we would expect a preference for early resolution of uncertainty when $\rho > \alpha$, and a preference for late resolution of uncertainty for $\rho < \alpha$. This inherent preference for information may or may not interfere with self-0's (instrumental) induced preference for information as we will see below.

3.6.1 Perfect elasticity of intertemporal substitution & flexible risk aversion

Setting $\rho = 1$ in the above utility functions to accommodate perfect elasticity of intertemporal substitution, self-1's lifetime utility becomes:

$$U_1(c;p) = c + \beta\delta CE_\alpha[R;p](1-c),$$

and self-0's lifetime utility evaluated at $t = 1$ becomes:

$$U_0(c;p) = c + \delta CE_\alpha[R;p](1-c).$$

Parallel to the linear case considered in section 3, perfect EIS implies that we can consider the binary action space $\mathcal{A}_b = \{0, 1\}$ without loss of generality, and

that both self-0 and self-1 compare their discounted certainty equivalent rate of return with 1, the return from consuming, to decide on their respective optimal consumption:

$$c_0^*(p) = \begin{cases} 1, & \text{if } CE_\alpha[R;p] < \frac{1}{\delta} \\ 0, & \text{if } CE_\alpha[R;p] \geqslant \frac{1}{\delta} \end{cases}, \quad c_1^*(p) = \begin{cases} 1, & \text{if } CE_\alpha[R;p] < \frac{1}{\beta\delta} \\ 0, & \text{if } CE_\alpha[R;p] \geqslant \frac{1}{\beta\delta} \end{cases}. \tag{3.8}$$

The belief thresholds now become functions of α:

$$\overline{p}^1(\alpha) = \frac{(\frac{1}{\beta\delta})^\alpha - \underline{R}^\alpha}{\overline{R}^\alpha - \underline{R}^\alpha}, \quad \overline{p}^0(\alpha) = \frac{(\frac{1}{\delta})^\alpha - \underline{R}^\alpha}{\overline{R}^\alpha - \underline{R}^\alpha}.$$

We maintain the assumption of $0 < \overline{p}^1(\alpha) < 1$ to guarantee a nontrivial information choice. We separate the discussion into the two distinctive cases (1) $\overline{p}^0(\alpha) < 0$ and (2) $\overline{p}^0(\alpha) > 0$, where self-0's preferred action differs in terms of its state dependency.[20] The departure from the benchmark lies in the certainty equivalent rate of return, which now varies in risk aversion.

Under the Epstein-Zin formulation, as we've shown above, self-1's problem for optimal consumption should be similarly analyzed as before, since his utility function has a similar functional form. Self-0's problem for optimal information, however, could be very different. We have to be careful in following Kamenica Gentzkow (2011) and applying their concavification approach, since we have adopted the Epstein-Zin formulation instead of the vNM formulation which they used in establishing the validity of this approach in their Proposition 1. In particular, self-0's utility from some pair of posterior beliefs $\{p_1, p_2\}$ being equal to the

[20] To determine the preferred action in each state, there is no uncertainty and hence risk aversion doesn't matter.

convex combination of the corresponding utilities, $U_0(c^*(p_1);p_1)$ and $U_0(c^*(p_2);p_2)$, relies on the linearity of expected utility in the probabilities. Whether the concavification approach still applies depends on whether this linearity property is maintained.

Consider self-0's objective function in $t = 0$. Denote her utility from a belief p at $t = 1$ by $U_0^1(c;p)$. At $t = 0$, there is no consumption and hence no flow utility in that period, but the consumption to be chosen at $t = 1$ by self-1 becomes stochastic to self-0, since it depends on the posterior belief which itself is random. Therefore, we can no longer take her expected utility in $t = 1$ as the objective function but have to go through another layer of certainty equivalence[21] to accommodate the Epstein-Zin formulation:

$$U_0^0(p_1,p_2;p^*,c^*(\cdot),\alpha) = \left\{\mu(p_1;p^*,p_2)U_0^1(c^*(p_1);p_1)^\alpha + (1-\mu(p_1;p^*,p_2))U_0^1(c^*(p_2);p_2)^\alpha\right\}^{\frac{1}{\alpha}},$$

which is self-0's utility evaluated at $t = 0$ scaled by a factor of $\frac{1}{\beta\delta}$. This is no longer a convex combination of the corresponding utilities and hence we cannot apply the concavification approach directly.

To circumvent this difficulty in applying the concavification approach, we consider a monotone transformation of the utilities:

$$f(x) = \frac{x^\alpha - 1}{\alpha},$$

[21] Technically, we need to first get the corresponding consumptions for $U_0^1(c^*(p_i);p_i)$ from the inverse function of $u(c) = \frac{c^\rho}{\rho}$, and then get the certainty equivalent consumption before finally evaluating it using $u(c)$. Here, since $\rho = 1$, $u(c) = c$ is an identity function, which is why we can take the certainty equivalent of $U_0^1(c^*(p_i);p_i)$ directly.

so that self-0's objective function becomes the *expected transformed* utility[22]:

$$U_0^T(p_1, p_2; p^*, c^*(\cdot), \alpha) = \mu(p_1; p^*, p_2) U_0^T(c^*(p_1); p_1) + (1 - \mu(p_1; p^*, p_2)) U_0^T(c^*(p_2); p_2),$$

where

$$U_0^T(c^*(p); p) \equiv \frac{U_0^1(c^*(p); p)^\alpha - 1}{\alpha} = \begin{cases} 0 & \text{, if } p \in [0, \overline{p}^1(\alpha)) \\ \frac{\delta^\alpha E_p[R^\alpha] - 1}{\alpha} & \text{, if } p \in [\overline{p}^1(\alpha), 1] \end{cases}, \forall \alpha \in (-\infty, 1].$$

Now that the transformed indirect utility function restores linearity in p and the transformed objective function becomes a convex combination of the transformed utilities, we can apply the concavification approach.

Essentially, this monotonic transformation implies that, as a result of perfect EIS, the current Epstein-Zin formulation with CRRA utility function on the state space Ω with discount factor δ is equivalent to the vNM expected utility formulation with linear utility function on the *transformed* state space $\Omega(\alpha) = \{\overline{R}^\alpha, \underline{R}^\alpha\}$ with the transformed discount factor δ^α. We want to emphasize the implication that the separation of our cases is independent of α in the sense that, for any α:

$$0 < \overline{p}^1(\alpha) < 1 \Leftrightarrow 0 < \overline{p}^1 < 1, \quad \overline{p}^0(\alpha) \lesseqgtr 0 \Leftrightarrow \overline{p}^0 \lesseqgtr 0.$$

Therefore, the characterization of the optimal information should resemble that for the linear case in Proposition 3.4.2.

That the "form" of the optimal information is independent of risk aversion is

[22]This convenience of transforming self-0's problem to a maximization problem of the expected value of the image of $t = 1$ stochastic utilities is a result of $f(\cdot)$ being a monotonic transformation and is independent of $t = 0$ flow utility being zero or not.

surprising, since the agent should demonstrate a preference for early resolution of uncertainty due to the Epstein-Zin formulation with $1 = \rho > \alpha$. But this is because there is another necessary condition for the agent to display this time preference, which is the existence of a nonzero deterministic consumption in the earlier time period. This is usually implicitly assumed, but is violated here. To see how it is violated, let us consider an example where self-0 consumes c in $t = 1$ and 1 in $t = 2$. If the uncertainty is resolved in $t = 2$, then self-0's utility in $t = 1$ is:

$$U_0^{t=2}(c, 1; p) = c + \delta CE_\alpha[R; p].$$

If, instead, the uncertainty is resolved in $t = 1$, then self-0's utility in $t = 1$ becomes:

$$U_0^{t=1}(c, 1; p) = CE_\alpha[c + \delta R; p].$$

We see that nonzero deterministic c in the earlier period is essential in inducing a preference for early resolution: we would have $U_0^{t=1}(c, 1; p) > U_0^{t=2}(c, 1; p)$ unless $c = 0$, in which csae $U_0^{t=1}(0, 1; p) = U_0^{t=2}(0, 1; p)$. In our model, while consuming $c = 0$ in $t = 1$ and 1 in $t = 2$ is a possible consumption path, we have shown in equation 3.8 that the only other possible consumption path is consuming $c = 1$ in $t = 1$ and 0 in $t = 2$ with $U_0^{t=1}(1, 0; p) > U_0^{t=2}(1, 0; p) = c$. In this case, the time of resolution of uncertainty also does not matter but for a different reason: there is no consumption and hence no uncertainty to resolve in $t = 2$.[23]

Since self-0 has no intrinsic preference for information, information is only of instrumental value to her, which, as we've discussed in section 3, comes from mak-

[23] Were there nonzero consumption in $t = 2$, we would restore self-0's intrinsic preference for information. If, due to some technical or institutional reason, the agent's maximal saving is bounded below 1 and hence self-1 has to incur a nonzero flow utility in $t = 1$, then self-0 will have a strict preference over the time of resolution of uncertainty, which would then influence the form of her optimal information.

ing self-1's actions depend on the state when she has a state-dependent preferred action and from persuading self-1 to save as often as she could when she prefers saving all the time. We have discussed earlier that risk aversion is irrelevant in determining self-0's or self-1's preferred action in either state. Therefore, the optimal information should take the same form as in the benchmark:

Proposition 3.6.1. *Suppose* $0 < \overline{p}^1 \leqslant 1$. *For* $\beta \in (0,1)$, *given arbitrary prior belief* p^*, *for any* $\alpha \in (-\infty, 1)$, *the equilibrium is exactly the same as that characterized in Proposition 1, with the only difference being that self-1's belief threshold* $\overline{p}^1$ *is replaced by the generalized* $\overline{p}^1(\alpha)$.

The *full information* equilibrium in the case of $\overline{p}^0 > 0$ is exactly the same as before, but the *partial* and *no information* equilibrium in the case of $\overline{p}^0 < 0$ needs to be adapted with the belief threshold $\overline{p}^1$ now changing in α to accommodate the flexible risk aversion. Note that the more risk averse the agent is, the smaller the certainty equivalent rate of return. Therefore, a more risk aversion self-1 with a smaller α is more difficult to persuade, since he would need to be more confident in the high state to be willing to save. If the agent is so risk averse that α is close to $-\infty$, self-1 would need to be almost certain in the high state to be willing to take on the risk and save. If, to the other extreme, the agent is only slightly risk averse that α is close to 1, self-1 should behave in a way closely resemble that of a risk neutral self-1. This gives rise to a conjecture on the monotone property of $\overline{p}^1(\alpha)$ in α.

Lemma 3.6.2. $\overline{p}^1(\alpha)$ *monotonically decreases in* α *on* $(-\infty, 1)$, *and it converges to 1 as* $\rho \to -\infty$ *and* $\frac{\frac{1}{\beta\delta} - \underline{R}}{\overline{R} - \underline{R}}$ *as* $\alpha \to 1$.

The major take-away from this section is that risk aversion has no influence on the form of the optimal information. The key reason behind this "irrelevance" of risk aversion in the persuasion problem that we've shown in this section is that

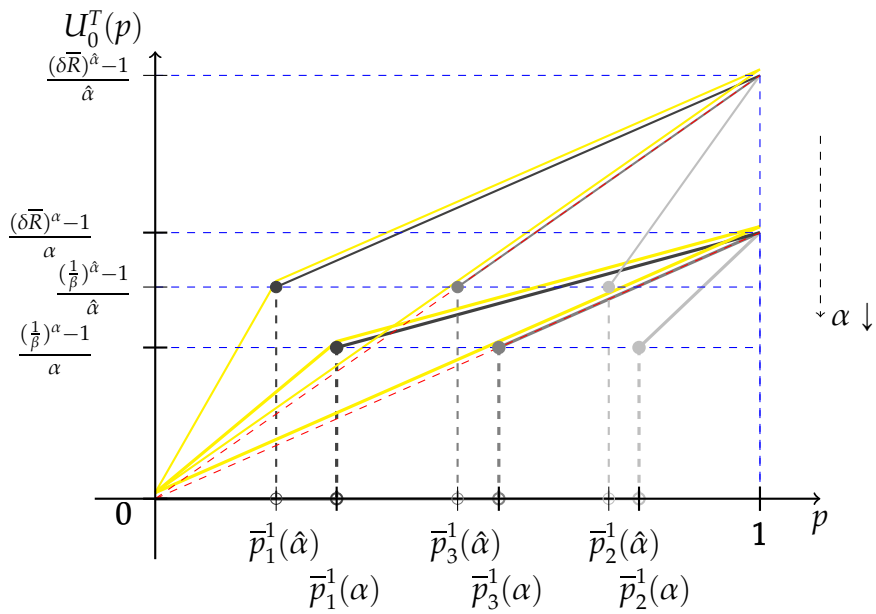

Figure 3.2: An illustration of the impact of changing α with $0 < \alpha < \hat{\alpha} < 1$. Fixed parameters $\beta, \delta, \overline{R}$,and $\underline{R}$.

whether self-0 prefers the same or different actions across the states is independent of the risk aversion. Essentially, when we consider self-0's optimal action in either state, there is no uncertainty involved and hence risk aversion has no role to play. But risk aversion does play a role in widening the region of p^* at which *partial information* is optimal, in the case where self-0 prefers to save independent of the state. This is because a more risk averse self-1 tends to be less confident in the high state and is therefore less likely to save unless persuaded, otherwise.[24]

3.7 Conclusion

In this paper, we study the use of informational tools in alleviating the dynamic inconsistency problem caused by present bias by formulating the problem as an information design problem and solving it using the concavification approach

[24]This an alternative intuition for the impact of α on the *partial information* equilibrium.

introduced in earlier papers in the literature. In our application, the misalignment in the sender and receiver's preferences is a result of present bias. We have characterized the optimal information choice when the agent is either risk neutral or risk averse. The risk aversion case is a slight departure from the standard examples we have seen in the information design literature where the receiver's action is essentially binary. For binary actions, partial information could be optimal when there is a discrete jump in the sender's utility as a function of the receiver's belief. In contrsat, in our model, the action space is continuous, and the role of the discrete jump is replaced by the inflection point of the sender's utility function.

Other than this extension of the standard result for binary action space to our example with a continuous action space, we have also discovered an interesting interaction of risk attitude and present bias. In our main model, risk attitude is captured by the CRRA utility function, which separates risk aversion from elasticity of intertemporal substitution. We consider a special case of the Epstein-Zin model and show that in that model risk aversion has no impact on the form of the optimal information choice.

Chapter 4

Information Design and Risk Dominance

4.1 Introduction

The success of a project is often impossible without the joint effort of all relevant parties. Developing an area needs the investments on both the real estate and the local infrastructure. Promoting electric cars in a new market needs the investment on both the charging stations and the electric car sales and services. Creating a new product in a tech company needs the input from both the engineering and the business teams. While the project could be hugely profitable with joint investment, each party may not be confident enough in the prospect of the project and hence reluctant to invest unless they know the other party has committed to investing. For the local government that wants to develop an area or the tech company that wants to introduce a new product, one way to induce joint investments is to make all parties become sufficiently confident that the investment is going to be profitable, by convincing them in the likely prospect of the project being profitable and their partners also investing.

The study of informational tools for inducing joint effort is made possible by the information design literature which extends the Bayesian persuasion problem pioneered by Aumann and Maschler (1995) and Kamenica and Gentzkow (2011) to a setting where the agents being persuaded interact strategically. While the general properties that need to be satisfied by the optimal information structure have been outlined in Bergemann and Morris (2016) and Mathevet et al. (2020).

The construction of the optimal information structure has been ad hoc in applications, which limits the usefulness of the information design framework in applied work. This chapter paves the way for providing a detailed construction of optimal information structure, which would be handy for applied researchers.

Two prominent constructions are pooling as in Li et al. (2019) and a specific construction which completely separates underlying uncertainty and strategic uncertainty introduced by Hoshino (2017). In this short paper, we find examples to illustrate their relative strength.

We first introduce the general setup for this type of information design in strategic problems, and present the two available constructions. After that, we show via two examples that the two constructions cannot be globally ranked in terms of their effectiveness. For practitioners this means that they should carefully compare both constructions in the context of the specific strategic persuasion problem they face.

4.2 Model

Two agents play an incomplete information game by choosing an action $a_i \in A = \{H, L\}$ for $i \in I = \{1, 2\}$. Their payoff is determined by their actions as well as an underlying state $\omega \in \Omega \subseteq \mathbb{R}$ which is captured by the utility function $u_i(a_i, a_j; \omega)$. The utility function has increasing differences which means that each agent benefits from a higher action of the opponent as well as a higher underlying state. We highlight three thresholds on ω corresponding to strict dominance,

strictly dominated, and risk dominance criteria:

$$\omega_i^{SD} = \inf_{u_i(H,a_j) \geqslant u_i(L,a_j), \forall a_j} \omega, \tag{4.1}$$

$$\omega_i^{SDd} = \sup_{u_i(H,a_j) \leqslant u_i(L,a_j), \forall a_j} \omega, \tag{4.2}$$

$$\omega^{RD} = \inf_{(u_1(H,H)-u_1(L,H))(u_2(H,H)-u_2(L,H)) \geqslant (u_1(L,L)-u_1(H,L))(u_2(L,L)-u_2(H,L))} \omega. \tag{4.3}$$

Given these thresholds, we know that at any state $\omega \geqslant \omega_i^{SD}$, agent i plays the pure strategy H; at any state $\omega \leqslant \omega_i^{SDd}$, agent i plays the pure strategy L; and at any state in the middle, their strategy depends on their belief on the opponent's strategy.

The information designer shares common prior $\mu_0 \in \Delta(\Omega)$ with the agents and can commit to a Bayes plausible information structure which determines the informational environment for the agents before they play. Specifically, the designer chooses a signal space S and a mapping from the underlying state to the distribution of signals $\pi : \Omega \to \Delta(S)$. Given the correlation of the signals, when observing the signals, the agents update their belief on the underlying state as well as the opponent's strategy and choose a strategy that maximizes their expected payoff $E[u_i(a_i, a_j; \omega)|s_i]$. In case there are multiple equilibria, the designer cares about the worst case scenario and selects the outcome with the lowest expected payoff for the agents. The designer prefers the outcome (H, H) independent of the state and therefore chooses (S, π) that maximizes the probability of (H, H) being the worst outcome which is then also the unique outcome given that (H, H) is the best outcome for the agents.

Two prominent information structures in the literature are *pooling* and *Hoshino (2017)*. Both constructions feature correlated signals which leverage the higher order beliefs to gradually increase the probability of agents playing the desirable action H. Their difference rests in the treatment of the underlying uncertainty and

strategic uncertainty. The intertwinement of these two uncertainties in solving incomplete information games has been discussed in Liu (2015). Briefly, a unique feature of the Hoshino (2017) construction is its sole reliance on strategic uncertainty after constructing the first signal. Here, we describe the two constructions:

Pooling:

1. Find the first cutoff ω_1^1 and set $\pi(s_1 = s_1^1|\omega) = 1, \forall\omega \geqslant \omega_1^1$: solve ω_1^1 from

$$E[\omega|\omega \geqslant \omega_1^1] = \omega_1^{SD} \tag{4.4}$$

2. Find the second cutoff ω_2^1 and set $\pi(s_2 = s_2^1|\omega) = 1, \forall\omega \geqslant \omega_2^1$: solve ω_2^1 as the smallest $\omega \in \Omega$ such that

$$E[u_2(H,H,\omega)|s_1 = s_1^1, s_2 = s_2^1] + E[u_2(L,H,\omega)|s_1 \neq s_1^1, s_2 = s_2^1] \tag{4.5}$$

$$\geqslant \quad E[u_2(H,L,\omega)|s_1 = s_1^1, s_2 = s_2^1] + E[u_2(L,L,\omega)|s_1 \neq s_1^1, s_2 = s_2^1] \tag{4.6}$$

3. Iteratively construct signals.

Hoshino:

1. Consider a partition of the state space $\Omega = \Omega_1 \bigcup \Omega_2 \bigcup \Omega_3$ such that $\omega > \omega^{SD}, \forall\omega \in \Omega_1$, $Prob(\omega \in \Omega_1) = \epsilon$, and $E_{\Omega_2}[\omega] \geqslant \omega^{RD}$.

2. Find $p \in (0,1)$ and set $\pi((s_1,s_2)|\omega) = 1, \forall\omega \in \Omega_1$ and $\pi((s_1,s_2)|\omega) = p, \forall\omega \in \Omega_2$ such that $E[w|(s_1,s_2)] \geqslant w^{SD}$.

3. Iteratively construct the following signals on Ω_2 such that:

$$\pi(s|\omega) = \pi(s|\omega'), \forall\omega,\omega' \in \Omega_2, \tag{4.7}$$

and the agent's belief on the opponent's playing H after observing any signal is high enough for them to by willing to play H.

4.3 Comparisons

In general, when full implementation is impossible, it is hard to tell which of the two constructions performs better and the fact that both constructions could involve infinitely many signals adds to the difficulty in comparing the different partial implementations. While Hoshino guarantees full implementation on states expected to be above the risk dominance threshold, the prescription necessarily leaves out all lower states where agents play the undesirable action L. The pooling construction, to the contrary, could proceed iteratively to include states that are below the risk dominance threshold, but whether it outperforms the Hoshino construction depends on when it stops. As we will see in the following examples, the prior distribution of the states plays an important role which sheds light on the relative strength of these constructions.

Observation 4.3.1. *If $\omega_1^1 < \omega_1^{SDd}$, compared to the pooling construction, the designer can induce (H, H) with a strictly higher probability using Hoshino's construction.*

Proof. $\omega_1^1 < \omega_1^{SDd}$ implies that:

$$E[u_1(H,H)|s_1] < E[u_1(L,H)|s_1], \forall s_1 \neq s_1^1.$$

Therefore, using pooling, the probability of inducing (H, H) is:

$$P(s_1^1) = \int_{\omega \geqslant \omega_1^1} \omega d\mu_0(\omega).$$

Since $\omega^{RD} < \min\{\omega_1^{SD}, \omega_2^{SD}\}$,

$$P(\Omega_2) = \int_{\omega \geqslant \omega^{RD}} \omega d\mu_0(\omega) > P(s_1^1).$$

■

This example highlights an extreme case where the pooling construction is least effective in the sense that we could only afford the construction of the first signal and the remaining states are already too low to induce further switching. Therefore, we could only induce playing H on states that are expected to be above the strict dominance threshold ω^{SD} using pooling. However, for Hoshino, we could induce playing H on states that are expected to be above the risk dominance threshold ω^{RD} which is necessarily lower than the strict dominance threshold and therefore at a higher probability.

Observation 4.3.2. *If $supp(\mu_0) = \{\bar{\omega}\} \bigcup \Omega_0$, where $\omega \in (\omega^{SDd}, \omega^{RD})$ for any $\omega \in \Omega_0$ and $\bar{\omega} > \omega^{SD}$, and $\mu_0(\bar{\omega}) = \epsilon$, compared to Hoshino's construction, the designer can induce (H, H) with a strictly higher probability using the pooling construction.*

Proof. Find ω_1^1 in the same way for both constructions. For Hoshino's construction, $\forall \omega < \omega_1^1$, $\omega \in \Omega_3$. Therefore, the probability of inducing (H, H) is:

$$1 - P(\Omega_3) = \int_{\omega \geqslant \omega_1^1} \omega d\mu_0(\omega). \tag{4.8}$$

For the pooling construction, since $\forall \omega \geqslant \omega_1^1$,

$$u_2(H, H, \omega) > u_2(H, L, \omega), \tag{4.9}$$

we can always find a $\omega_2^1 < \omega_1^1$ such that the inequality constraint in the step 2 of the

pooling construction is still satisfied. Therefore, the probability of inducing (H, H) is at least

$$\int_{\omega \geqslant \omega_2^1} \omega d\mu_0(\omega) > \int_{\omega \geqslant \omega_1^1} \omega d\mu_0(\omega). \tag{4.10}$$

■

This example highlights another extreme case where the Hoshino construction is least effective in the same sense as before where we could only afford the construction of the first signal. Here, while we could construct the first single which leverages the strict dominance using either pooling or Hoshino, we could proceed the construction with pooling but not Hoshino due to the remaining states being lower than the risk dominance threshold.

Chapter 5

Conclusion

In the chapter on quality control, we showed the optimality of sequential inspection among a selection of protocols. We further provided some insights on how to order the agents using sequential inspection when they are different in terms of their ability to acquire information. We leave a complete analysis of a model with more general assumption on differences between agents for future studies.

In the chapter on self persuasion of a present biased agent, we characterized the optimal information choice for a risk averse present biased agent and provided an example that shows that intertemporal substitution would play a role in deciding the information choice, while risk aversion may not.

In the chaper on information design and risk dominance, we provided two examples to illustrate the relative strength of the two prominent constructions of optimal information structure currently available in the literature.

Appendix A

Omitted Proofs in Chapter 3

A.1 Proof of Proposition 3.4.2:

Self-1's problem Recall that for any consumption choice $c \in \mathcal{A}_b$ and any feasible posterior belief p on the high state $\overline{R}$, self-1's lifetime expected utility is:

$$U_1(c;p) = c + \beta\delta E_p[R](1-c) \tag{A.1}$$

Therefore, self-1's optimal action given a posterior belief p is:

$$c^*(p) = \begin{cases} 1 & \text{, if } p < \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}} \\ 0 & \text{, if } p \geqslant \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}} \end{cases} \tag{A.2}$$

Self-0's problem We follow the concavification approach from Kamenica and Gentzkow (2011).[1] Anticipating self-1 choosing $c^*(p)$, self-0's indirect utility evaluated at $t = 1$ is:

$$U_0(c^*(p);p) = c^*(p) + \delta E_p[R](1-c^*(p)) \tag{A.3}$$

$$= \begin{cases} 1 & \text{, if } p < \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}} \\ \delta E_p[R] & \text{, if } p \geqslant \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}} \end{cases} \tag{A.4}$$

[1] We present an algebraic rather than pictorial proof, but we still follow the concavification approach by considering convex combinations of $U_0(c^*(p_1);p_1)$ and $U_0(c^*(p_2);p_2)$. We want to emphasize that it is essential to have the expected utility formulation to use this approach.

Denote $\overline{U_0(c^*(p);p)}$ as the concave envelope of self-0's indirect utility function.

Lemma A.1.1. *Consider the following two functions of p:*

$$L_1(p) = 1 + (\delta\overline{R} - 1)p \tag{A.5}$$

$$L_2(p) = \begin{cases} 1 + \frac{\frac{1}{\beta}-1}{\frac{1}{\beta\delta}-\underline{R}}(\overline{R}-\underline{R})p & \text{, if } p \in [0, \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}}] \\ \delta E_p[R] & \text{, if } p > \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}} \end{cases} \tag{A.6}$$

$\overline{U_0(c^*(p);p)}$ *is either* $L_1(p)$ *or* $L_2(p)$ *and cannot be anything else.*

Proof. By definition of concave envelope, for any point $\tilde{p}$,

$$\overline{U_0(c^*(\tilde{p});\tilde{p})} = max_{p_1\in[0,\tilde{p}],p_2\in[\tilde{p},1]}\frac{\tilde{p}-p_2}{p_1-p_2}U_0(c^*(p_1);p_1) + \frac{p_1-\tilde{p}}{p_1-p_2}U_0(c^*(p_2);p_2) \tag{A.7}$$

1. If $\tilde{p} \in [0, \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}})$:
   $\forall p_1 \in [0,\tilde{p}]$, $U_0(c^*(p_1);p_1) = 1$, which helps to simplify $\overline{U_0(c^*(\tilde{p});\tilde{p})}$ to:

$$\overline{U_0(c^*(\tilde{p});\tilde{p})} = max_{p_1\in[0,\tilde{p}],p_2\in[\tilde{p},1]}\frac{\tilde{p}-p_2}{p_1-p_2} + \frac{p_1-\tilde{p}}{p_1-p_2}U_0(c^*(p_2);p_2) \tag{A.8}$$

   Given that $U_0(c^*(p_2);p_2)$ is strictly increasing in p_2 on $[\frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}}, 1]$, we can further simplify it:

$$\overline{U_0(c^*(\tilde{p});\tilde{p})} = max_{p_1\in[0,\tilde{p}],p_2\in[\frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}},1]}\frac{\tilde{p}-p_2}{p_1-p_2} + \frac{p_1-\tilde{p}}{p_1-p_2}U_0(c^*(p_2);p_2) \tag{A.9}$$

   p_1 should be optimally set at 0, since for any $p_2 \in [\frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}}, 1]$, the objective

function is strictly decreasing in p_1:

$$\begin{aligned}\frac{\partial}{\partial p_1}\Big[\frac{\tilde{p}-p_2}{p_1-p_2}+\frac{p_1-\tilde{p}}{p_1-p_2}U_0(c^*(p_2);p_2)\Big] &= \frac{p_2-\tilde{p}}{(p_1-p_2)^2}(1-U_0(c^*(p_2);p_2)) \\ &< 0, \quad \text{since } U_0(c^*(p_2);p_2)>1.\end{aligned}$$

$\overline{U_0(c^*(\tilde{p});\tilde{p})}$ can then be further simplified to:

$$\overline{U_0(c^*(\tilde{p});\tilde{p})} = max_{p_2\in[\frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}},1]} \frac{\tilde{p}-p_2}{(-p_2)}+\frac{\tilde{p}}{p_2}U_0(c^*(p_2);p_2) \tag{A.10}$$

$$= max_{p_2\in[\frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}},1]} 1+\frac{\tilde{p}}{p_2}(-1+\delta E_{p_2}[R]) \tag{A.11}$$

Note that:

$$\frac{d}{dp_2}[1+\frac{\tilde{p}}{p_2}(-1+\delta E_{p_2}[R])] = \frac{\tilde{p}}{(p_2)^2}(1-\delta\underline{R}) \tag{A.12}$$

We consider the following cases:

(a) If $\frac{1}{\delta} < \underline{R}$:
$p_2^* = \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}}$ is the unique maximizer of $1+\frac{\tilde{p}}{p_2}(-1+\delta E_{p_2}[R])$, and

$$\begin{aligned}\overline{U_0(c^*(\tilde{p});\tilde{p})} &= 1+\frac{\tilde{p}}{p_2^*}(-1+\delta E_{p_2^*}[R]) \\ &= 1+\frac{\frac{1}{\beta}-1}{\frac{1}{\beta\delta}-\underline{R}}(\overline{R}-\underline{R})\tilde{p} \\ &= L_2(\tilde{p})\end{aligned}$$

(b) If $\frac{1}{\delta} > \underline{R}$:

$p_2^* = 1$ is the unique maximizer of $1 + \frac{\tilde{p}}{p_2}(-1 + \delta E_{p_2}[R])$, and

$$\begin{aligned}
\overline{U_0(c^*(\tilde{p});\tilde{p})} &= 1 + \frac{\tilde{p}}{p_2^*}(-1 + \delta E_{p_2^*}[R]) \\
&= 1 + (\delta\overline{R} - 1)\tilde{p} \\
&= L_1(\tilde{p})
\end{aligned}$$

2. If $\tilde{p} \in [\frac{\frac{1}{\beta\delta} - \underline{R}}{\overline{R} - \underline{R}}, 1]$:

   $\forall p_2 \in [\tilde{p}, 1]$, $U_0(c^*(p_2); p_2) = \delta E_{p_2}[R]$, which helps to simplify $\overline{U_0(c^*(\tilde{p});\tilde{p})}$ to:

$$\begin{aligned}
\overline{U_0(c^*(\tilde{p});\tilde{p})} &= max_{p_1 \in [0,\tilde{p}], p_2 \in [\tilde{p},1]} \frac{\tilde{p} - p_2}{p_1 - p_2} U_0(c^*(p_1); p_1) + \frac{p_1 - \tilde{p}}{p_1 - p_2} \delta E_{p_2}[R] \\
&= max\Big\{ max_{p_1 \in [0, \frac{\frac{1}{\beta\delta} - \underline{R}}{\overline{R} - \underline{R}}), p_2 \in [\tilde{p},1]} \frac{\tilde{p} - p_2}{p_1 - p_2} + \frac{p_1 - \tilde{p}}{p_1 - p_2} \delta E_{p_2}[R], \delta E_{\tilde{p}}[R] \Big\}
\end{aligned}$$

   To determine the maximum, let us first discuss the maximization of $\frac{\tilde{p} - p_2}{p_1 - p_2} + \frac{p_1 - \tilde{p}}{p_1 - p_2} \delta E_{p_2}[R]$ with respect to p_1.

$$\begin{aligned}
\frac{\partial}{\partial p_1} \Big\{ \frac{\tilde{p} - p_2}{p_1 - p_2} + \frac{p_1 - \tilde{p}}{p_1 - p_2} \delta E_{p_2}[R] \Big\} &= \frac{p_2 - \tilde{p}}{(p_1 - p_2)^2} \big\{ 1 - \delta E_{p_2}[R] \big\} \\
&< 0, \quad \text{since } \delta E_{p_2}[R] > 1, \forall p_2 \geqslant \tilde{p} \geqslant \frac{\frac{1}{\beta\delta} - \underline{R}}{\overline{R} - \underline{R}}
\end{aligned}$$

   This implies the optimality of setting $p_1 = 0$ for $p_1 \in [0, \frac{\frac{1}{\beta\delta} - \underline{R}}{\overline{R} - \underline{R}})$ and hence the following simplification:

$$\begin{aligned}
max_{p_1 \in [0, \frac{\frac{1}{\beta\delta} - \underline{R}}{\overline{R} - \underline{R}}), p_2 \in [\tilde{p},1]} \frac{\tilde{p} - p_2}{p_1 - p_2} + \frac{p_1 - \tilde{p}}{p_1 - p_2} \delta E_{p_2}[R] &= max_{p_2 \in [\tilde{p},1]} \frac{\tilde{p} - p_2}{(-p_2)} + \frac{\tilde{p}}{p_2} \delta E_{p_2}[R] \\
&= max_{p_2 \in [\tilde{p},1]} 1 + \frac{\tilde{p}}{p_2} (\delta E_{p_2}[R] - 1)
\end{aligned}$$

$p_2 \geqslant \tilde{p} \geqslant \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}}$ implies that this simplified objective function is increasing in p_2 and hence the optimality of setting $p_2 = 1$, which further implies that:

$$\begin{aligned}\overline{U_0(c^*(\tilde{p});\tilde{p})} &= max\{1+\tilde{p}(\delta\overline{R}-1), \delta E_{\tilde{p}}[R]\} \\ &= \tilde{p}\delta\overline{R} + (1-\tilde{p})max\{1, \delta\underline{R}\}\end{aligned}$$

Again, we consider the following cases:

(a) If $\frac{1}{\delta} < \underline{R}$:

$$\overline{U_0(c^*(\tilde{p});\tilde{p})} = \tilde{p}\delta\overline{R} + (1-\tilde{p})\delta\underline{R} = L_2(\tilde{p})$$

Since in this case, $\delta E_{\tilde{p}}[R] > 1 + \tilde{p}(\delta\overline{R} - 1)$, it is only necessary to have $p_1 \geqslant \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}}$. Indeed, $\overline{U_0(c^*(\tilde{p});\tilde{p})}$ is independent of the p_1 and p_2 as long as $p_1, p_2 \in [\frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}}, 1]$, since the objective function:

$$\begin{aligned}\frac{\tilde{p}-p_2}{p_1-p_2}U_0(c^*(p_1);p_1) + \frac{p_1-\tilde{p}}{p_1-p_2}U_0(c^*(p_2);p_2) &= \frac{\tilde{p}-p_2}{p_1-p_2}\delta E_{p_1}[R] \\ &\quad + \frac{p_1-\tilde{p}}{p_1-p_2}\delta E_{p_2}[R] \\ &= \delta E_{\tilde{p}}[R]\end{aligned}$$

(b) If $\frac{1}{\delta} > \underline{R}$:

$$\overline{U_0(c^*(\tilde{p});\tilde{p})} = \tilde{p}\delta\overline{R} + (1-\tilde{p}) = L_1(\tilde{p})$$

■

We've shown in the above proof to the lemma that (1) if $\frac{1}{\delta} < \underline{R}$, $\overline{U_0(c^*(p);p)} = L_2(p)$; (2) if $\frac{1}{\delta} > \underline{R}$, $\overline{U_0(c^*(p);p)} = L_1(p)$. We've also shown that:

1. If $\frac{1}{\delta} < \underline{R}$, when the prior is (1) $p^* \in [0, \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}})$, $p_1 = 0$ and $p_2 = \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}}$; (2) $p^* \in [\frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}}, 1]$, $p_1 = p_2 = p^*$[2].

2. If $\frac{1}{\delta} > \underline{R}$, $p_1 = 0$ and $p_2 = 1$.

p_1^* is the posterior belief on $\overline{R}$ after observing a low state message, and p_2^* is the posterior belief on $\overline{R}$ after observing a high state message. Bayes rule implies that, for any p_0,

$$p_1^* = \frac{p_0\pi(\overline{R})(\underline{r})}{p_0\pi(\overline{R})(\underline{r}) + (1-p_0)\pi(\underline{R})(\underline{r})}; \quad p_2^* = \frac{p_0\pi(\overline{R})(\overline{r})}{p_0\pi(\overline{R})(\overline{r}) + (1-p_0)\pi(\underline{R})(\overline{r})}$$

Since we've known the p_1^* and p_2^*, we can back out self-0's optimal strategy:

1. If $\frac{1}{\delta} < \underline{R}$:

   (a) $\forall p_0 \in [0, \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}})$,

$$\begin{aligned} p_1^* = 0 \quad &\Rightarrow \quad \pi(\overline{R})(\underline{r}) = 0, \pi(\overline{R})(\overline{r}) = 1. \\ \text{Also, since } p_2^* = \frac{\frac{1}{\beta\delta} - \underline{R}}{\overline{R} - \underline{R}} \quad &\Rightarrow \quad \pi(\underline{R})(\overline{r}) = \frac{1-p_2^*}{p_2^*}\frac{p_0}{1-p_0}, \\ &\qquad\quad \pi(\overline{R})(\overline{r}) = \frac{p_2^* - p_0}{p_2^*(1-p_0)}. \end{aligned}$$

   (b) $\forall p_0 \in [\frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}}, 1]$,

$$p_1^* = p_2^* = p_0 \Rightarrow \pi(\overline{R})(\underline{r}) = \pi(\underline{R})(\underline{r}), \pi(\overline{R})(\overline{r}) = \pi(\underline{R})(\overline{r}).$$

This condition characterizes all the uninformative signal structures. By

[2] This is the case where self-0 is actually indifferent among infinitely many feasible signal structures which satisfy that $p_1, p_2 \geqslant \frac{\frac{1}{\beta\delta}-\underline{R}}{\overline{R}-\underline{R}}$.

uninformative, we mean that self-1 could not update his belief to some $p \neq p_0$.

2. If $\frac{1}{\delta} > \underline{R}$:

$$p_1^* = 0 \quad \Rightarrow \quad \pi(\overline{R})(\underline{r}) = 0, \pi(\overline{R})(\overline{r}) = 1.$$
$$\text{Also, since } p_2^* = 1 \quad \Rightarrow \quad \pi(\underline{R})(\overline{r}) = 0, \pi(\overline{R})(\overline{r}) = 1.$$

Alternative representation of Self-1's strategy Self-1 observes a signal realization, $r \in \mathcal{M} = \{\overline{R}, \underline{R}\} = \Omega$, and then best responds to it with his Bayesian updated belief. Consider self-1's strategy as the mapping $C : \mathcal{M} \to \mathcal{A}_b$. Denote $p(r; p_0, \pi(\cdot)(\overline{R}))$ as self-1's posterior belief after observing the signal realization r. Then, as discussed above,

$$p(\underline{r}; p_0, \pi(\cdot)(\overline{r})) = p_1^*(p_0, \pi(\cdot)(\overline{r})), \quad p(\overline{r}; p_0, \pi(\cdot)(\overline{r})) = p_2^*(p_0, \pi(\cdot)(\overline{r}))$$

Therefore, for each case with respect to p_0 and $\pi(r)(\overline{r})$,

$$C(\overline{r}) = C(p(\overline{r}; \cdot)) = C(p_1^*), \quad C(\underline{r}) = C(p(\underline{r}; \cdot)) = C(p_2^*)$$

Specifically,

1. If $\frac{1}{\delta} < \underline{R}$:
   (a) $\forall p_0 \in [0, \frac{\frac{1}{\beta\delta} - \underline{R}}{\overline{R} - \underline{R}})$, $C(p_1^*) = 1$ and $C(p_2^*) = 0$;
   (b) $\forall p_0 \in [\frac{\frac{1}{\beta\delta} - \underline{R}}{\overline{R} - \underline{R}}, 1]$, $C(p_1^*) = C(p_2^*) = 0$;
2. If $\frac{1}{\delta} > \underline{R}$:
   $C(p_1^*) = 1$ and $C(p_2^*) = 0$.

A.2 Proof of Proposition 3.5.1:

For any belief p, Self-0's indirect utility function is:

$$U_0(t^*(p);p) = \frac{(1+t^*(p))^{-\rho}}{\rho}(1+\beta^{-1}t^*(p)), \tag{A.13}$$

where the optimal consumption ratio:

$$t^*(p) \equiv \left\{\beta\delta E_p[R^\rho]\right\}^{\frac{1}{1-\rho}}. \tag{A.14}$$

To apply the concavification approach, we first take a look at the curvature of this indirect utility function:

$$\begin{aligned}
\frac{\partial^2 U_0(t^*(p);p)}{\partial p^2} &= \left\{(\beta-\rho)t^*(p)+(1-\beta\rho)\right\}(1+t^*(p))^{-\rho-2}t^*(p)^{2\rho-1}(1-\rho)^{-2} \\
&\quad \times\beta\delta^2(\overline{R}^\rho-\underline{R}^\rho)^2 \\
&= \begin{cases} >0 & \text{if } \rho\leqslant\beta \text{ or } \{\rho>\beta \,\&\, t^*(p)<\frac{1-\beta\rho}{\rho-\beta}\}, \\ =0 & \text{if } \rho>\beta \,\&\, t^*(p)=\frac{1-\beta\rho}{\rho-\beta}, \\ <0 & \text{if } \rho>\beta \,\&\, t^*(p)>\frac{1-\beta\rho}{\rho-\beta}. \end{cases}
\end{aligned}$$

Since when $\rho \leqslant \beta$, we have global convexity, full information will be optimal. When $\rho > \beta$, since $t^*(p)$ increases in p monotonically as:

$$t^{*'}(p) = \frac{1}{1-\rho}t^*(p)^\rho\beta\delta(\overline{R}^\rho-\underline{R}^\rho) > 0, \tag{A.15}$$

which implies that $t^*(p) \leqslant t^*(1)$. If $\overline{R}$ is small enough such that:

$$t^*(1) = \left\{\beta\delta\overline{R}^\rho\right\}^{\frac{1}{1-\rho}} < \frac{1-\beta\rho}{\rho-\beta}, \tag{A.16}$$

then $t^*(p) < \frac{1-\beta\rho}{\rho-\beta}$ is satisfied for all p and hence $U_0(t^*(p);p)$ is global convex and full information is optimal. Similarly, if $\underline{R}$ is large enough such that:

$$t^*(0) = \left\{\beta\delta\underline{R}^\rho\right\}^{\frac{1}{1-\rho}} > \frac{1-\beta\rho}{\rho-\beta}, \tag{A.17}$$

then $t^*(p) > \frac{1-\beta\rho}{\rho-\beta}$ is satisfied for all p and hence $U_0(t^*(p);p)$ is global concave and no information is optimal. Besides these two extreme cases, the monotonicity of $t^*(p)$ also implies that, for any $\underline{R}, \overline{R}$ such that $\underline{R} \leqslant R_0 \leqslant \overline{R}$, where

$$R_0 \equiv [\frac{1}{\beta\delta}(\frac{1-\beta\rho}{\rho-\beta})^{1-\rho}]^{\frac{1}{\rho}}, \tag{A.18}$$

then, as we increase p from 0 to 1, $U_0(t^*(p);p)$ starts out being convex till some point $p(R_0)$ and then becomes concave all the way until $p = 1$. The turning point $p(R_0)$ is solved from:

$$E_{p(R_0)}[R^\rho] = R_0^\rho. \tag{A.19}$$

Next, we want to characterize the full information case where the line connecting the two endpoints $U_0(t^*(0);0)$ and $U_0(t^*(1);1)$ happens to be tangent to $U_0(t^*(p);p)$ at $p = 1$. We will see that this is a marginal case as discussed below. The slope of the line representing full information is:

$$\frac{U_0(t^*(1);1) - U_0(t^*(0);0)}{1} = (1+t^*(1))^{-\rho}\rho^{-1}\beta^{-1}\left\{t^*(1)+\beta-(\frac{1+t^*(0)}{1+t^*(1)})^{-\rho}(\beta+t^*(0))\right\}. \tag{A.20}$$

The slope of $U_0(t^*(p);p)$ at $p=1$ is:

$$\begin{aligned}\frac{\partial U_0(t^*(1);1)}{\partial p}|_{p=1} &= (1+t^*(1))^{-\rho}\rho^{-1}\beta^{-1} \\ &\times\{t^*(1)+\beta-\frac{1}{1+t^*(1)}\{[t^*(1)\frac{\beta-\rho}{1-\rho}+\beta] \\ &+t^*(1)(\frac{R}{\underline{R}})^{\rho}[t^*(1)+\frac{1-\rho\beta}{1-\rho}]\}\}.\end{aligned}$$

Rearranging terms, the slopes are equal if and only if:

$$(1+t^*(0))^{\rho}=(1+t^*(1))^{1+\rho}\frac{\beta+t^*(0)}{[t^*(1)\frac{\beta-\rho}{1-\rho}+\beta]+t^*(1)(\frac{R}{\underline{R}})^{\rho}[t^*(1)+\frac{1-\rho\beta}{1-\rho}]}. \tag{A.21}$$

Also, $\frac{\partial U_0(t^*(1);1)}{\partial p}|_{p=1} \geqslant \frac{U_0(t^*(1);1)-U_0(t^*(0);0)}{1}$ if and only if either one of the following two cases is satisfied:

1. $[t^*(1)\frac{\beta-\rho}{1-\rho}+\beta]+t^*(1)(\frac{R}{\underline{R}})^{\rho}[t^*(1)+\frac{1-\rho\beta}{1-\rho}] \leqslant 0$;
2. $[t^*(1)\frac{\beta-\rho}{1-\rho}+\beta]+t^*(1)(\frac{R}{\underline{R}})^{\rho}[t^*(1)+\frac{1-\rho\beta}{1-\rho}] > 0$, and $(1+t^*(0))^{\rho} \leqslant (1+t^*(1))^{1+\rho}\frac{\beta+t^*(0)}{[t^*(1)\frac{\beta-\rho}{1-\rho}+\beta]+t^*(1)(\frac{R}{\underline{R}})^{\rho}[t^*(1)+\frac{1-\rho\beta}{1-\rho}]}$.

Denote LHS, RHS as the left and right hand side of the above equation respectively. Note first that $\underline{R} \in [0,\overline{R}]$. If we take $\underline{R}=0$, then $t^*(0)=0$, and we have:

$$LHS = 1 \tag{A.22}$$

$$RHS = (1+t^*(1))^{1+\rho}\frac{\beta}{t^*(1)\frac{\beta-\rho}{1-\rho}+\beta} < 0. \tag{A.23}$$

To see the last inequality, since $t^*(1) \geqslant \frac{1-\beta\rho}{\rho-\beta}$ and $\rho > \beta$,

$$t^*(1)\frac{\beta-\rho}{1-\rho}+\beta \leqslant \frac{1-\beta\rho}{\rho-\beta}\frac{\beta-\rho}{1-\rho}+\beta=\frac{\beta-1}{1-\rho}<0. \tag{A.24}$$

If we take $\underline{R} = \overline{R}$, then:

$$
\begin{aligned}
LHS &= (1 + t^*(1))^{\rho} && \text{(A.25)} \\
RHS &= (1 + t^*(1))^{1+\rho} \frac{\beta + t^*(1)}{\left[t^*(1)\frac{\beta-\rho}{1-\rho} + \beta\right] + t^*(1)\left[t^*(1) + \frac{1-\rho\beta}{1-\rho}\right]} && \text{(A.26)} \\
&= (1 + t^*(1))^{\rho} \frac{\beta + (1+\beta)t^*(1) + t^*(1)^2}{\beta + \frac{(1-\rho\beta)+(\beta-\rho)}{1-\rho} t^*(1) + t^*(1)^2} < LHS, && \text{(A.27)}
\end{aligned}
$$

since $1 + \beta < \frac{(1-\rho\beta)+(\beta-\rho)}{1-\rho}$. Next, we want to show that there exists $\underline{R}_0$ such that:

$$
\left[t^*(1)\frac{\beta-\rho}{1-\rho} + \beta\right] + t^*(1)\left(\frac{\underline{R}}{\overline{R}}\right)^{\rho}\left[t^*(1) + \frac{1-\rho\beta}{1-\rho}\right] = 0. \tag{A.28}
$$

We've shown above that value of the left hand side terms is strictly negative at $\underline{R} = 0$ and strictly positive at $\underline{R} = \overline{R}$. It is also easy to see that it monotonically increases in $\underline{R}$, which implies the existence of a unique $\underline{R}_0$ that solves this equation. This further implies that as $R \to \underline{R}_0^+$, $RHS \to \infty$ and hence is larger than LHS. Together with $RHS < LHS$ at $\underline{R} = \overline{R}$, we know there exists a R_0 such that $LHS = RHS$. A sufficient condition for the uniqueness of this R_0 is the convexity of the RHS in $\underline{R}$:

$$
\begin{aligned}
\frac{\partial^2 RHS}{\partial p^2} &= \left[t^*(1)\frac{\beta-\rho}{1-\rho} + \beta\right]\frac{\rho}{1-\rho} t^*(0)\underline{R}^{-2}\left[\frac{\rho}{1-\rho} - 1\right] \\
&\quad + t^*(1)\left(\frac{\underline{R}}{\overline{R}}\right)^{\rho}\left[t^*(1) + \frac{1-\rho\beta}{1-\rho}\right]\underline{R}^{-2}\left[\frac{\rho^2}{1-\rho} t^*(0)\left(\rho - 1 + \frac{\rho}{1-\rho}\right) - (\rho-1)\beta\rho\right] \\
&> t^*(1)\left(\frac{\underline{R}}{\overline{R}}\right)^{\rho}\left[t^*(1) + \frac{1-\rho\beta}{1-\rho}\right]\underline{R}^{-2}\left\{\frac{\rho}{1-\rho} t^*(0)(1-\rho)^3 + (1-\rho)\beta\rho\right\}, \\
&\quad \forall \underline{R} \in (R_0, \overline{R}) \\
&> 0.
\end{aligned}
$$

We summarize the comparisons of the two slopes:

1. for $\underline{R} \in [0, R_0)$, $\frac{\partial U_0(t^*(1);1)}{\partial p}|_{p=1} > \frac{U_0(t^*(1);1)-U_0(t^*(0);0)}{1}$;

2. for $\underline{R} = R_0$, $\frac{\partial U_0(t^*(1);1)}{\partial p}|_{p=1} = \frac{U_0(t^*(1);1)-U_0(t^*(0);0)}{1}$;

3. for $\underline{R} \in (R_0, \overline{R}]$, $\frac{\partial U_0(t^*(1);1)}{\partial p}|_{p=1} < \frac{U_0(t^*(1);1)-U_0(t^*(0);0)}{1}$.

These comparisons together with the curvature of $U_0(t^*(p);p)$ we've discussed above implies that full information is optimal for all $\underline{R} \leqslant R_0$, and partial information is optimal for all $\underline{R} \in (R_0, \overline{R}]$.

A.3 Proof of Proposition 3.6.1:

Note that $U_0^T(p; c^*(p), \alpha)$ is a piecewise linear function that is constant at 0 on $p \in [0, \overline{p}^1(\alpha; \cdot)]$ and increases linearly in p on $[\overline{p}^1(\alpha; \cdot), 1]$ from $\frac{(\frac{1}{\beta})^\alpha - 1}{\alpha}$ to $\frac{\delta^\alpha \overline{R}^\alpha - 1}{\alpha}$. Therefore, we have a parallel lemma to that one used in proving Proposition 3.4.2:

Lemma A.3.1. *Consider the following two functions of p:*

$$L_1(p; \alpha) = \frac{\delta^\alpha \overline{R}^\alpha - 1}{\alpha} p \tag{A.29}$$

$$L_2(p; \alpha) = \begin{cases} \frac{(\frac{1}{\beta})^\alpha - 1}{\alpha \overline{p}^1(\alpha; \cdot)} p & , \text{if } p \in [0, \overline{p}^1(\alpha; \cdot)] \\ \frac{\delta^\alpha E_p[R^\alpha] - 1}{\alpha} & , \text{if } p > \overline{p}^1(\alpha; \cdot) \end{cases}$$

$\overline{U_0^T(p; c^*(p), \alpha)}$ *is either* $L_1(p; \alpha)$ *or* $L_2(p; \alpha)$ *and cannot be anything else.*

Whether $L_1(p; \alpha)$ or $L_2(p; \alpha)$ is the concave envelope hinges on the comparison of their values at the belief threshold $\overline{p}^1(\alpha; \cdot)$.

$$L_1(\overline{p}^1(\alpha; \cdot); \alpha) = \frac{\delta^\alpha \overline{R}^\alpha - 1}{\alpha} \overline{p}^1(\alpha; \cdot) = \frac{\delta^\alpha \overline{R}^\alpha - 1}{\alpha} \frac{(\frac{1}{\beta\delta})^\alpha - \underline{R}^\alpha}{\overline{R}^\alpha - \underline{R}^\alpha} \tag{A.31}$$

$$L_2(\overline{p}^1(\alpha; \cdot); \alpha) = \frac{(\frac{1}{\beta})^\alpha - 1}{\alpha \overline{p}^1(\alpha; \cdot)} \overline{p}^1(\alpha; \cdot) = \frac{(\frac{1}{\beta})^\alpha - 1}{\alpha} \tag{A.32}$$

Therefore,

$$L_1(\overline{p}^1(\alpha; \cdot); \alpha) - L_2(\overline{p}^1(\alpha; \cdot); \alpha) = \frac{(\overline{R}^\alpha - (\frac{1}{\beta\delta})^\alpha)(1 - \delta^\alpha \underline{R}^\alpha)}{\alpha(\overline{R}^\alpha - \underline{R}^\alpha)} \propto \frac{1 - \delta^\alpha \underline{R}^\alpha}{\alpha}. \tag{A.33}$$

We are back to the two familiar cases:

1. $\frac{1}{\delta} < \underline{R}$: $\forall \alpha$, $\overline{U_0^T(p; c^*(p), \alpha)} = L_2(\overline{p}^1(\alpha; \cdot); \alpha)$.

2. $\frac{1}{\delta} > \underline{R}$: $\forall \alpha$, $\overline{U_0^T(p; c^*(p), \alpha)} = L_1(\overline{p}^1(\alpha; \cdot); \alpha)$.

The optimal signal structure depends on how $\frac{1}{\delta}$ compares with $\underline{R}$ but not α. We summarize the findings: when $\frac{1}{\delta} < \underline{R}$, the optimal signal structure is partially informative; when $\frac{1}{\delta} > \underline{R}$, the optimal signal structure is (fully) informative. In either case, it's exactly the same as that characterized in Proposition 3.4.2.

A.4 Proof of Lemma 3.6.2:

STEP1: MONOTONICITY OF $(E_p[R^\alpha])^{\frac{1}{\alpha}}$ IN α Since $R \in \Omega = \{\underline{R}, \overline{R}\}$ and $\underline{R}, \overline{R} \in \mathbb{R}$, for any $p \in [0,1]$, we know that $E|X| < \infty$. Consider $\phi_1(x) = x^{\frac{\alpha_1}{\alpha_2}}$, with $-\infty < \alpha_1 < \alpha_2 < 0$. $\phi_1(\cdot)$ is a strictly convex function, as $\phi_1''(x) = \frac{\alpha_1}{\alpha_2}(\frac{\alpha_1}{\alpha_2} - 1)x^{\frac{\alpha_1}{\alpha_2}-2} > 0$. Also, $\forall p \in (0,1)$, $\forall \alpha$, the random variable R is not almost surely constant. Jensen's Inequality implies that:

$$(E_p[R^{\alpha_2}])^{\frac{\alpha_1}{\alpha_2}} < E_p[R^{\alpha_2 \frac{\alpha_1}{\alpha_2}}] \Leftrightarrow (E_p[R^{\alpha_2}])^{\frac{1}{\alpha_2}} > (E_p[R^{\alpha_1}])^{\frac{1}{\alpha_1}}, \text{as } \alpha_1 < 0. \tag{A.34}$$

Similarly, consider $\phi_2(x) = x^{\frac{\alpha_1}{\alpha_2}}$, with $1 > \alpha_1 > \alpha_2 > 0$. $\phi_2(\cdot)$ is also a convex function. Again, by Jensen's Inequality, we have:

$$(E_p[R^{\alpha_2}])^{\frac{\alpha_1}{\alpha_2}} < E_p[R^{\alpha_2 \frac{\alpha_1}{\alpha_2}}] \Leftrightarrow (E_p[R^{\alpha_2}])^{\frac{1}{\alpha_2}} < (E_p[R^{\alpha_1}])^{\frac{1}{\alpha_1}}, \text{as } \alpha_1 > 0. \tag{A.35}$$

Therefore, we've shown that the effective expected rate of return, $(E_p[R^\alpha])^{\frac{1}{\alpha}}$, is increasing in α on both $(-\infty, 0)$ and $(0,1)$. To see what happens at $\alpha = 0$, we take natural log of $(E_p[R^\alpha])^{\frac{1}{\alpha}}$:

$$ln[(E_p[R^\alpha])^{\frac{1}{\alpha}}] = \frac{ln[p\overline{R}^\alpha + (1-p)\underline{R}^\alpha]}{\alpha}. \tag{A.36}$$

Now that both numerator and denominator are equal to 0 at $\alpha = 0$, we apply L'Hospital's rule:

$$\begin{aligned} lim_{\alpha \to 0} \frac{ln[p\overline{R}^\alpha + (1-p)\underline{R}^\alpha]}{\alpha} &= lim_{\alpha \to 0} \frac{\frac{p\overline{R}^\alpha ln(\overline{R}) + (1-p)\underline{R}^\alpha ln(\underline{R})}{p\overline{R}^\alpha + (1-p)\underline{R}^\alpha}}{1} && \text{(A.37)} \\ &= pln(\overline{R}) + (1-p)ln(\underline{R}). && \text{(A.38)} \end{aligned}$$

Given that $\overline{R}, \underline{R} \in \mathbb{R}$, the limit of $(E_p[R^\alpha])^{\frac{1}{\alpha}}$ at $\alpha = 0$, which is equal to $e^{pln(\overline{R})+(1-p)ln(\underline{R})}$, is well defined.

At $\alpha = 0$, the CRRA utility function takes the special form of the log function. Therefore, instead of $(E_p[R^\alpha])^{\frac{1}{\alpha}}$, we would have:

$$e^{E_p[ln(R)]} = e^{pln(\overline{R})+(1-p)ln(\underline{R})} = lim_{\alpha\to 0}(E_p[R^\alpha])^{\frac{1}{\alpha}}, \tag{A.39}$$

which proves the continuity of $(E_p[R^\alpha])^{\frac{1}{\alpha}}$ in α.

Combined with the monotonicity on both $(-\infty, 0)$ and $(0, 1)$, for any $-\infty < \alpha_1 < 0 < \alpha_2 < 1$, we have $(E_p[R^{\alpha_1}])^{\frac{1}{\alpha_1}} < lim_{\alpha\to 0}(E_p[R^\alpha])^{\frac{1}{\alpha}} < (E_p[R^{\alpha_2}])^{\frac{1}{\alpha_2}}$. Therefore, it has to be that $(E_p[R^\alpha])^{\frac{1}{\alpha}}$ is increasing in α everywhere on $(-\infty, 1)$.

Step2: monotonicity of $\overline{p}^1(\alpha; \beta, \delta, \overline{R}, \underline{R})$ in α We plan to use the implicit function theorem. First, we want to show that $(E_p[R^\alpha])^{\frac{1}{\alpha}}$ is monotonically increasing in p.

$$\frac{\partial}{\partial p}(E_p[R^\alpha])^{\frac{1}{\alpha}} = \frac{1}{\alpha}(E_p[R^\alpha])^{\frac{1}{\alpha}-1}(\overline{R}^\alpha - \underline{R}^\alpha) > 0. \tag{A.40}$$

The Euler equation for self-1 is:

$$1 = \beta\delta(E_p[R^\alpha])^{\frac{1}{\alpha}}. \tag{A.41}$$

If we assume $\underline{R} < \frac{1}{\beta\delta} < \overline{R}$, then there exists an interior $\overline{p}^1(\alpha; \beta, \delta, \overline{R}, \underline{R}) \in (0, 1)$ such that $(E_{\overline{p}^1(\alpha;\beta,\delta,\overline{R},\underline{R})}[R^\alpha])^{\frac{1}{\alpha}} = \frac{1}{\beta\delta}$.

Consider a function $F(\alpha, p; \beta, \delta, \overline{R}, \underline{R}) \equiv (E_p[R^\alpha])^{\frac{1}{\alpha}} - \frac{1}{\beta\delta}$. At any arbitrary $\alpha \in (-\infty, 0) \bigcup (0, 1)$ and $p = \overline{p}^1(\alpha; \beta, \delta, \overline{R}, \underline{R})$, $F(\alpha, \overline{p}^1(\alpha; \cdot); \cdot) = 0$ is satisfied by rearranging the Euler equation.

Next, we check the properties of the partial derivative of $F(\alpha, p; \cdot)$ with respect

to p. We've shown that:

$$F_p(\alpha, p; \cdot) = \frac{\partial}{\partial p}(E_p[R^\alpha])^{\frac{1}{\alpha}} > 0. \tag{A.42}$$

Since it is strictly positive everywhere, it is nonzero at $\overline{p}^1(\alpha; \cdot)$ in particular. Also, it is continuous since it is a monotonic transformation of $E_p[R^\alpha]$, which is linear in p.

Finally, we check the properties of the partial derivative of $F(\alpha, p; \cdot)$ with respect to α.

$$\begin{aligned} F_\alpha(\rho, p; \cdot) &= \frac{\partial}{\partial \alpha}(E_p[R^\alpha])^{\frac{1}{\alpha}} && \text{(A.43)} \\ &= \frac{1}{\alpha}(E_p[R^\alpha])^{\frac{1}{\alpha}-1}\big[p\overline{R}^\alpha ln(\overline{R}) + (1-p)\underline{R}^\alpha ln(\underline{R})\big] && \text{(A.44)} \\ &\quad +(E_p[R^\alpha])^{\frac{1}{\alpha}} ln\big[p\overline{R}^\alpha + (1-p)\underline{R}^\alpha\big](-1)\alpha^{-2}, && \text{(A.45)} \end{aligned}$$

the functional form of which implies that it is necessarily continuous except at $\alpha = 0$. First, we need to check the differentiability of $(E_p[R^\alpha])^{\frac{1}{\alpha}}$ at $\alpha = 0$. Recall that, at $\alpha = 0$, the form of $e^{E_p[ln(R)]} = e^{pln(\overline{R})+(1-p)ln(\underline{R})}$.

$$
\begin{aligned}
& lim_{h\to 0}\frac{(E_p[R^h])^{\frac{1}{h}} - e^{pln(\overline{R})+(1-p)ln(\underline{R})}}{h} \\
= \; & lim_{h\to 0}\frac{e^{\frac{1}{h}ln(p\overline{R}^h+(1-p)\underline{R}^h)} - e^{pln(\overline{R})+(1-p)ln(\underline{R})}}{h} \\
= \; & lim_{h\to 0}\frac{\partial}{\partial h}e^{\frac{1}{h}ln(p\overline{R}^h+(1-p)\underline{R}^h)}, \text{ by L'Hospital's Rule} \\
= \; & lim_{\alpha\to 0}\frac{\partial}{\partial \alpha}(E_p[R^\alpha])^{\frac{1}{\alpha}} \\
= \; & lim_{\alpha\to 0}(E_p[R^\alpha])^{\frac{1}{\alpha}} lim_{\alpha\to 0}\{\frac{p\overline{R}^\alpha ln(\overline{R}) + (1-p)\underline{R}^\alpha ln(\underline{R})}{\alpha E_p[R^\alpha]} - \frac{ln(E_p[R^\alpha])}{\alpha^2}\} \\
= \; & e^{E_p[ln(R)]}\{lim_{\alpha\to 0}\frac{p\overline{R}^\alpha ln(\overline{R}) + (1-p)\underline{R}^\alpha ln(\underline{R})}{\alpha E_p[R^\alpha]} \\
& -lim_{\alpha\to 0}\frac{p\overline{R}^\alpha ln(\overline{R}) + (1-p)\underline{R}^\alpha ln(\underline{R})}{2\alpha E_p[R^\alpha]}\}, \text{ by L'Hospital's Rule} \\
= \; & e^{E_p[ln(R)]} lim_{\alpha\to 0}\frac{p\overline{R}^\alpha ln(\overline{R}) + (1-p)\underline{R}^\alpha ln(\underline{R})}{2\alpha E_p[R^\alpha]} \\
= \; & \infty \text{ or } -\infty
\end{aligned}
$$

Therefore, $(E_p[R^\alpha])^{\frac{1}{\alpha}}$ is not differentiable at $\alpha = 0$ and hence $F_\alpha(\alpha, p; \cdot)$ is not continuous at $\alpha = 0$.

We proceed with applying the implicit function theorem[3]. The implicit function $\overline{p}^1(\alpha; \beta, \delta, \overline{R}, \underline{R})$ is well defined, continuous, and has continuous partial derivatives, except at $\alpha = 0$. This implies:

$$
\begin{aligned}
& \frac{\partial}{\partial \alpha}F(\alpha, \overline{p}^1(\alpha; \beta, \delta, \overline{R}, \underline{R}); \cdot) = F_\alpha(\alpha, \overline{p}^1(\alpha; \cdot); \cdot) + F_p(\alpha, \overline{p}^1(\alpha; \cdot); \cdot)\frac{\partial \overline{p}^1(\alpha; \cdot)}{\partial \alpha} = 0 \\
\Leftrightarrow \quad & \frac{\partial \overline{p}^1(\alpha; \cdot)}{\partial \alpha} = -\frac{F_\alpha(\alpha, \overline{p}^1(\alpha; \cdot); \cdot)}{F_p(\alpha, \overline{p}^1(\alpha; \cdot); \cdot)},
\end{aligned}
$$

[3]We use the version of the implicit function theorem as stated on p. 195 in Chiang (2005).

which is well defined and strictly negative, as it holds everywhere except at $\alpha = 0$ that:

$$F_p(\alpha, \overline{p}^1(\alpha;\cdot);\cdot) > 0, \text{ and } F_\alpha(\alpha, \overline{p}^1(\alpha;\cdot);\cdot) > 0. \tag{A.46}$$

To see that $\overline{p}^1(\alpha;\cdot)$ is also well-defined and continuous at $\alpha = 0$, we can solve for $\overline{p}^1(0;\cdot)$ from the indifferent condition implied by the Euler equation $(E_{\overline{p}^1(\alpha;\cdot)}[R^\alpha])^{\frac{1}{\alpha}} - \frac{1}{\beta\delta} = 0$, where $(E_{\overline{p}^1(\alpha;\cdot)}[R^\alpha])^{\frac{1}{\alpha}}|_{\alpha=0} = e^{E_p[ln(R)]}$:

$$\overline{p}^1(0;\cdot) = \frac{ln(\frac{(\beta\delta)^{-1}}{\underline{R}})}{ln(\frac{\overline{R}}{\underline{R}})}. \tag{A.47}$$

Since:

$$lim_{\alpha\to 0}\overline{p}^1(\alpha;\cdot) = lim_{\alpha\to 0}\frac{(\frac{(\beta\delta)^{-1}}{\underline{R}})^\alpha - 1}{(\frac{\overline{R}}{\underline{R}})^\alpha - 1} = lim_{\alpha\to 0}\frac{(\frac{(\beta\delta)^{-1}}{\underline{R}})^\alpha ln(\frac{(\beta\delta)^{-1}}{\underline{R}})}{(\frac{\overline{R}}{\underline{R}})^\alpha ln(\frac{\overline{R}}{\underline{R}})} = \frac{ln(\frac{(\beta\delta)^{-1}}{\underline{R}})}{ln(\frac{\overline{R}}{\underline{R}})} = \overline{p}^1(0;\cdot), \tag{A.48}$$

we have shown that $\overline{p}^1(\alpha;\cdot)$ is well-defined and continuous at $\alpha = 0$. Therefore, we've shown that $\overline{p}^1(\alpha;\cdot)$ is continuously decreasing in α.

Step3: limits of $\overline{p}^1(\alpha;\beta,\delta,\overline{R},\underline{R})$ as $\alpha \to 1$ and $\alpha \to -\infty$ We can solve for $\overline{p}^1(\alpha;\cdot)$ from $F(\alpha, \overline{p}^1(\alpha;\cdot);\cdot) = (E_{\overline{p}^1(\alpha;\cdot)}[R^\alpha])^{\frac{1}{\alpha}} - \frac{1}{\beta\delta} = 0$:

$$\overline{p}^1(\alpha;\cdot) = \frac{(\frac{1}{\beta\delta})^\alpha - \underline{R}^\alpha}{\overline{R}^\alpha - \underline{R}^\alpha} = \frac{(\frac{(\beta\delta)^{-1}}{\underline{R}})^\alpha - 1}{(\frac{\overline{R}}{\underline{R}})^\alpha - 1}. \tag{A.49}$$

$$
\begin{aligned}
lim_{\alpha\to 1}\overline{p}^1(\alpha;\cdot) &= lim_{\alpha\to 1}\frac{(\frac{1}{\beta\delta})^\alpha - \underline{R}^\alpha}{\overline{R}^\alpha - \underline{R}^\alpha} = \frac{(\frac{1}{\beta\delta}) - \underline{R}}{\overline{R} - \underline{R}} \\
lim_{\alpha\to -\infty}\overline{p}^1(\alpha;\cdot) &= lim_{\alpha\to -\infty}\frac{(\frac{(\beta\delta)^{-1}}{\underline{R}})^\alpha - 1}{(\frac{\overline{R}}{\underline{R}})^\alpha - 1} = 1,
\end{aligned}
$$

since we've assumed that $\underline{R} < \dfrac{1}{\beta\delta} < \overline{R}$.

Bibliography

Amador, M., I. Werning, and G.-M. Angeletos (2006). Commitment vs. flexibility. *Econometrica 74*(2), 365–396.

Aumann, R. J. and M. Maschler (1995). *Repeated games with incomplete information*. MIT press.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance 59*(4), 1481–1509.

Beauchêne, D., J. Li, and M. Li (2019). Ambiguous persuasion. *Journal of Economic Theory 179*, 312–365.

Bénabou, R. and J. Tirole (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics 117*(3), 871–915.

Bergemann, D. and U. Hege (1998). Venture capital financing, moral hazard, and learning. *Journal of Banking & Finance 22*(6-8), 703–735.

Bergemann, D. and U. Hege (2005). The financing of innovation: Learning and stopping. *RAND Journal of Economics*, 719–752.

Bergemann, D. and S. Morris (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics 11*(2), 487–522.

Bergemann, D. and S. Morris (2019). Information design: A unified perspective. *Journal of Economic Literature 57*(1), 44–95.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The annals of mathematical statistics*, 265–272.

Carrillo, J. D. and T. Mariotti (2000). Strategic ignorance as a self-disciplining device. *The Review of Economic Studies 67*(3), 529–544.

Carroll, G. (2015). Robustness and linear contracts. *American Economic Review 105*(2), 536–63.

Dworczak, P. and G. Martini (forthcoming). The simple economics of optimal persuasion. *Journal of Political Economy*.

Gerardi, D. and L. Maestri (2012). A principal–agent model of sequential testing. *Theoretical Economics 7*(3), 425–463.

Gul, F. and W. Pesendorfer (2001). Temptation and self-control. *Econometrica 69*(6), 1403–1435.

Halac, M., N. Kartik, and Q. Liu (2016). Optimal contracts for experimentation. *The Review of Economic Studies 83*(3), 1040–1091.

Hörner, J. and L. Samuelson (2013). Incentives for experimenting agents. *The RAND Journal of Economics 44*(4), 632–663.

Hoshino, T. (2017). Multi-agent persuasion: Leveraging strategic uncertainty. *International Economic Review*.

Jakobsen, A. (2018). Dynamic (in)consistency and the value of information. working paper.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review 101*(6), 2590–2615.

Kolotilin, A. (2018). Optimal information disclosure: A linear programming approach. *Theoretical Economics 13*(2), 607–635.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics 112*(2), 443–478.

Li, F., Y. Song, and M. Zhao (2019). Global manipulation by local obfuscation. *Available at SSRN 3471491*.

Lipnowski, E. and L. Mathevet (2017). Simplifying bayesian persuasion. Technical report, Mimeo, New York University.

Lipnowski, E. and L. Mathevet (2018). Disclosure to a psychological audience. *American Economic Journal: Microeconomics 10*(4), 67–93.

Liu, Q. (2015). Correlation and common priors in games with incomplete information. *Journal of Economic Theory 157*, 49–75.

Manso, G. (2011). Motivating innovation. *The journal of finance 66*(5), 1823–1860.

Mariotti, T., N. Schweizer, N. Szech, and J. von Wangenheim (2018). Information nudges and self-control.

Mathevet, L., J. Perego, and I. Taneva (2020). On information design in games. *Journal of Political Economy 128*(4), 1370–1404.

Shmaya, E. and L. Yariv (2009). Foundations for bayesian updating. *Unpublished*

paper, CalTech.[989].

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies 23*(3), 165–180.