Essays on Information and Communication in Microeconomic Theory

Inaugural-Dissertation
zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften
durch die
Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von
INGA DEIMEN
aus Duisburg

Bonn 2016

Dekan:	Prof. Dr. Rainer Hüttemann
Erstreferent:	Prof. Dr. Dezső Szalay
Zweitreferent:	Prof. Sven Rady, Ph.D.

Tag der mündlichen Prüfung: 03.06.2016

Acknowledgements

Many people supported me during the process of preparing this thesis. First, I would like to thank my supervisor Dezső Szalay not only for countless interesting discussions and continuous feedback, but also for his encouragement and support. Further, I would like to use this opportunity to thank my second supervisor Sven Rady for his advice and encouragement.

I want to thank my coauthors Felix Ketelaar, Mark T. Le Quement, and Julia Wirtz for their effort and patience. I am grateful not only to my colleagues but also to my friends for helpful comments, interesting discussions, and their support. This thesis has greatly benefited from discussions and comments made by the audiences of numerous conferences and seminar participants at various universities.

Moreover, I would like to thank the Bonn Graduate School of Economics (BGSE) for providing an excellent research environment. Financial support by the German Research Foundation (DFG) through the BGSE and through SFB/ TR 15 is gratefully acknowledged.

I thank my family for their understanding and their support, and for being there for me.

Contents

Introduction

Making decisions in an uncertain environment is a complex task and information can help to reach a sound decision. For example, a committee may aggregate individual knowledge about a candidate before deciding whom to hire, a firm may conduct a market study before choosing a product design, or a student may learn over time how much effort is necessary to achieve a success.

This dissertation studies the role of information in three different microeconomic environments and introduces for each an independent theoretical model. In all three models agents face uncertainty about the state of the world and seek information to take actions that maximize their expected outcome. A common feature of all three models is that the state is two-dimensional whereas the action is one-dimensional. As will be seen, this difference allows interesting effects to come into play. In the first two chapters the information has to be communicated between different players. In Chapter 1 members of a committee share their two-dimensional information before voting on a binary outcome. In Chapter 2 one division of a firm obtains two-dimensional information and communicates with another division that has to take the one-dimensional action. In the third chapter a single agent faces two-dimensional uncertainty about her type and the production function. In each period, the agent gathers information by experimentation with the level of effort; the effort choice as well as the output are binary. The remainder of the introduction explains the modeling assumptions and the findings of each chapter in more detail.

In the first chapter, members of a committee deliberate before voting on an outcome. The question is whether truthful communication of individual information is possible, given that committee members have heterogeneous preferences. This chapter, “Consistency and Communication in Committees”, is based on joint work with Felix Ketelaar and Mark T. Le Quement (Deimen et al. (2015)). We generalize the classical binary Condorcet jury model by introducing a richer state and signal space, thereby generating a concern for consistency in the evaluation of aggregate information. For concreteness, a jury aims at determining whether a defendant is guilty or innocent. If guilty, he must have committed the crime at one specific point in time. If innocent, he must have been engaged in some activity at the moment of the crime. Different scenarios constitute different variants of guilt and innocence. Jurors receive private signals, generated by the trial hearing, and before voting whether to acquit or convict the defendant, they share their information. More consistent aggregate information yields stronger evidence and is therefore more convincing.

Formally, we consider a deliberation and voting model in which rich state and signal spaces combine with a binary action space. In our model there are two basic states, each of which splits into a set of sub-states. Each signal is informative with respect to a basic state and a particular sub-state. More consistent signals provide stronger evidence for the corresponding state. Members of a heterogeneous committee communicate via cheap talk before voting on a binary outcome. In contrast to results obtained in the classical binary signal setup (e.g. Coughlan (2000)), we find that the truthful communication and sincere voting equilibrium is virtually always compatible with a positive probability of ex-post conflict among agents. Key is the aspect of consistency in the information structure; individual information is interpreted within the context of aggregate information.

By contrast, in the second chapter individual pieces of information are communicated as aggregate information; key is the choice of the information structure. The second chapter, “Information, Authority, and Smooth Communication in Organizations”, is based on joint work with Dezső Szalay (Deimen and Szalay (2015)). Two divisions, overarched by a headquarters, need to reach a decision that affects the payoffs of all three parties involved. The organization and the roles of the divisions are exogenously fixed: division one gets to observe information about the ideal decisions from both divisions’ perspectives. Division two retains the right to determine the common decision. Hence, the divisions are bound to communicate with each other. Headquarters’ influence on the decision process is reduced to controlling the informational environment that division one gets to observe.

Formally, we study strategic information transmission in a sender-receiver game. We frame our game in a set-up, where the information structure is chosen by a third party, the headquarters who has the objective to maximize the joint surplus of sender and receiver, the two divisions. The uninformed receiver has to take a decision and seeks advice from the informed sender. Optimally, the headquarters chooses the information structure such that equilibrium communication is smooth: the sender precisely suggests his preferred policy and the receiver takes the sender’s advice at face value. We show that communication of unverifiable information without commitment may be a constrained-efficient decision mechanism - unsurpassed even by optimal delegation. The crucial conditions under which this works are that conflicts are stochastic and unsystematic and that information is endogenous.

Opposed to the two previous setups where information is shared between different players, in the third chapter a single agent experiments with the level of effort to gather information. The third chapter, "A Bandit Model of Two-Dimensional Uncertainty – Rationalizing Mindsets" is based on joint work with Julia Wirtz. We analyze a new type of bandit where an agent is confronted with two-dimensional uncertainty. The agent does not know whether ability or effort is required to succeed at a given task. Moreover, the agent does not know her own ability level. In each period, after deciding whether to exert effort or not, the agent observes a success or a failure and updates her beliefs about both the task and her ability accordingly. In contrast to a standard bandit model, the agent gains information even when she is not exerting effort. In our setting we find that different agents react to failure in different ways; while some agents find it optimal to resign others prefer to increase their effort. Given an infinite-time horizon agents may start and stop exerting effort repeatedly.

A possible interpretation of different reactions to failure is given by the educational psychologist Carol Dweck. In particular, Dweck (2006) attributes diverging behavior in response to failure to different mindsets. In this literature, agents that have a "fixed mindset" believe that success is based on innate ability, whereas agents with "growth mindset" believe that success comes from hard work. Consequently, when facing a failure fixed types stop exerting effort whereas growth types start exerting more effort. We show that different effort costs and beliefs about the own ability and the production function together with Bayesian updating can explain the differences in behavior.

1 Consistency and Communication in Committees

1.1 Introduction

This chapter considers a deliberation and voting model in which rich state and signal spaces combine with a binary action space. A committee consisting of privately informed agents with known heterogeneous preferences engages in simultaneous information exchange prior to voting. Our information structure generates a concern for consistency in the aggregation of individual signals; a given signal is interpreted differently depending on how it matches other available evidence. We find that in contrast to the classical model featuring binary state and signal spaces, full information sharing and sincere voting can constitute an equilibrium although agents with some probability disagree ex post.

Consider the example of a jury aiming at determining whether a defendant is guilty or innocent. If guilty, he must have committed the crime at one specific point in time, for example on one particular day of a given week. If innocent, he must have been engaged in some activity at the moment of the crime; for example working, watching TV, or doing sports. Different days of the week constitute mutually exclusive variants of the guilty state while different activities constitute mutually exclusive variants of the innocent state.

Jurors gather evidence through a trial hearing which generates private signals. Before deciding whether to acquit or convict, jurors retire to deliberate and share their private signals. Consider two possible scenarios. In the first scenario, half of the jurors received a signal indicating that the defendant committed the crime on Monday, while the other half received a signal indicating that he committed the crime on Wednesday. In the second scenario, all jurors received a signal indicating Monday. The latter scenario is more consistent than the first and therefore provides more convincing evidence of guilt. Jurors do not as such care about the time at which the crime was committed but wish to establish with sufficient certainty whether the defendant is guilty or innocent. More consistent profiles yield stronger evidence.

The core elements of the above description apply to many other situations. Consider a group of investment bankers that contemplates investing in shares of a large manufacturer, e.g. Chrysler. Committee members need to assess whether Chrysler will avoid bankruptcy in the near future. This may happen if either the US Federal State provides a bailout package or if some private company (e.g. Fiat) decides to step in. On the other hand, if Chrysler does go bankrupt, this may happen according to different chapters of the bankruptcy code. Another example is that of a board of directors that seeks to predict whether a Democrat or a Republican will win the next US presidential election. Different Democratic (Republican) candidates constitute different variants of the Democratic (Republican) state.

We incorporate the key features of the above examples into a model of collective decision making. There are two basic states, each of which splits into a set of substates. Each signal is informative with respect to a basic state and a particular substate. In this context, the consistency of signals matters as illustrated above; more consistent signals provide stronger evidence for the corresponding state. Members of a heterogeneous committee communicate via cheap talk before voting on a binary outcome. In contrast to results obtained in the classical binary signal setup,[1] we find that the truthful communication and sincere voting equilibrium (TS equilibrium) is virtually always compatible with a positive probability of ex post conflict among agents.

The intuition for our result comes out clearly when compared to the classical binary signal model.[2] In the latter, in the putative TS equilibrium, pivotality at the communication stage pins down uniquely the information held by remaining committee members. Disagreement about the optimal decision rule implies that there is always at least one agent for whom this pivotal profile implies a suboptimal decision on the equilibrium path. Consequently, this agent profitably deviates and bends the decision rule in his favored direction.

In our model, the set of pivotal profiles is not a singleton anymore: different signal profiles can yield similar posteriors because the conditional probability of guilt depends on two aspects, the total number of signals indicating respectively guilt or innocence as well as the consistency of signals within these subsets. A smaller total number of guilty signals can be compensated by a higher degree of consistency among guilty signals. The multiplicity of pivotal scenarios in our model allows two effects to come into play. First, there exist pivotal profiles for which all agents agree

---

1 See Coughlan (2000), Austen-Smith and Feddersen (2006), Meirowitz (2007), Van Weelden (2008).

2 See in particular Coughlan (2000).

with the decision taken on the equilibrium path. At these profiles a deviation is disadvantageous (*consensus effect*). Second, the impact of an agent's announcement depends on the signals of the whole committee. He faces uncertainty as to which announcement is more or less consistent with other agents' signals and is thus suitable to shift the outcome in a desired direction (*uncertainty effect*).

Building on the theory of strategic voting as information aggregation (E.g. Austen-Smith and Banks (1996), Feddersen and Pesendorfer (1998), Bhattacharya (2013)), different classes of contributions have analyzed communication in heterogeneous committees by modifying the baseline model (Coughlan (2000)). Austen-Smith and Feddersen (2006), Meirowitz (2007), and Le Quement (2013) examine the implications of preference uncertainty. Van Weelden (2008), Hummel (2012), and Le Quement and Yokeeswaran (2015) analyze alternative communication protocols. Gerardi et al. (2009), Gerardi and Yariv (2007), and Wolinsky (2002) adopt a mechanism design approach. Additionally, experimental work has analyzed communication behavior in groups (E.g. Goeree and Yariv (2011) and Dickson et al. (2008)). Our contribution lies in the introduction of a novel information structure that captures the idea that individual information is interpreted within the context of aggregate information.

1.2 Model

A jury of n agents, $n \in \mathbb{N}$, $n \geq 3$ has to decide whether to acquit (A) or convict (C) a defendant. The defendant is either innocent (I) or guilty (G). Both innocence and guilt occur in finitely many different variants $i_1, ..., i_{m_I}$ and $g_1, ..., g_{m_G}$, respectively, with $m_I, m_G \in \mathbb{N}$, $m_I, m_G \geq 2$.[3] The state space is hence given as $\Omega = I \cup G$ where $I = \{i_1, ..., i_{m_I}\}$ and $G = \{g_1, ..., g_{m_G}\}$. We denote the state by ω and say that the defendant is innocent if $\omega \in I$ and guilty if $\omega \in G$. The state is drawn from a publicly known prior distribution such that $P(\omega \in I) + P(\omega \in G) = 1$, $P(i_l) = \frac{P(\omega \in I)}{m_I}$ $\forall l \in \{1, ..., m_I\}$, and $P(g_l) = \frac{P(\omega \in G)}{m_G}$ $\forall l \in \{1, ..., m_G\}$.

The jury implements an action $a \in \{A, C\}$ by voting according to a voting rule $k \in \{1, \ldots, n\}$. Each agent $j \in \{1, \ldots, n\}$ casts a vote in favor of one of the two actions. If the number of votes cast for conviction is greater than or equal to k, the defendant is convicted while otherwise he is acquitted.

---

3 The case of $m_I = m_G = 1$ corresponds to the classical model analyzed by Coughlan (2000) and others.

The utility of agent j from action a conditional on $\omega \in \Omega$ is given by

$$u_j(a,\omega) = \begin{cases} 0 & \text{if } (a,\omega) \in \{(A,I)\,(C,G)\} \\ -q_j & \text{if } (a,\omega) = (C,I) \\ -\left(1-q_j\right) & \text{if } (a,\omega) = (A,G). \end{cases}$$

The commonly known preference parameter $q_j \in (0,1)$ characterizes the relative importance assigned to the two types of errors.[4] As agent j maximizes expected utility, he prefers conviction over acquittal if and only if the probability of the defendant being guilty exceeds the cut-off q_j. We mainly focus on the case of a committee featuring two preference types $q_H < q_D$ referred to as hawks and doves.

Prior to the voting stage, each agent receives a private signal $s \in S = \Omega$. Signals are i.i.d. across agents conditional on ω. If $\omega = i_l$ for some $l \in \{1,...,m_I\}$, then

$$\begin{aligned} P(s = i_l|\omega = i_l) &= \lambda \cdot \frac{p}{\lambda + (m_I - 1)}, \\ P(s = i_r|\omega = i_l) &= \frac{p}{\lambda + (m_I - 1)} \quad \forall r \in \{1,...,m_I\},\ r \neq l, \\ P\left(s = g_r|\omega = i_l\right) &= \frac{1-p}{m_G} \quad \forall r \in \{1,...,m_G\}, \end{aligned}$$

with $p \in \left(\frac{1}{2},1\right)$ and $\lambda > 1$. If $\omega = g_l$ for some $l \in \{1,...,m_G\}$, respective expressions apply after permuting i and g as well as I and G.[5] Applying Bayes' law, p measures up to priors the probability that the signal correctly reveals whether the defendant is innocent or guilty. $\lambda = \frac{P(s=i_l|\omega=i_l)}{P(s=i_l|\omega=i_r)} = \frac{P(s=g_l|\omega=g_l)}{P(s=g_l|\omega=g_r)}$ measures the relative informativeness of signals with respect to the particular variant of innocence or guilt. We finally assume that

$$\begin{aligned} \frac{P(s = i_l|\omega = i_r)}{P\left(s = i_l|\omega = g_t\right)} &\geq 1 \quad \forall l,r \in \{1,...,m_I\},\ t \in \{1,...,m_G\} \quad \text{and} \\ \frac{P\left(s = g_l|\omega = g_r\right)}{P\left(s = g_l|\omega = i_t\right)} &\geq 1 \quad \forall l,r \in \{1,...,m_G\},\ t \in \{1,...,m_I\}, \end{aligned}$$

which is equivalent to requiring a lower bound $p \geq \max\left\{\frac{m_G-1+\lambda}{2m_G-1+\lambda}, \frac{m_I-1+\lambda}{2m_I-1+\lambda}\right\}$ on the

---

4 Existing results (Austen-Smith and Feddersen (2006)) indicate that truthful communication is easier to achieve if there is uncertainty about preference types. By assuming observable preference types, we isolate the truth-telling incentives that are specific to our model.

5 Note that whenever $\lambda > 1$ our signal generating process is not reducible to a process that generates i.i.d. signals conditional on I and G. It is however reducible to a process that generates correlated signals conditional on I and G.

informativeness of signals with respect to innocence and guilt.

A collection of signals constitutes a *signal profile* $\sigma = \left(x_1, ..., x_{m_I}, y_1, ..., y_{m_G}\right)$ where x_r denotes the number of i_r-signals, $r \in \{1, ..., m_I\}$ and y_l denotes the number of g_l-signals, $l \in \{1, ..., m_G\}$. Conditional on signal profile σ, the posterior probability of guilt is given by

$$\beta(\sigma) \equiv P(\omega \in G|\sigma) = \frac{P(\omega \in G) \cdot P(\sigma|\omega \in G)}{P(\omega \in G) \cdot P(\sigma|\omega \in G) + P(\omega \in I) \cdot P(\sigma|\omega \in I)}.$$

In terms of utilities agents only care whether $\omega \in I$ or $\omega \in G$, hence the number $\beta(\sigma)$ is a sufficient statistic for the favored action of each individual agent for any signal profile σ.

The timing of the game is as follows. Nature draws the state ω. Each agent receives a private signal s. Subsequently, agents simultaneously send a public cheap talk message $m \in M = S$.[6] Finally, each agent casts a vote, the action $a \in \{A, C\}$ is implemented according to the voting rule, and payoffs are realized.

Our equilibrium concept is Perfect Bayesian Equilibrium. We focus on the so-called *TS equilibrium*: agents *truthfully* reveal their private information by sending a message $m = s$ at the communication stage and (correctly) believe that others also communicate truthfully. Subsequently, agents vote *sincerely*, i.e. they vote for conviction if and only if they favor conviction given their beliefs.

1.3 Simple Example

Consider a three persons committee consisting of two doves and one hawk. The voting rule is $k = 2$, i.e. simple majority. Let $I = \{i_1, i_2\}$ and $G = \left\{g_1, g_2\right\}$. Referring to the example given in the introduction, think of i_1 as "innocent and working", i_2 as "innocent and doing sports", g_1 as "guilty on Monday", and g_2 as "guilty on Wednesday". Aggregate signal profiles are ordered as follows

$$\begin{matrix} \beta(3,0,0,0) \\ \beta(0,3,0,0) \end{matrix} < \ldots < \begin{matrix} \beta(1,0,0,2), \beta(1,0,2,0) \\ \beta(0,1,0,2), \beta(0,1,2,0) \end{matrix} < \begin{matrix} \beta(0,0,1,2) \\ \beta(0,0,2,1) \end{matrix} < \begin{matrix} \beta(0,0,0,3) \\ \beta(0,0,3,0) \end{matrix}.$$

Suppose q_H, q_D are such that a dove favors conviction if and only if the aggregate signal profile is either $(0,0,0,3)$ or $(0,0,3,0)$ while a hawk in addition favors conviction for profiles $(0,0,1,2)$ and $(0,0,2,1)$, so that the committee disagrees for these two profiles. While both types require three g-signals to prefer conviction, doves fur-

[6] See end of Section 1.3 for a comment on sequential communication.

thermore require these g-signals to be consistent. In this setting the TS equilibrium exists, as we now show.

An agent never has an incentive to deviate from sincere voting as this weakly decreases the probability that his favored decision ensues. This holds true independently of his announcement at the communication stage. Moreover, given that the voting rule is simple majority doves can always enforce their favored decision and thus have no incentive to deviate from truth-telling.

We now analyze the truth-telling incentives of the hawk. The hawk's announcement is pivotal if the remaining two agents hold signal profiles $(0,0,2,0)$ or $(0,0,0,2)$. In the first (second) case, a g_1- (g_2-) announcement triggers conviction while any of the remaining announcements causes acquittal. Given the symmetry of the model, conditions ensuring truth-telling when the hawk holds an i_1- or an i_2-signal are identical modulo an exchange of subscripts. We can therefore without loss of generality focus on deviation incentives conditional on an i_1-signal. An equivalent argument applies when the hawk holds a g_1- or a g_2-signal.

Let the hawk hold an i_1-signal and be pivotal at the communication stage. The aggregate signal profile is then $(1,0,2,0)$ or $(1,0,0,2)$. In either case, the committee decision given the true signal profile is acquittal and coincides with the decision favored by the hawk. Accordingly, he has no incentive to deviate from truth-telling. Here, the consensus effect is the source of truth-telling: a hawk holding an i_1-signal agrees with the doves on the preferred action in all pivotal scenarios.

Assume next that the hawk holds a g_1-signal and is pivotal at the communication stage. The signal profile of the entire committee is then either $(0,0,1,2)$ or $(0,0,3,0)$. The hawk disagrees with the acquittal ensuing from truth-telling at $(0,0,1,2)$ while he agrees with the conviction ensuing from truth-telling at $(0,0,3,0)$. If the hawk deviates to announcing some i-signal, the signal profile observed by others at the voting stage is either $(1,0,2,0)$, $(1,0,0,2)$, $(0,1,0,2)$ or $(0,1,2,0)$, in all cases leading to an undesired acquittal. If the hawk deviates to a g_2-announcement, the signal profile observed by the remaining agents at the voting stage is $(0,0,0,3)$ or $(0,0,2,1)$. The deviation beneficially overturns an acquittal in the first case but adversely overturns a conviction in the second case. The hawk thus faces uncertainty about the impact of his statement. Among the two pivotal profiles, $(0,0,0,2)$ incentivizes lying while $(0,0,2,0)$ incentivizes truth-telling. We call this the uncertainty effect. Which incentive dominates depends on the relative likelihood assigned to these two profiles, the latter itself depending on the probability assigned to the states g_1 and g_2. An agent holding a g_1-signal assigns a higher probability to state g_1 than to state g_2 and accordingly to profile $(0,0,2,0)$ than to profile $(0,0,0,2)$. The signal profile that incentivizes truth-telling is thus considered more likely than the one that incentivizes lying. Hence the hawk does not prefer to announce a g_2-signal. We conclude that the TS equilibrium exists despite the existence of signal profiles generating ex post conflict.

We close with two remarks. By the same arguments as above, the TS equilibrium also exists under unanimity when $k = 3$. Moreover, the TS equilibrium continues to exist under sequential communication if the hawk speaks first. Indeed, the hawk's incentives then replicate those arising under simultaneous communication while doves still determine the outcome and hence have no incentives to deviate.[7]

1.4 Analysis of the TS Equilibrium

The example of Section 1.3 shows that the TS equilibrium can exist despite potential disagreement after full pooling of information. In what follows, we provide an equilibrium analysis for the general model.

For any signal $s \in S$, let σ_s denote the signal profile that consists of one signal s only. Moreover, for a given agent j, we denote the signal profile of all other agents by σ_{-j}. The following lemma addresses the effect of shifting mass from one entry of σ to another. This replicates the change in other agents' beliefs achievable by misreporting a signal in the putative TS equilibrium.

Lemma 1.1. *For any signal profile $\sigma = \left(x_1, ..., x_{m_I}, y_1, ..., y_{m_G}\right)$, the function $\beta(\sigma)$ is invariant under any permutation of x-entries and any permutation of y-entries of σ. Moreover, the following inequalities hold:*

$$\beta\left(\sigma + \sigma_{g_r}\right) > \beta\left(\sigma + \sigma_{i_l}\right) \quad \forall\, l \in \{1, ..., m_I\}, r \in \{1, ..., m_G\}, \tag{1.4.1}$$

$$\beta\left(\sigma + \sigma_{g_l}\right) \geq \beta\left(\sigma + \sigma_{g_r}\right) \quad \forall\, l, r \in \{1, ..., m_I\},\ y_r \leq y_l, \tag{1.4.2}$$

$$\beta\left(\sigma + \sigma_{i_l}\right) \leq \beta\left(\sigma + \sigma_{i_r}\right) \quad \forall\, l, r \in \{1, ..., m_I\},\ x_r \leq x_l. \tag{1.4.3}$$

Conditions (1.4.2) *and* (1.4.3) *hold with equality if and only if respectively $y_l = y_r$ and $x_l = x_r$.*

Proof. See Appendix. □

Lemma 1.1 shows that three factors determine the posterior probability of guilt; an increase in the total number of g-signals and in the consistency of the profile of

[7] The existence of the TS equilibrium under unanimity and sequential communication stands in contrast to the impossibility results of Austen-Smith and Feddersen (2006) and Van Weelden (2008).

g-signals leads to an increase in the posterior probability of guilt. An increase in the consistency of the profile of i-signals has the opposite effect.

Note that if $\lambda = 1$ or $m_G = m_I = 1$ the impossibility result shown in Coughlan (2000) applies: the TS equilibrium exists if and only if either 1) at least k agents ($n-k$ agents) favor conviction (acquittal) for any realization of signals or 2) all agents favor the same action for any realization of signals. The existence of the TS equilibrium beyond these trivial cases thus requires both an enlarged state and signal space as well as the assumption that signals are informative with respect to the variant of innocence or guilt that applies.

Introducing terminology, we say that hawks have *critical mass* if the number of hawks is weakly greater than k, so that hawks are sufficiently many to impose conviction whenever they wish. Otherwise, doves have critical mass. We call a signal profile σ a *conflict profile* if conditional on σ hawks and doves favor different actions, that is if $q_H < \beta(\sigma) < q_D$. We impose the following simple assumption on preferences of hawks and doves.

Assumption 1.1 (No partisans). *The preferred action of each agent depends on the aggregate signal profile. Moreover, hawks require less than the maximal possible evidence of guilt to prefer conviction and doves require less than the maximal possible evidence of innocence to prefer acquittal.*

Let q denote the threshold of the type that has critical mass and consider an agent j of the type that does not have critical mass. We denote the set of signal profiles σ_{-j} at which an i_r-report by agent j triggers an acquittal while a g_l-report triggers a conviction by

$$Piv_{i_r,g_l}(q) \equiv \{\sigma_{-j} : \beta(\sigma_{-j} + \sigma_{i_r}) < q \wedge \beta(\sigma_{-j} + \sigma_{g_l}) \geq q\}.$$

Our main result reads as follows.

Theorem 1.1. *Let k be non-unanimous and impose Assumption 1.1.*
a) Assume hawks have critical mass. The TS equilibrium exists if and only if

$$q_D \leq \hat{q}_D(q_H) \quad \equiv \quad \frac{1}{P(\sigma_{-j} \in Piv_{i_r,g_l}(q_H)|s_j = g_l)} \sum_{\sigma_{-j} \in Piv_{i_r,g_l}(q_H)} P(\sigma_{-j}|s_j = g_l) \cdot \beta(\sigma_{-j} + \sigma_{g_l}).$$

b) Assume doves have critical mass. The TS equilibrium exists if and only if

$$q_H \geq \hat{q}_H(q_D) \quad \equiv \quad \frac{1}{P(\sigma_{-j} \in Piv_{i_r,g_l}(q_D)|s_j = i_r)} \sum_{\sigma_{-j} \in Piv_{i_r,g_l}(q_D)} P(\sigma_{-j}|s_j = i_r) \cdot \beta(\sigma_{-j} + \sigma_{i_r}).$$

c) If hawks have critical mass and $q_D = \hat{q}_D(q_H)$ there exists at least one conflict profile. If doves have critical mass and $q_H = \hat{q}_H\left(q_D\right)$ there exists at least one conflict profile.

Proof. See Appendix. □

Theorem 1.1 provides a general existence result for the TS equilibrium. Part a) states the existence of a critical dove type $\hat{q}_D(q_H)$ such that the TS equilibrium exists if and only if $q_D \in \left(q_H, \hat{q}_D(q_H)\right]$. The threshold $\hat{q}_D(q_H)$ corresponds to the probability of guilt conditional on all pivotal profiles σ_{-j} where a truthful g_l-report of agent j leads to conviction while an i_r-report leads to acquittal. Part b) states the corresponding result for the case where doves have critical mass. Part c) yields the fundamental qualitative statement that the TS equilibrium is compatible with the existence of conflict profiles. It stands in stark contrast to Coughlan's impossibility result.

We outline the main steps of the proof of Theorem 1.1 in what follows. A first observation is that agents of the type that has critical mass never have an incentive to deviate as their preferred action given aggregate information is always implemented. A second observation is that under a non-unanimous voting rule, an agent of the type that does not have critical mass is never pivotal at the voting stage, irrespective of whether he deviated at the communication stage. Assuming hawks have critical mass, given sincere voting a dove's vote can only influence the outcome if hawks vote for acquittal. The reciprocal argument holds for the case of doves having critical mass. In both cases, however, *all* other agents will vote unanimously and thus the agent's vote cannot be pivotal. A third observation is that a hawk never has an incentive to misreport a g-signal as an i-signal while a dove never has an incentive to misreport an i-signal as a g-signal. This immediately follows from (1.4.1). A deviation of the above described type would only worsen the implemented decision rule.

A fourth observation is that no agent has an incentive to misreport an i-signal as a different i-signal or to misreport a g-signal as a different g-signal. Here, the uncertainty effect is key. Consider a juror holding signal i_r and contemplating announcing $i_l \neq i_r$ instead. The set of signal profiles σ_{-j} splits into pairs of profiles that are identical up to a simple permutation of the numbers of i_r- and i_l-signals. If both reports trigger identical actions for both these profiles, the reporting decision is irrelevant. Otherwise, a truthful i_r-report will lead to an acquittal for the profile that features more i_r-signals and lead to a conviction for the other profile. Deviating to an i_l -report overturns both outcomes. If the defendant is guilty, both profiles are equally likely to occur and agent j is thus indifferent between the two reports. If the defendant is innocent, the profile that is more consistent with agent j's own signal is more likely than the other one given $\lambda > 1$. Hence agent j has an incentive to trigger acquittal for the former profile rather than the latter and thus to tell the truth. Deviations from one g-signal to another are ruled out by a similar argument.

Given the four above observations, the only deviations that remain to be excluded involve doves reporting an i-signal instead of a g-signal and hawks reporting a g-signal instead of an i-signal. These deviations have a clear effect on the outcomes via (1.4.1). Here, the multiplicity of pivotal profiles allows the consensus effect to provide truth-telling incentives. While hawks and doves have diverging interests for some pivotal profiles, their preferred outcome coincides for others. This (partial) consensus is more pronounced the smaller $q_D - q_H$. Accordingly, we get an upper bound for q_D in Part a) and a lower bound for q_H in Part b). As for Part c), a type who is indifferent between truth-telling and lying necessarily faces pivotal profiles that incentivize truth-telling and pivotal profiles that incentivize deviating. For the sake of concreteness, consider a dove holding a g_l-signal and let hawks have critical mass. Let σ be a signal profile such that a hawk prefers conviction precisely for those signal profiles that yield at least as much evidence for the defendant being guilty as σ does. Whenever $q_D > \beta(\sigma)$, the profile σ is a conflict profile and incentivizes lying. On the other hand, Assumption 1.1 guarantees the existence of another signal profile $\tilde{\sigma}$ that satisfies $\beta(\tilde{\sigma}) > \beta(\sigma)$ and $\beta\left(\tilde{\sigma} - \sigma_{g_l} + \sigma_{i_r}\right) < \beta(\sigma)$. An example of such a profile $\tilde{\sigma}$ is one that has the same total number of i- and g-signals as σ but is either slightly less consistent with respect to its i-signals or slightly more consistent with respect to its g-signals. Now, if $\beta(\tilde{\sigma}) > q_D > \beta(\sigma)$ then $\tilde{\sigma}$ incentivizes truth-telling and this incentive dominates the deviation incentive from σ if q_D is sufficiently close to $\beta(\sigma)$.[8]

We finally sketch how our results generalize to arbitrary preference types. For $j \in \{1, ..., n\}$ let q_j be juror j's preference parameter and assume without loss of generality $q_1 \leq ... \leq q_n$. Let $k \in \{2, ..., n-1\}$ denote the voting rule. The TS equilibrium then implements juror k's optimal decision rule. For any juror $j < k$, the implemented decision rule is (weakly) "dovish" as $q_j \leq q_k$ while for any juror $j > k$ the decision rule is (weakly) "hawkish" as $q_j \geq q_k$. The insights from the two type case suggest that the TS equilibrium exists if and only if for all $j \in \{1, ..., n\}$ we have

[8] Numerical simulations show that the parameter area that is compatible with the existence of the TS equilibrium in our model is typically larger than in the binary model in Coughlan (2000). Moreover, the number of conflict profiles compatible with the TS equilibrium becomes large when committee size increases, contrasting e.g. Le Quement (2013).

$q_j \in [\hat{q}_H(q_k), \hat{q}_D(q_k)]$ with $\hat{q}_H(\cdot)$ and $\hat{q}_D(\cdot)$ defined as in Theorem 1.1.

The problem is that a juror is now pivotal at the voting stage if exactly $k-1$ other jurors favor conviction given the aggregate signal profile.[9] Consider a juror $j < k$ holding an i-signal. If this juror reports a g-signal instead, an irreversible conviction is triggered only if in addition to juror k also juror $k+1$ prefers to convict based on the reported evidence. Indeed, if only jurors 1 to k prefer conviction based on reports, juror j can veto a conviction. As a consequence, for profiles σ_{-j} where a particular g-announcement causes jurors 1 to k to favor a conviction based on reported evidence, a deviation to this announcement can only be advantageous as juror j can implement his favored decision at the voting stage.

Given these considerations, the following result holds. The TS equilibrium exists if and only if $q_j \in \left[\underline{q}(q_k, q_{k+1}), \overline{q}(q_k, q_{k-1})\right]$ for all $j \in \{1, ..., n\}$ where the noteworthy aspect is that the bounds now depend on two preference types instead of only one. Furthermore, if jurors k and $k+1$ favor the same action for each signal profile we have $\underline{q}(q_k, q_{k+1}) = \hat{q}_H(q_k)$. Likewise, if jurors k and $k-1$ favor the same action for each signal profile we have $\overline{q}(q_k, q_{k-1}) = \hat{q}_D(q_k)$. It follows that if jurors $k-1$, k, and $k+1$ share the same optimal decision rule the TS equilibrium exists if and only if $q_j \in [\hat{q}_H(q_k), \hat{q}_D(q_k)]\ \forall j \in \{1, ..., n\}$. These insights provide some guidance regarding the optimal composition of heterogeneous committees with an eye to maximizing truth-telling incentives. In a committee too polarized for the TS equilibrium to exist, the inclusion of moderate agents endowed with decision power through a suitably chosen voting rule can help to overcome lying incentives. However, a single moderate agent will not suffice to ensure truth-telling.

1.5 Conclusion

In our collective decision model with pre-vote communication, a positive probability of ex post disagreement among agents is frequently compatible with the existence of the truthful communication and sincere voting equilibrium. The driving forces underlying our positive result are the consensus and uncertainty effects, both of which originate in the multiplicity of pivotal scenarios at the communication stage. The latter feature follows from the role played by consistency given our information structure. From a conceptual perspective, the key and novel feature of our information structure is that a given signal is interpreted in the light of other available information; meaning is determined in context. We find this aspect worth

[9] This is ruled out in the two-type setup by excluding unanimous voting rules.

exploring within other communication games.

1.6 Appendix

Proof of Lemma 1.1. We write $x \equiv \sum_{t=1}^{m_I} x_t$ and $y \equiv \sum_{t=1}^{m_G} y_t$. The statements follow immediately from $p \geq \max\left\{\frac{m_G-1+\lambda}{2m_G-1+\lambda}, \frac{m_I-1+\lambda}{2m_I-1+\lambda}\right\}$ and

$$\beta(\sigma) = \left[1 + \frac{P(\omega \in I)}{P(\omega \in G)} \cdot \frac{m_G}{m_I} \cdot \left(\frac{p \cdot m_I}{\left(1-p\right) \cdot (\lambda + m_I - 1)}\right)^x \cdot \left(\frac{\left(1-p\right) \cdot (\lambda + m_G - 1)}{p \cdot m_G}\right)^y \cdot \frac{\sum_{r=1}^{m_I} \lambda^{x_r}}{\sum_{l=1}^{m_G} \lambda^{y_l}}\right]^{-1}.$$

Note that if $\lambda = 1$ then $\beta(\sigma)$ only depends on x and y. □

Proof of Theorem 1.1. Observations 1 to 3 from the main text are obvious. To show observation 4, suppose first that agent j who is not of the type having critical mass holds a signal $s = i_l$ and considers to report $m = i_r$ with $r \neq l$, $l, r \in \{1, ..., m_I\}$. Consider two candidates for σ_{-j}, namely $\hat{\sigma} = \left(x_1, ..., x_l, ..., x_r, ..., x_{m_I}, y_1, ..., y_{m_G}\right)$ and $\hat{\sigma}_{x_l \longleftrightarrow x_r} = \left(x_1, ..., x_r, ..., x_l, ..., x_{m_I}, y_1, ..., y_{m_G}\right)$ and assume without loss of generality that $x_l \geq x_r$. We compare the expected utility of the reports $m = i_l$ and $m = i_r$ conditional on $\sigma_{-j} \in \left\{\hat{\sigma}, \hat{\sigma}_{x_l \longleftrightarrow x_r}\right\}$. If both reports $m = i_l$ and $m = i_r$ trigger identical actions, the reporting decision does not matter. In particular, this is the case if $x_l = x_r$ by (1.4.3). If the reports trigger different actions, then $x_l > x_r$ and thus $m = i_l$ will trigger acquittal for $\sigma_{-j} = \hat{\sigma}$ and conviction for $\sigma_{-j} = \hat{\sigma}_{x_l \longleftrightarrow x_r}$ while $m = i_r$ will trigger conviction for $\sigma_{-j} = \hat{\sigma}$ and acquittal for $\sigma_{-j} = \hat{\sigma}_{x_l \longleftrightarrow x_r}$, again by (1.4.3). Hence

$$\begin{aligned}
& Eu\left[m = i_l | \sigma_{-j} \in \left\{\hat{\sigma}, \hat{\sigma}_{x_l \longleftrightarrow x_r}\right\}\right] - Eu\left[m = i_r | \sigma_{-j} \in \left\{\hat{\sigma}, \hat{\sigma}_{x_l \longleftrightarrow x_r}\right\}\right] \\
= & -\sum_{t=1}^{m_G} P\left(\omega = g_t | s = i_l\right) \cdot P\left(\sigma_{-j} = \hat{\sigma} | \omega = g_t, \sigma_{-j} \in \left\{\hat{\sigma}, \hat{\sigma}_{x_l \longleftrightarrow x_r}\right\}\right) \cdot \left(1-q\right) \\
& + \sum_{t=1}^{m_G} P\left(\omega = g_t | s = i_l\right) \cdot P\left(\sigma_{-j} = \hat{\sigma}_{x_l \longleftrightarrow x_r} | \omega = g_t, \sigma_{-j} \in \left\{\hat{\sigma}, \hat{\sigma}_{x_l \longleftrightarrow x_r}\right\}\right) \cdot \left(1-q\right) \\
& - \sum_{t=1}^{m_I} P\left(\omega = i_t | s = i_l\right) \cdot P\left(\sigma_{-j} = \hat{\sigma}_{x_l \longleftrightarrow x_r} | \omega = i_t, \sigma_{-j} \in \left\{\hat{\sigma}, \hat{\sigma}_{x_l \longleftrightarrow x_r}\right\}\right) \cdot q \\
& + \sum_{t=1}^{m_I} P\left(\omega = i_t | s = i_l\right) \cdot P\left(\sigma_{-j} = \hat{\sigma} | \omega = i_t, \sigma_{-j} \in \left\{\hat{\sigma}, \hat{\sigma}_{x_l \longleftrightarrow x_r}\right\}\right) \cdot q \\
= & \frac{\frac{P(\omega \in I)}{m_I} \cdot \frac{p}{\lambda - 1 + m_I}}{\frac{P(\omega \in I)}{m_I} \cdot p + \frac{P(\omega \in G)}{m_G} \cdot \left(1-p\right)} \cdot \frac{-\lambda^{x_r+1} - \lambda^{x_l} + \lambda^{x_r} + \lambda^{x_l+1}}{\lambda^{x_r} + \lambda^{x_l}} \cdot q \\
> & \; 0.
\end{aligned}$$

As the set of signal profiles possibly held by other agents splits into pairs of the form

$\{\hat{\sigma}, \hat{\sigma}_{x_l \longleftrightarrow x_r}\}$ this shows that the proposed deviation is not profitable. Deviations from one g-signal to another are ruled out in the same way. It remains to analyze under which circumstances a dove holding some g-signal wants to deviate by reporting some i-signal instead (Part a) and under which circumstances a hawk holding some i-signal wants to deviate by reporting some g-signal (Part b).

a) + b) Assume that agent j is a dove holding signal $s_j = g_l$ for some $l \in \{1, ..., m_G\}$ and considers reporting $m = i_r$ for some $r \in \{1, ..., m_I\}$. By (1.4.1), for any profile $\sigma_{-j} \in Piv_{i_r, g_l}(q_j)$ a truthful report $m = g_l$ will trigger conviction while reporting $m = i_r$ will trigger acquittal. Truthful reporting hence constitutes an equilibrium iff

$$\begin{aligned} 0 &\leq Eu(m_j = g_l) - Eu(m_j = i_r) \\ &= \sum_{\sigma_{-j} \in Piv_{i_r, g_l}(q_H)} P(\sigma_{-j} | s_j = g_l) \cdot \beta(\sigma_{-j} + \sigma_{g_l}) - P(\sigma_{-j} \in Piv_{i_r, g_l}(q_H) | s_j = g_l) \cdot q_D \end{aligned}$$

which proves Part a). Part b) follows similarly.

c) Consider the case where hawks have critical mass. Let σ_H be a profile such that $\beta(\sigma_H) \geq q_j$ and $\beta(\tilde{\sigma}) < q_j$ for all signal profiles $\tilde{\sigma}$ satisfying $\beta(\tilde{\sigma}) < \beta(\sigma_H)$. Such a profile exists by Assumption 1.1. We need to show that $\hat{q}_D(q_H) > \beta(\sigma_H)$ in which case σ_H is a conflict profile. Suppose agent j is a dove holding a g_1-signal and considers deviating by reporting $m_j = i_1$. By Part a) it suffices to show that there exists $\sigma_{-j} \in Piv_{i_1, g_1}(q_H)$ such that

$$\beta(\sigma_{-j} + \sigma_{i_1}) < \beta(\sigma_H) < \beta(\sigma_{-j} + \sigma_{g_1}). \tag{A1.1}$$

After reshuffling x- and y-entries, we may assume without loss of generality that σ_H satisfies $x_1 \geq ... \geq x_{m_I}$, $y_1 \geq ... \geq y_{m_G}$. First, assume that $y_2 > 0$. Then the profile $\sigma_{-j} = \sigma_H - \sigma_{g_2}$ satisfies (A1.1). Similarly, if $x_2 > 0$ then profile $\sigma_{-j} = \sigma_H - \sigma_{i_2}$ satisfies (A1.1). Suppose $x_2 = y_2 = 0$. If $y_1 = 0$ then $x_1 = n$ and $\sigma_H = (n, 0, ..., 0)$, so hawks would want to convict irrespective of any information. If $x_1 = 0$ then $y_1 = n$ and $\sigma_H = (0, ..., 0, y_1 = n, 0, ..., 0)$, so hawks would want to convict only given the maximal possible evidence of guilt. Both cases contradict Assumption 1.1. Finally, assume $x_1 \neq 0 \neq y_1$. Since $n \geq 3$ we must have $x_1 \geq 2$ or $y_1 \geq 2$. In the former case, $\sigma_{-j} = \sigma_H - \sigma_{i_1} - \sigma_{g_1} + \sigma_{i_2}$ satisfies (A1.1) while in the latter case $\sigma_{-j} = \sigma_H - \sigma_{i_1} - \sigma_{g_1} + \sigma_{g_2}$ satisfies (A1.1). This concludes the proof for the case of hawks having critical mass. The proof for the case of doves having critical mass is alike. □

2 Information, Authority, and Smooth Communication in Organizations

2.1 Introduction

Sound decision making requires good information. The success of organizations depends crucially on the quality of information their decision-makers have and on the alignment of interests within the organization. Most of the time, organizations do not have automatic access to information but must actively acquire it prior to decision-making. The search for information is subject to choices and must be considered as part of the decision-making process. How and where the information enters the organization is by and large determined by the organization's existing structure. Inside the organization the information needs to be communicated to the decision-maker. Such communication is prone to strategic manipulation; on the way towards the decision-maker, inferences are drawn, details can be dropped, things can be swept under the rug. The present chapter tries to shed light on how organizations with given communication channels can cope with such problems. We show that an appropriate acquisition of information can ensure sound decision-making despite strategic communication, provided that a priori known conflicts are eliminated.

Our way to demonstrate this result builds on the following insights. Conflicts within the organization depend critically on the information available. As a result of this, feeding better information into the organization does not necessarily imply better decision-making; it may instead result in more relevant things being swept under the rug. Due to their ability to withhold information, those who can filter information along its way have a significant influence on the decisions that are made. Rational information acquisition by the organization takes all these factors into account and eliminates conflicts to the point where this is possible. As a consequence, based on the information that reaches the decision-maker, all parties would make the same decision; communication and delegation are outcome equivalent. In other words, everything is as if the informed party were formally legitimized to make choices.

Our conclusions stem from an analysis of a stylized model of a multidivisional organization with a common headquarters. For concreteness, we have in mind a car manufacturer with a European and a US-American division. The firm wishes to build a new model that will be sold on both markets. To learn about customers' tastes, the firm launches a market study. Due to economies of scale, market research is directed by the US division, production is directed by the European division. The tastes of Europeans and US-Americans are positively but not perfectly correlated. No systematic taste differences between the continents are expected; prior conflicts are absent. But, depending on the results of the study, different designs could turn out to be optimal for each market and conflicts could arise between the divisions ex post. Hence, communication is strategic and prone to manipulation.

We show that at most the inferences from the study, the optimal design from the US-division's perspective, but not the observed results themselves can be communicated in equilibrium. What inferences the European receiver infers from the US sender's inference depends critically on the nature of information that enters the organization. Headquarters shapes the communication process with a view to reaching the highest feasible joint surplus for the two divisions. Its only influence is through the attention devoted to the two markets in the study. The optimal way to do this is to equalize the residual uncertainty that remains for each division when information is used optimally from the receiver's point of view. Communication is not completely honest, as details are dropped, but unbiased: based on the optimal information, the sender's recommendation is an unbiased estimate of the receiver's preferred course of action. To achieve this unbiasedness, the market study devotes relatively more attention to the receiver's market. The sender needs to be forced to base his inference more heavily on receiver relevant facts and so sender relevant facts need to be observed with noise. Since optimal information eliminates biases, it does not matter where the decision is taken; it is always the same. Moreover, as the sender transmits the maximal amount of information he is willing to provide, the mechanism reaches the highest payoff for the organization among all direct communication protocols where the sender gets to see the information first. Hence, decision-making is arguably sound, as claimed.

Several lines of thought in our theory appear already, without a formal model, in March and Simon (1958). In their description of problem-solving, the authors note that: "The design of the search process is itself often an object of rational decision." (p.140). In their discussion of communication processes inside an organization, the authors coin the term *uncertainty absorption* and describe its consequences as follows:

> "Uncertainty absorption takes place when inferences are drawn from a body of evidence and the inferences, instead of the evidence itself, are then communicated. [...] Both the amount and the *locus of uncertainty absorption affect the influence structure of the organization.* Because of this, uncertainty absorption is frequently used, consciously or unconsciously, as a technique for acquiring and exercising power. [...] Whatever may be the position in the organization holding the formal authority to legitimize the decision, to a considerable extent the effective discretion is exercised at the points of uncertainty absorption."
> (March and Simon (1958), pp 165–167, emphasis in original)

In our model, uncertainty absorption corresponds to the sender drawing a unidimensional inference - a conditional expectation - from multiple signals. And indeed, although the receiver is formally legitimized to make the decision, the effective discretion is in fact exercised by the sender. This goes so far that communication and delegation become outcome equivalent. Given optimal information, allocating formal authority to the informed sender or bringing the information to the receiver are two ways to reach exactly the same outcome.

Our opening lines are inspired by the picture of organizations drawn by Cyert and March (1963), in particular their insightful discussion of communication and information acquisition (chapter 4). The ideas that information needs to be acquired, that the search for information is endogenous, and that the communication system influences the information that is acquired, all appear in their work. Our contribution is to offer a formal model that puts these elements together and hopefully advances our understanding of them. Our main result is that decisions can be steered indirectly by choosing what issues to look into and how deeply to probe into them. While it may be surprising how well this works in principle, it seems obvious that it does work in practice. Indeed, Cyert et al. (1958) offer case study evidence consistent with our theory. The authors followed a medium-large manufacturing concern in the 1950s in the process of installing an electronic data-processing system. It was quickly decided that an outside consulting firm was needed. An offer was obtained from a consulting firm named Alpha in the study. There was an important person in the manufacturing concern, named the controller. After Alpha had made its offer, the controller decided that a competing offer should be requested from another firm; he selected a firm Beta out of a list of candidates that had been prepared beforehand. Beta delivered its offer. A memorandum was prepared at the request of the controller that listed the criteria that should be looked at to compare the offers and reach a decision. The final staff memorandum on the decision clearly recommended to hire Beta, a recommendation that the controller accepted. The controller is cited with the words: "I asked the boys to set down the pros and cons. The decision was Beta. It was entirely their decision."(Cyert et al. (1958), p.332)

Of course, we will never know why the boys favored Beta over Alpha; it could be that they wanted to please the controller or that Beta made the better offer. However, there is no account of explicit manipulation in the study. The point is that the controller can steer the decision indirectly to the point that it doesn't really matter who takes the decision. This chapter shows that this is precisely how a benevolent controller should act.

We are not the first to take up Simon's concept of authority. Aghion and Tirole (1997) distinguish formal from real authority. The allocation of formal authority has important effects on initiative and participation when there are private costs of information acquisition. In contrast, we abstract from such costs and information is acquired by the organization itself. On top of this, our concept of real authority is different, allowing the receiver to amend proposals as in Crawford and Sobel (1982), the seminal paper on strategic information transmission between a sender and a receiver. Dessein (2002) studies the allocation of formal authority in the Crawford-Sobel model and shows that delegating decision rights to the informed sender is always better than communicating whenever meaningful communication is possible at all. The essential differences to the present chapter are the nature of biases and information. In the Crawford-Sobel model, the sender wishes to induce an action that exceeds the ideal action of the receiver in each state of the world by some constant. Moreover, the sender's information is exogenously given. In our model, the information is endogenously determined by the organization and influences the magnitude and direction of biases, which both depend on the realized state of the world. If the organization can adapt to the situation along the informational margin, then delegation and communication become perfect substitutes.

Alonso et al. (2008) study the allocation of formal authority in an organization where two divisions interact with a headquarters. Both divisions have some information and need to make choices, preferably in a coordinated way. The organization can choose between vertical communication where all information flows upwards to a headquarters or horizontal communication where one division communicates with the other and the latter is in charge of decision making for both divisions. Depending on the relative importance of coordinating actions and of adapting choices to local conditions either one or the other form of communication is optimal.[1] We study the same organization, however, in a quite different situation where the form of the organization is exogenously given and information instead is endogenous. Allowing the firm to choose the information that enters the organization makes different allocations of formal authority perfect substitutes in our model.

Communication works so well in our model, because the organization acquires information that eliminates conflicts to the point where this is possible. Although communication is not completely honest about the observed evidence, it is honest about the inference drawn from the evidence. Following the sender's advice one-for-one is optimal for the receiver, because remaining conflicts are orthogonal to the sender's recommendation. Battaglini (2002) studies a multi-sender multidimensional cheap talk problem and uses an orthogonal construction to elicit perfect information from the senders. Although we rely on orthogonality as Battaglini does, we cannot apply his construction because there is only one sender in our model and a unidimensional choice needs to be made. Instead, we need to adjust the information that the sender obtains to ensure orthogonality.[2] Preferences over information are not studied in Battaglini (2002).[3]

Controlling the access to information in a communication game is first studied in Ivanov (2010), showing that communicating with an expert who has partial information is better for the receiver than talking to a an expert who is perfectly informed whenever meaningful communication is possible. Moreover, communication with controlled information can even outperform optimal delegation to a per-

---

[1] An important difference between Dessein (2002) and Alonso et al. (2008) and the current chapter are that biases are state dependent in the latter. Such biases have also been analyzed, e.g., by Stein (1989), Ottaviani and Sørensen (2006a), Ottaviani and Sørensen (2006b), Kawamura (2015), and in the most general model by Gordon (2010).

[2] In a model with a privately known bias, Li and Madarász (2008) remark that communication works well if the bias is independent of the state and symmetrically distributed around zero. However, these authors study mandatory disclosure of given biases, whereas biases arise from information in our model. For further analyses of privately known unidirectional biases, see Morgan and Stocken (2003) and Dimitrakas and Sarafidis (2005).

[3] For other approaches to multidimensional cheap talk, see Meyer et al. (2013), Chakraborty and Harbaugh (2007), Chakraborty and Harbaugh (2010), and Levy and Razin (2007). These papers are not concerned with the impact of the quality of multidimensional information on communication.

fectly informed expert from the receiver's point of view. The common ground with the present chapter is the comparison of institutions, one of which involves controlling the quality of information, broadly speaking.[4] However, there are substantial differences, the most important one is that Ivanov (2010) analyzes senders who are systematically biased in one direction. Moreover, we study noisy information structures within a class that induces smooth posteriors whereas Ivanov (2010) investigates partitional information structures. Our main result is the outcome equivalence of optimal delegation and communication, which does not arise in Ivanov's model.[5]

The optimal information structure is noisy in our model, in order to make the sender willing to share his information and the receiver willing to use it; so, noise helps to facilitate communication as in Blume et al. (2007) or Goltsman et al. (2009). However, in Blume et al. (2007), the sender has perfect information and noise is added to the sender's message, while our sender is endowed with noisy information but communicates without further noise. Goltsman et al. (2009) compare the outcomes of different decision protocols and show, among other results, that the noise-mechanism of Blume et al. (2007) is an optimal mediation mechanism. Moscarini (2007) assumes Gaussian noisy information and noiseless communication to study central bank competence. Communication equilibria are partitional in his analysis and information is exogenously given in his approach; our comparative statics predictions are similar. Gordon and Nöldeke (2013) combine Gaussian noise in information and communication. Similar to this chapter, the communication equilibria are in linear strategies. However, Gordon and Nöldeke (2013) restrict attention to the class of equilibria in linear strategies a priori and use the resulting equilibrium strategies to explain figures of speech, such as exaggeration, understatement, and irony. The existence of these equilibria depends on the noise that is added exogenously to the sender's message. In contrast, communication is noiseless in our model and we are interested in an unrestricted optimum of our game.

Preferences over information in markets have been studied extensively. Vives (1999) surveys the literature on product market competition with information frictions, Vives (2008) the literature on financial markets. A more recent overview is

---

4 Argenziano et al. (2013) compare delegation and communication when the sender has a one-sided bias and acquires costly information.

5 A further difference is that Ivanov (2010) studies information structures that are optimal for the receiver whereas we study optimality from the perspective of joint surplus. For an analysis of sender optimal information structures, see, e.g., Szalay (2005) and Eső and Szalay (2015). Kamenica and Gentzkow (2011) analyze sender optimal persuasion rules; the difference to the present problem is the commitment to information that reaches the decision-maker.

given in Pavan and Vives (2015). Angeletos and Pavan (2007) investigate the social value of information in large markets with strategic complementarity or substitutability, externalities, and heterogeneous information. This literature relies on Gaussian noise and, depending on the context, either CARA preferences or quadratic payoffs, to find equilibria in linear strategies. Nöldeke and Tröger (2006) prove the existence of linear strategy equilibria in a market microstructure model for the wider class of elliptical distributions, which contains the Normal distribution as a special case. We allow at the same time for general payoff functions and elliptical distributions. We are not aware of any other contribution that does so too. Note also that in the literature on strategic market interactions, agents observe information and choose actions directly. Cheap talk communication of unverifiable information followed by a common action that affects the payoffs of a sender and a receiver is not analyzed in this literature.

The remainder of the chapter is organized as follows. In Section two, we present the model. In Section three, we analyze communication and derive an upper bound on the amount of information that can be transmitted in any equilibrium. In Section four, we analyze optimal information acquisition from the organization's perspective. A final section concludes and discusses extensions. Lengthy proofs are gathered in the Appendix.

2.2 Model

We consider a firm, comprised of two divisions with a common headquarters. A decision $x \in \mathbb{R}$ needs to be taken that affects the payoffs of all three parties. Division one has preferences described by

$$u^S\left(x,\eta\right) = -\ell\left(x-\eta\right);$$

division two has preferences

$$u^R\left(x,\omega\right) = -\ell\left(x-\omega\right).$$

The loss function $\ell\left(q\right)$ is symmetric around its minimizer, $q = 0$, twice differentiable, and at least as convex as the quadratic function. More precisely, we assume that the Arrow-Pratt measure of relative curvature of the loss function satisfies $\frac{q\ell''(q)}{\ell'(q)} \geq 1$ for all $q \neq 0$.[6] In addition, ℓ rises sufficiently slowly to make expected utility well-

[6] Examples include $\ell\left(q\right) = q^{2n}$ for $n \in \mathbb{N}$.

defined. η and ω are random variables - the tastes of consumers that are served by the two divisions - whose realizations describe the ideal policies from each division's point of view. These ideal policies are given by $x^R(\omega) = \omega$ and $x^S(\eta) = \eta$, respectively. The realizations of ω and η are unknown at the outset. Headquarters is interested in joint surplus[7]

$$u^H\left(x, \eta, \omega\right) = -\ell\left(x - \eta\right) - \ell\left(x - \omega\right).$$

The decision process in the firm is organized as follows. Division one, henceforth the sender, gets to observe noisy signals

$$s_\omega = \omega + \varepsilon_\omega \quad \text{and} \quad s_\eta = \eta + \varepsilon_\eta,$$

where ε_ω and ε_η are uncorrelated noise terms. Division two, henceforth the receiver, is in charge of making the decision. Headquarters shapes the communication between the divisions by controlling the research that division one conducts. Formally, headquarters chooses the amount of noise in the sender's signals, that is the variances $\sigma^2_{\varepsilon_\omega}$ and $\sigma^2_{\varepsilon_\eta}$ of the noise terms ε_ω and ε_η. This choice is publicly observable. However, the realizations of signals s_ω and s_η are privately observed by the sender. The sender communicates with the receiver, who finally chooses x. There is no cost of sending messages and the receiver is unable to commit to the action x as a function of the information he receives, so communication is modeled as cheap talk in the sense of Crawford and Sobel (1982).

To make the updating about the underlying states tractable we place restrictions on the joint distribution of $\omega, \eta, \varepsilon_\omega$ and ε_η. We focus on an environment where conditional means are linear functions of the observed information. Moreover, linear transformations of the underlying random variables follow the same class of distribution as the underlying random variables do. As is well known, these assumptions are satisfied, e.g., if $\omega, \eta, \varepsilon_\omega$ and ε_η are jointly normally distributed. However, these assumptions are generally fulfilled by all members of the class of elliptical distributions, which includes the Normal distribution as a special case. In what follows, we term the joint distribution of ω, η, s_ω and s_η the *information structure*. An information structure is feasible if it belongs to the elliptical class, has a density function, finite first and second moments, and if the marginal joint distribution of ω and η equals the prior distribution. Given these assumptions, the joint density of a random vector $\mathbf{Y}$ of dimension n can be written as $f_{\mathbf{Y}}\left(\mathbf{y}\right) =$

7 As shown by Alonso et al. (2008), profit sharing between headquarters and the divisions gives rise to such headquarters preferences.

$c_n |\Sigma|^{-\frac{1}{2}} \phi\left(\left(\mathbf{y}-\mu\right)' \Sigma^{-1} \left(\mathbf{y}-\mu\right)\right)$, where μ is the mean vector, Σ is up to a constant factor equal to the covariance matrix, $\phi(\cdot)$ is a given function, and c_n a scale factor, which we simply denote $c = c_1$ in the one-dimensional case.[8]

We assume that all the differences in preferences are unsystematic and random. Formally, we assume that $\mathbb{E}[\omega] = \mathbb{E}[\eta]$. This amounts to saying that systematic differences in preferences - where one division wishes to push the decision in a particular direction relative to the other division's preferred choice - have been eliminated prior to the current interaction. This does not imply that preferences are aligned. It only implies that based on prior information no differences of opinions are expected. In addition, we impose the innocuous normalization that $\mathbb{E}[\omega] = \mathbb{E}[\eta] = \mathbb{E}[\varepsilon_\omega] = \mathbb{E}[\varepsilon_\eta] = 0$. The covariance matrix is described by $\sigma_\omega^2 \equiv Var(\omega)$, $\sigma_\eta^2 \equiv Var(\eta)$, $\sigma_{\varepsilon_i}^2 \equiv Var(\varepsilon_i)$ for $i = \omega, \eta$, and $\sigma_{\omega\eta} \equiv Cov(\omega, \eta)$. The covariances involving the noise terms are zero by assumption. The coefficient of correlation between ω and η is defined as

$$\rho \equiv \frac{\sigma_{\omega\eta}}{\sigma_\omega \sigma_\eta}.$$

To complete the description of the model, consider the ideal policies from each division's perspective if each of them had access to the information s_ω and s_η.

Lemma 2.1. *As functions of the underlying signal realizations, s_ω, s_η, the ideal choice functions of the receiver and the sender are*

$$x^R(s_\omega, s_\eta) \equiv \arg\max_x \mathbb{E}\left[u^R(x,\omega) \middle| s_\omega, s_\eta\right] = \mathbb{E}\left[\omega | s_\omega, s_\eta\right] = \alpha^R s_\omega + \beta^R s_\eta$$

and

$$x^S(s_\omega, s_\eta) \equiv \arg\max_x \mathbb{E}\left[u^S(x,\eta) \middle| s_\omega, s_\eta\right] = \mathbb{E}\left[\eta \middle| s_\omega, s_\eta\right] = \alpha^S s_\omega + \beta^S s_\eta,$$

where α^i, β^i for $i = R, S$ are weights, independent of s_ω, s_η.
Unless $\sigma_\omega^2 = \sigma_\eta^2 = \sigma_{\omega\eta}$, $x^R(s_\omega, s_\eta) \neq x^S(s_\omega, s_\eta)$ for all $s_\omega, s_\eta \neq 0$.

The optimal choice functions correspond to the conditional expectations and conditional expectations are linear in our statistical framework. The intuition is familiar from the Normal distribution-quadratic loss case; we state the result as a lemma, because we prove the generalization both with respect to a wider class of distributions and loss functions.

[8] The Normal distribution corresponds to the case $\phi(u) = e^{-\frac{u}{2}}$ and Σ identically equal to the covariance matrix. The factor c_n depends on n to make f a density. Other members of the elliptical class include, e.g., the exponential power distribution (and as a special case the Laplace) or the logistic distribution. For more details on elliptical distributions see, e.g., Fang et al. (1990).

The divisions disagree on the optimal course of action for almost all signal realizations unless the tastes of their customers are perfectly correlated with identical marginal distributions, in which case their customers are essentially identical. The coefficient of correlation captures the alignment of interests in an intuitive way. It is easy to show that no meaningful communication is possible if $\rho \leq 0$.[9] To focus on the interesting case, we assume that $0 < \rho < 1$.

It is worth pausing for a minute to discuss the crucial assumptions and differences to other approaches in the literature. The main difference is the way we capture conflicts of interests. We assume identical loss functions for sender and receiver and capture all the differences between them by the random variables ω and η and their distributions. The first moments describe ideal policies, the second moments shape expected utilities. Assuming equal prior expectations amounts to saying that differences of opinion prior to the current interaction have been eliminated. The remaining conflicts are random and unsystematic, in the sense that their expected value is zero. We make these assumptions, because it is by now well known that communication does not work well with systematic differences of opinions. In contrast, it is not yet known how well communication can work with unsystematic differences of opinions.

We analyze the game proceeding backwards, starting with the inference that the sender draws from observing facts and the ensuing communication continuation games. We then reduce the model to one where communication is about inferences instead of facts and discuss the receiver's inferences drawn from the sender's inference. Building on this analysis, we discuss the optimal organizational response to filtering information this way, the optimal amount and kind of information that the organization acquires.

2.3 The Sender as a Strategic Information Channel

Suppose that headquarters has chosen a research policy - formally, an information structure - and the sender gets to observe the results of the research. What part of the observed information is the sender willing to share with the receiver at all?

2.3.1 Limits to Communication

We focus on Bayesian equilibria in the communication game. After observing signal realizations s_ω, s_η, the sender sends a message $m \in \mathbb{M}$ to the receiver. The message

[9] A formal proof of this statement is available from the authors upon request.

space is sufficiently rich; we do not impose any restrictions on $\mathbb{M}$. It is enough to consider pure message strategies for the sender.[10] A pure sender strategy maps the sender's information into messages $M : \mathbb{R}^2 \to \mathbb{M}$, $(s_\omega, s_\eta) \mapsto m$. A pure receiver strategy maps messages into actions, $X : \mathbb{M} \to \mathbb{R}$, $m \mapsto x$. The receiver updates his belief about the sender's type after observing the sender's message and acts optimally against this belief. The following lemma derives an upper bound on the information that can be communicated in any equilibrium of the communication game. In particular, the sender is willing to share his inference but not the underlying facts.

Lemma 2.2. *In any equilibrium, all sender types s_ω, s_η such that $\alpha^S s_\omega + \beta^S s_\eta =$ constant induce the same action.*

Define the statistic

$$\theta \equiv \alpha^S s_\omega + \beta^S s_\eta.$$

All sender types with signal realizations s_ω, s_η adding up to θ share the same ideal policy, θ. Moreover, with symmetric loss functions, the sender's preferences over distinct actions depend only on the distance of these actions to θ. Hence, the set of types who share the same θ induce at most two distinct actions, and these actions need to be equidistant from θ in any equilibrium. However, any attempt to separate sender types whose signals aggregate to θ into subsets that induce distinct actions gives some other types, whose signals aggregate to some value close to θ, a strict incentive to lie. Hence, no such equilibrium can exist. Obviously, the lemma also implies that it is impossible to elicit the information s_ω, s_η from the sender, unless the ideal policies of sender and receiver coincide altogether.

Corollary 2.1. *Truthful communication of the underlying information, s_ω, s_η, is an equilibrium if and only if $\sigma_\omega^2 = \sigma_\eta^2 = \sigma_{\omega\eta}$.*

Since induced actions depend only on the realization of θ, the sender is willing to reveal at most the inference he draws from the facts, that is θ, but never the underlying facts. Hence, we can characterize any equilibrium of the communication game in terms of communication about the sender's inference, θ, only.[11]

---

[10] More specifically, it is standard in the literature to look at the most informative equilibria and these equilibria involve pure strategies in our game. Therefore, we abstain from introducing the notational clutter to deal formally with mixed strategies.

[11] Note the close connection between this result and the process of uncertainty absorption described in March and Simon (1958).

2.3.2 Inference from Inference and Conflicts

>From the ex ante perspective, before the signals are realized, the sender's inference is random itself. Any given choice of information structure gives rise to a joint distribution of ω, η, and θ. Given that $\omega, \eta, \varepsilon_\omega$ and ε_η follow a joint elliptical (Normal) distribution, the random variables ω, η, and θ follow a joint elliptical (Normal) distribution as well.[12] One can show that the moments involving θ are given by $\mathbb{E}[\theta] = 0$ as well as

$$Var(\theta) = \sigma_\eta^2 \frac{\frac{\sigma_{\varepsilon\omega}^2}{\sigma_\omega^2} + \frac{\sigma_{\varepsilon\eta}^2}{\sigma_\eta^2}\rho^2 + 1 - \rho^2}{\left(1 + \frac{\sigma_{\varepsilon\omega}^2}{\sigma_\omega^2}\right)\left(1 + \frac{\sigma_{\varepsilon\eta}^2}{\sigma_\eta^2}\right) - \rho^2}, \tag{2.3.1}$$

$$Cov(\omega,\theta) = \sigma_{\omega\eta} \frac{\frac{\sigma_{\varepsilon\eta}^2}{\sigma_\eta^2} + \frac{\sigma_{\varepsilon\omega}^2}{\sigma_\omega^2} + 1 - \rho^2}{\left(1 + \frac{\sigma_{\varepsilon\omega}^2}{\sigma_\omega^2}\right)\left(1 + \frac{\sigma_{\varepsilon\eta}^2}{\sigma_\eta^2}\right) - \rho^2}, \tag{2.3.2}$$

and

$$Cov(\eta,\theta) = Var(\theta). \tag{2.3.3}$$

Equations (2.3.1) and (2.3.2) depend crucially on the normalized noise variances, $\frac{\sigma_{\varepsilon\omega}^2}{\sigma_\omega^2}$ and $\frac{\sigma_{\varepsilon\eta}^2}{\sigma_\eta^2}$. By construction of θ, only covariance matrices with $Cov(\eta,\theta) = Var(\theta)$ are possible. For all that matters in terms of induced choices and payoffs, we can analyze our model in terms of this reduced form joint distribution of inference and underlying states.

What inference would the receiver draw if the sender communicated his inference? Since the joint distribution of ω, η, and θ is firmly within the class that has linear conditional means, the receiver's ideal policy conditional on observing θ is

$$\mathbb{E}[\omega|\theta] = \frac{Cov(\omega,\theta)}{Var(\theta)} \cdot \theta. \tag{2.3.4}$$

The conditional expectation corresponds to the linear regression of the unknown state on the observed information. To understand the slope of the regression, note that the regression of η on θ is simply

$$\mathbb{E}[\eta|\theta] = \frac{Cov(\eta,\theta)}{Var(\theta)} \cdot \theta = \theta. \tag{2.3.5}$$

Clearly, given that θ is the conditional expectation of η given the underlying facts, the sender does not revise his conditional expectation if shown θ again. In contrast,

[12] The proof of this statement follows from Fang et al. (1990) Theorem 2.16.

the receiver's inference corrects for the relative informational content of the sender's inference, θ, with respect to the underlying states ω and η: by equation (2.3.3), the slope $\frac{Cov(\omega,\theta)}{Var(\theta)}$ corresponds to $\frac{Cov(\omega,\theta)}{Cov(\eta,\theta)}$. If the sender gets to observe information that is relatively more informative about ω than about η, then $Cov(\omega,\theta) > Cov(\eta,\theta)$ and the receiver's ideal policy attaches a higher weight to the information θ than the sender's ideal policy. The situation is reversed if the sender gets to see information that is relatively more useful to the sender. The regressions have identical slopes if the sender's inference is equally informative about ω and η.

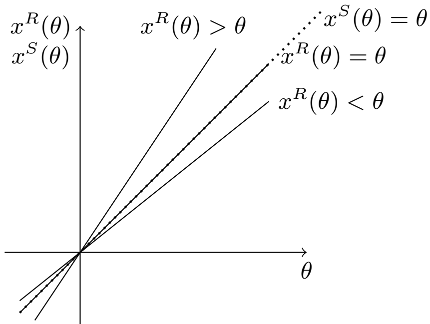

Figure 2.3.1. Conflicts with respect to θ between sender and receiver.

The difference $\theta - \mathbb{E}[\omega|\theta]$ describes the bias of the sender relative to the receiver. If the sender observes information that is relatively more informative about η, then the sender has incentives to exaggerate. If the sender's information is relatively more informative about ω, then the sender has incentives to downplay. Finally, there is no bias when communicating about the sender's inference when the sender's inference θ is equally informative about ω and η. For convenience, the three cases are depicted in Figure 2.3.1.

2.4 Optimal Information Structures

We now address headquarters' problem of choosing an optimal information structure. What information should the sender get to observe about each of the underlying taste parameters, ω and η? The sender's information impacts on payoffs through two channels. Firstly, assuming honest transmission of the sender's inference, the relative informational content of θ impacts directly on the sender's and the receiver's expected payoff from making a receiver-optimal decision based on θ. Secondly, the relative informational content determines the sender's bias in the communication game and thus impacts on the amount of information that is transmitted through communication. It is helpful to look at the two margins separately. Therefore, we begin our analysis with the clearly unrealistic case where the sender's inference θ

becomes publicly available.[13] In a second step, in Section 2.4.2, we look at the case of main interest, where θ is private information.

To streamline the exposition, we present our analysis first assuming that marginals are identical. That is, we assume $\sigma^2_\omega = \sigma^2_\eta$. We discuss the role of this assumption and abandon it in Section 2.4.4 below.

2.4.1 Public Inferences

2.4.1.1 Headquarters Problem

If the receiver observes the sender's inference θ, then he follows the policy $x^R(\theta) = \mathbb{E}[\omega|\theta] = \frac{Cov(\omega,\theta)}{Var(\theta)} \cdot \theta$, resulting in a loss of $\ell\left(\frac{Cov(\omega,\theta)}{Var(\theta)}\theta - \omega\right)$ for the receiver and a loss of $\ell\left(\frac{Cov(\omega,\theta)}{Var(\theta)}\theta - \eta\right)$ for the sender. Both losses depend only on sums of the underlying random variables, $\zeta \equiv \frac{Cov(\omega,\theta)}{Var(\theta)}\theta - \omega$ and $\tau \equiv \frac{Cov(\omega,\theta)}{Var(\theta)}\theta - \eta$, which are again elliptical (Normal). Let σ^2_ζ and σ^2_τ denote the variances of ζ and τ and let $z \equiv \frac{\zeta}{\sigma_\zeta}$ and $t \equiv \frac{\tau}{\sigma_\tau}$ denote the standardized arguments of the loss functions. As demonstrated formally in the Appendix, we can write headquarters' problem as

$$\max_{Cov(\omega,\theta),Var(\theta)} -\int \ell\left(\sigma_\zeta z\right) c\phi(z)\,dz - \int \ell\left(\sigma_\tau t\right) c\phi(t)\,dt$$
$$s.t.\ Cov(\omega,\theta), Var(\theta) \text{ feasible.}$$

where z and t follow a spherical (Standard Normal) distribution with density $c\phi(\cdot)$.

Each division's expected utility depends negatively on a residual variance that measures the residual uncertainty after using θ optimally from the receiver's perspective. Naturally, the residual uncertainty for the receiver is

$$\sigma^2_\zeta = \sigma^2_\omega - \frac{Cov(\omega,\theta)^2}{Var(\theta)} = Var(\omega|\theta), \tag{2.4.1}$$

where the second equality holds because θ is used optimally from the receiver's perspective.[14] In contrast, θ is in general not used optimally from the sender's perspective. The residual uncertainty that the sender faces when θ is used according to the policy $x^R(\theta)$ is

$$\sigma^2_\tau = \sigma^2_\eta - \left(2Cov(\omega,\theta) - \frac{Cov(\omega,\theta)^2}{Var(\theta)}\right), \tag{2.4.2}$$

---

[13] We can think of this as some form of mediated information transmission; the sender's information s_ω, s_η is aggregated to $\alpha^S s_\omega + \beta^S s_\eta = \theta$ and then mechanically transmitted to the receiver.

[14] For the derivation of the conditional second moments see Lemma 2.4 in the Appendix.

which differs from $Var\left(\eta|\theta\right) = \sigma_\eta^2 - Var\left(\theta\right)$ unless (2.3.5) and (2.3.4) are identically equal to each other.

Consider now the feasible set of information structures. Not any joint distribution of ω,η,θ is a feasible reduced form information structure, because θ must be derived from Bayesian updating by the sender about η, conditioning on the information that the sender gets to see. Thus, a joint distribution of ω,η and θ is feasible only if there are noise variances $\sigma^2_{\varepsilon_\omega}$ and $\sigma^2_{\varepsilon_\eta}$ that, together with the prior distribution, induce the joint distribution. The following lemma makes the restrictions from Bayesian updating explicit.

Lemma 2.3. *A joint distribution of ω,η,θ can be generated through Bayesian updating if and only if $Cov\left(\omega,\theta\right) \in \left[0,\sigma_{\omega\eta}\right]$ and for any given $Cov\left(\omega,\theta\right) = C$, $Var\left(\theta\right) \in \left[\frac{\sigma_\eta}{\sigma_\omega}\rho C, \frac{\sigma_\eta}{\sigma_\omega}\frac{1}{\rho}C\right]$.*

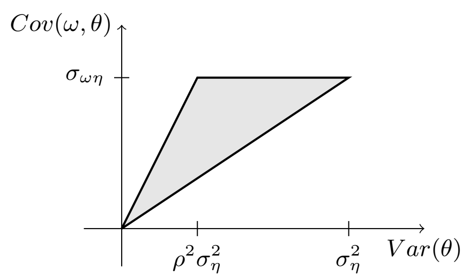

Figure 2.4.1. The feasible set of information structures, Γ.

$Var\left(\theta\right)$ and $Cov\left(\omega,\theta\right)$ are jointly constrained to lie in the triangle described in Figure 2.4.1. We call the feasible set Γ. To understand the shape of Γ, note that any pair of normalized noise variances, $\frac{\sigma^2_{\varepsilon\eta}}{\sigma^2_\eta}, \frac{\sigma^2_{\varepsilon\omega}}{\sigma^2_\omega} \geq 0$, results in a $Cov\left(\omega,\theta\right) \leq \sigma_{\omega\eta}$. The covariance is maximal if at least one of the signals is perfectly precise. In the limiting case of infinitely noisy signals, the sender does not revise his prior at all and so both $Var\left(\theta\right)$ and $Cov\left(\omega,\theta\right)$ are zero. If the sender observes a signal s_η without noise, $\sigma^2_{\varepsilon_\eta} = 0$, then his posterior mean becomes identically equal to η and the resulting variance is $Var\left(\theta\right) = \sigma^2_\eta$. If the sender observes s_ω without noise, $\sigma^2_{\varepsilon_\omega} = 0$, and the signal s_η is infinitely noisy, $\sigma^2_{\varepsilon_\eta} \to \infty$, then $Var\left(\theta\right) = \rho^2\sigma^2_\eta$, because the sender's posterior mean rises less than one for one with the sender's observation. By continuity, any pair of covariance and variance in the interior of the triangle can be generated by some pair of noise variances. Finally, Γ is always nonempty, because the lowest feasible $Var\left(\theta\right)$ for any given $Cov\left(\omega,\theta\right)$ is below the highest feasible $Var\left(\theta\right)$ by the Cauchy-Schwarz inequality, $\sigma^2_{\omega\eta} \leq \sigma^2_\eta\sigma^2_\omega$.[15]

[15] We include edges and vertices in the feasible set that result from taking limits. The limiting pos-

2.4.1.2 Equalizing Residual Uncertainty

We can now restate headquarters' problem as

$$\max_{Cov(\omega,\theta),Var(\theta)} -\int \ell\left(\sigma_\zeta z\right) c\phi(z)\,dz - \int \ell\left(\sigma_\tau t\right) c\phi(t)\,dt$$
$$s.t.\ Cov(\omega,\theta), Var(\theta) \in \Gamma,$$

where σ_ζ and σ_τ are defined in (2.4.1) and (2.4.2). Headquarters maximizes a continuous objective function on a compact domain, so the problem is well defined and a solution exists. The solution takes the following form:

Theorem 2.1. *Suppose that the sender and the receiver are equally uncertain ex ante, $\sigma_\omega^2 = \sigma_\eta^2$. If the loss function satisfies $\frac{q\ell''(q)}{\ell'(q)} > 1$ for all $q \neq 0$, then headquarters' problem of choosing an optimal information structure has a unique solution, which is given by $Var(\theta)^* = Cov(\omega,\theta)^* = \sigma_{\omega\eta}$. If the loss function satisfies $\frac{q\ell''(q)}{\ell'(q)} = 1$ for all $q \neq 0$ (corresponding to the quadratic case), then any information structure satisfying $Cov(\omega,\theta) = \sigma_{\omega\eta}$ is optimal.*

We solve the problem by maximizing sequentially with respect to $Var(\theta)$ and $Cov(\omega,\theta)$. For a given level of $Cov(\omega,\theta)$, headquarters' problem resembles a risk sharing problem. Both divisions dislike higher residual uncertainty and an increase of $Var(\theta)$ increases (2.4.1), the residual uncertainty the receiver faces, and decreases (2.4.2), the residual uncertainty the sender faces. For a sufficiently convex loss function, the problem is single-peaked in $Var(\theta)$ and has a unique maximum at the point where the residual uncertainty for both divisions is equalized. Equating (2.4.1) and (2.4.2) and solving for $Var(\theta)$, we obtain

$$Var(\theta)^* = Cov(\omega,\theta).$$

The residual uncertainty for both divisions is then equal to the residual uncertainty that the receiver faces, $Var(\omega|\theta) = \sigma_\omega^2 - Cov(\omega,\theta)$. Since this is a decreasing function of $Cov(\omega,\theta)$, it is optimal to choose $Cov(\omega,\theta)$ as high as possible,

$$Cov(\omega,\theta)^* = \sigma_{\omega\eta}.$$

The unique optimum corresponds to the intersection of the dashed and the solid

terior distributions and moments when one noise variance goes out of bounds converge to the distribution when only one signal is received; the limiting case when both noise variances go out of bounds converges to the distribution when no signal at all is received, the prior.

line in Figure 2.4.2. The role of the curvature condition is to guarantee uniqueness of the optimal $Var(\theta)$. For the quadratic loss function, headquarters' payoff becomes linear in the residual variances, which implies that the receiver's loss from increasing $Var(\theta)$ just offsets the sender's gain and thus the sum of their payoffs becomes independent of $Var(\theta)$. Hence, any information structure with the highest feasible $Cov(\omega,\theta)$, depicted as the solid line in the figure, is optimal.

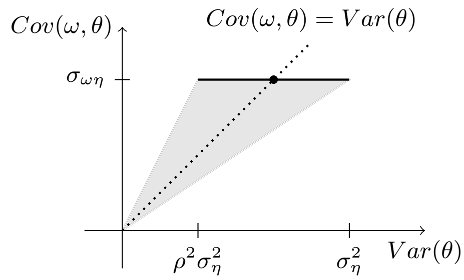

Figure 2.4.2. The optimal information structure maximizes $Cov(\omega,\theta)$. For sufficiently convex loss functions it is unique and satisfies $Cov(\omega,\theta) = Var(\theta)$.

The optimum can be understood by decomposing information into its common and idiosyncratic content. Since $Cov(\omega,\eta|\theta) = \sigma_{\omega\eta} - Cov(\omega,\theta)$, $Cov(\omega,\theta)$ measures the amount of common information. Naturally, the optimal information structure contains all the common information there is,

$$Cov(\omega,\eta|\theta)^* = \sigma_{\omega\eta} - Cov(\omega,\theta)^* = 0,$$

implying that conditional on θ, the taste parameters become uncorrelated. $Var(\theta)$ measures the amount of idiosyncratic information. Since there is only one signal, θ, idiosyncratic information necessarily involves a trade-off: $Var(\omega|\theta)$ is increasing in $Var(\theta)$, while $Var(\eta|\theta)$ is decreasing in $Var(\theta)$.

In terms of the underlying signals, headquarters allows the sender to observe ω without noise, $\sigma^2_{\varepsilon_\omega} = 0$, but adds noise $\sigma^2_{\varepsilon_\eta} = \frac{1-\rho^2}{\rho}\sigma^2_\eta$ to the signal about η. If the signal s_η were perfectly precise, then the sender would not pay any attention to the signal s_ω. While θ would still contain the maximum amount of common information, θ would not be informative enough about ω and so the receiver would face too much residual uncertainty. Hence, noise is needed to keep the sender from using the signal that is of primary importance to him exclusively.

2.4.2 Private Inferences

We now consider the case of main interest where the sender has private information about θ and thus is free to make up any statement he likes. As is standard in the literature, we assume that the sender and the receiver are able to coordinate on the ex ante Pareto optimal equilibrium in the communication game. The optimal information structure eliminates conflicts in a certain, well defined sense:

Theorem 2.2. *Let the sender and the receiver face equal prior uncertainty, $\sigma_\omega^2 = \sigma_\eta^2$. Then, the unique optimal information structure chosen by headquarters satisfies $Var(\theta)^* = Cov(\omega,\theta)^* = \sigma_{\omega\eta}$. The Pareto best equilibrium of the ensuing continuation game involves smooth strategies; the sender truthfully announces θ, $m^*(\theta) = \theta\ \forall\theta$, and the receiver takes the sender's advice at face value, $x^*(m) = m\ \forall m$. All parties' payoffs are the same as if the sender were given the right to choose the action x directly.*

The theorem is a straightforward implication of our preceding results in conjunction with a verification that the described strategies constitute an equilibrium of the communication game. Since headquarters cannot improve upon its payoff compared to the case where θ is public information, the situation corresponds to an optimum if this payoff is reached. Suppose the receiver believes that the sender plays the message strategy $m(\theta) = \theta$ for all θ. Then, his best reply is the action strategy $x^*(m) = \frac{Cov(\omega,\theta)^*}{Var(\theta)^*} \cdot m = m$ for all m. The sender, who anticipates this policy, induces his ideal policy by being truthful about θ, so the construction is indeed an equilibrium. Note that in this equilibrium the strategies of both players are smooth - in fact, linear - functions.

Since $x^*(m^*(\theta)) = \theta$ for all θ, the sender's optimal policy is implemented for all θ. Consequently, whether the sender communicates with the receiver or whether the sender is given the right to choose the policy, the payoffs of all parties involved are exactly the same.[16] The intuition is that, for equal marginals, an information structure that equalizes residual uncertainty automatically eliminates any bias in the use of information. Formally, $Cov(\omega,\theta)^* = Var(\theta)^*$ implies

$$x^R(\theta) - x^S(\theta) = \left(\frac{Cov(\omega,\theta)^*}{Var(\theta)^*} - 1\right)\cdot\theta = 0 \quad \forall\theta.$$

Note that there remains a conflict between sender and receiver with respect to using the underlying signals, s_ω and s_η. However, the receiver simply cannot do better than follow the sender's advice, because based on observing the sender's inference θ, a garbled piece of information, the receiver's ideal choice coincides with the sender's ideal choice based on observing the underlying signals. The sender is

[16] Note that the problem of multiple solutions for the quadratic loss case if θ is public information is eliminated, because truthful communication now requires that $\frac{Cov(\omega,\theta)^*}{Var(\theta)^*} = 1$.

willing to share his inference despite disagreement too. The sender knows that the receiver would ideally like to choose an action that matches the state ω, not θ. However, under the optimal information structure, the sender's recommendation θ and the difference $\omega - \theta$ become uncorrelated. Put differently, the optimal information structure *orthogonalizes* the conflict between the divisions and the recommendation and hence removes any impediments to communication.

Communication is in fact unsurpassed by any form of delegation, even *optimal delegation*. Even if headquarters or the receiver had the right to constrain the sender's discretion under delegation, they would not want to make use of this right. The sender's optimal choice is necessarily a function of his inference θ only, and the sender uses this inference in the receiver's best interest. Hence, constraining the sender's discretion under delegation decreases the receiver's payoff and joint surplus.

2.4.3 The Quality of Decision Making

Under the optimal information structure information is lost because only inferences are transmitted. How much is lost by such garbling and how does this depend on the underlying conflicts?

We can measure the amount of information transmitted in equilibrium by the variance of induced choices; the higher this variance, the more information is transmitted. Headquarters throws in just enough noise to ensure that $Var(\theta) = \sigma_{\omega\eta}$. For identical priors, the variance of the induced choice is thus

$$Var(\theta) = \rho\sigma_\eta^2.$$

The higher is ρ, the more variable the induced choice. In the limit as $\rho \to 1$, the sender truthfully announces η and the variance of choices approaches σ_η^2. There are two reasons why increasing ρ results in an improvement of information transmission. Recall that the sender always observes ω without noise. The higher is ρ, the higher the attention the sender pays to this signal and the more this signal is reflected in the sender's preferred choice. Moreover, the sender observes η with an amount of noise equal to $\sigma_{\varepsilon_\eta}^2 = \frac{1-\rho^2}{\rho}\sigma_\eta^2$, a decreasing function of ρ. The higher is ρ, the more precise the sender's signal about the sender-relevant random variable η. So, senders with better aligned interests are more trustworthy to begin with and get endowed with more precise information, rendering their advice even more valuable.[17]

2.4.4 Extensions: Unequal Priors

We now drop the assumption of equal prior uncertainty and allow for $\sigma_\omega^2 \neq \sigma_\eta^2$. For quadratic loss functions, the canonical case studied in the literature, our result generalizes to asymmetric priors.

Proposition 2.1. *Assume quadratic loss functions and suppose that* $\min\left\{\sigma_\omega^2, \sigma_\eta^2\right\} \geq \sigma_{\omega\eta}$. *Then, headquarters' optimal choice of information structure is unique and given by* $Var(\theta)^* = Cov(\omega,\theta)^* = \sigma_{\omega\eta}$. *All parties receive the same expected payoff, regardless of who has the right to choose* x.

Recall that by Theorem 2.1 any information structure satisfying $Cov(\omega,\theta) = \sigma_{\omega\eta}$ is optimal for quadratic losses if θ is public. Hence, to show that headquarters can reach the same expected payoff under communication of unverifiable information - and under delegation - it suffices to show that the admissible set of information structures contains the element $Var(\theta) = Cov(\omega,\theta) = \sigma_{\omega\eta}$. The condition in the proposition is equivalent to

$$\frac{\sigma_\eta}{\sigma_\omega}\rho \leq 1 \leq \frac{\sigma_\eta}{\sigma_\omega}\frac{1}{\rho},$$

which guarantees that the 45° line is an element of the feasible set, Γ. We need to rule out very asymmetric priors where $\sigma_\omega^2 > \sigma_{\omega\eta} > \sigma_\eta^2$ or $\sigma_\eta^2 > \sigma_{\omega\eta} > \sigma_\omega^2$ that would render the solution $Var(\theta)^* = Cov(\omega,\theta)^* = \sigma_{\omega\eta}$ infeasible. Note that by Cauchy-Schwarz, it is impossible that both prior variances exceed the covariance.

Note that nonverifiability makes the solution unique. While the sum of residual variances is constant for all information structures with the highest feasible $Cov(\omega,\theta)$, there is only one information structure among them that makes the signal θ equally useful for both sender and receiver and thus ensures that truthful communication about θ is an equilibrium.

2.5 Conclusions

Two divisions, overarched by a headquarters, need to reach a decision that affects the payoffs of all parties involved. Division one privately gets to observe informa-

[17] In the limit, the feasible set of information structures converges to the 45° line and any piece of information is equally informative about ω and η. Hence endowing the sender with perfect information becomes optimal. At the same time the underlying interests of sender and receiver become perfectly correlated.

tion about ideal policies from both divisions' perspectives. Division one draws inferences from the information and communicates them to division two. Division two, who retains the right to make the decision, draws its own inferences from division one's inferences. Anticipating the chain of inferences within the organization, headquarters chooses what information to acquire at the outset. Choosing what to look into is a powerful tool. When properly done, conflicts within the organization are diminished, making it less important who has the right to make decisions: communication and delegation become outcome equivalent.

Almost by definition, an equivalence result raises nearly as many questions as it answers. In particular, one may wonder what happens if not headquarters but the sender has discretion over the acquisition of information. Quite clearly, it seems, that the benchmark of no loss through communication cannot be reached. Much to our surprise, we show in companion work that this conclusion is unwarranted. For a special case of the current environment, we are able to show that the sender acquiring orthogonalized information remains an equilibrium of the game. There are also other equilibria, but the sender cannot gain from making the information more useful to himself - precisely, because its usefulness would be lost in communication. Many other interesting questions can be pursued in our environment. We leave these for future work.

2.6 Appendix

Lemma 2.4. *Let Y follow an elliptical distribution, $Y \sim EC_n(\mu, \Sigma, \phi)$. Further let*

$$Y = (Y_1, Y_2), \quad \mu = (\mu_1, \mu_2), \quad \Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix},$$

where the dimensions of Y_1, μ_1 and Σ_{11} are m, m, and $m \times m$.

i) The elliptical distribution is symmetric about μ.

ii) Linear combinations of elliptically distributed random variables are again elliptical.

iii) The conditional distribution of $(Y_1 | Y_2 = y_2)$ is elliptical, with conditional mean vector

$$\mathbb{E}\left[Y_1 | Y_2 = y_2\right] = \mu_1 + \left(y_2 - \mu_2\right) \Sigma_{22}^{-1} \Sigma_{21} \tag{A2.1}$$

and conditional covariance matrix satisfying

$$\Sigma^* = \Sigma_{11} - \Sigma_{12} \Sigma_{22}^{-1} \Sigma_{21}. \tag{A2.2}$$

Proof of Lemma 2.4. i) by definition, ii) Fang et al. (1990) Theorem 2.16, iii) Fang et al. (1990) Theorem 2.18. □

Proof of Lemma 2.1. Let $u \equiv u^R = u^S$ and $z = \omega, \eta$. Consider the problem

$$\max_x \int_{-\infty}^{\infty} u\,(x - z)\, f\left(z|\, s_\omega, s_\eta\right) dz,$$

where $f\left(z|\, s_\omega, s_\eta\right)$ is the conditional density of $z = \omega, \eta$ given the signals. Since the utility depends only on the distance between x and z we have $u'\,(x - z) > 0$ for $z < x$, $u'\,(x - z) = 0$ for $x = z$, and $u'\,(x - z) < 0$ for $z > x$.

Consider the candidate solution $x^* = \mu_z \equiv \mathbb{E}\left[z|\, s_\omega, s_\eta\right]$. The first-order condition can be written as

$$\int_{-\infty}^{\infty} u'\left(x^* - z\right) f\left(z|\, s_\omega, s_\eta\right) dz = \int_{-\infty}^{\infty} u'\left(\mu_z - z\right) f\left(z|\, s_\omega, s_\eta\right) dz = 0.$$

Consider two points $z_1 = \mu_z - \Delta$ and $z_2 = \mu_z + \Delta$ for arbitrary $\Delta > 0$. By symmetry of u around its bliss point and symmetry of the distribution around μ_z, we have

$$u'\,(\Delta)\, f\left(\mu_z - \Delta\middle|\, s_\omega, s_\eta\right) = -u'\,(-\Delta)\, f\left(\mu_z + \Delta\middle|\, s_\omega, s_\eta\right).$$

Since this holds point-wise for each Δ, it also holds if we integrate over Δ. Thus, the first-order condition is satisfied at $x^* = \mu_z$. By concavity of u in x, only one value of x satisfies the first-order condition.

Applying equation (A2.1), the conditional expectations are

$$\mathbb{E}\left[\eta|\, s_\omega, s_\eta\right] = \alpha^S s_\omega + \beta^S s_\eta \tag{A2.3}$$

and

$$\mathbb{E}\left[\omega|\, s_\omega, s_\eta\right] = \alpha^R s_\omega + \beta^R s_\eta, \tag{A2.4}$$

where the weights in the sender's ideal choice are

$$\alpha^S = \sigma^2_{\varepsilon_\eta} \frac{\rho\sigma_\omega\sigma_\eta}{(\sigma^2_\omega + \sigma^2_{\varepsilon_\omega})(\sigma^2_\eta + \sigma^2_{\varepsilon_\eta}) - (\rho\sigma_\omega\sigma_\eta)^2}$$

and

$$\beta^S = \sigma^2_\eta \frac{\sigma^2_{\varepsilon_\omega} - \sigma^2_\omega\rho^2 + \sigma^2_\omega}{(\sigma^2_\omega + \sigma^2_{\varepsilon_\omega})(\sigma^2_\eta + \sigma^2_{\varepsilon_\eta}) - (\rho\sigma_\omega\sigma_\eta)^2}$$

and the weights in the receiver's ideal choice are

$$\alpha^R = \sigma_\omega^2 \frac{\sigma_{\varepsilon_\eta}^2 + \sigma_\eta^2 - \sigma_\eta^2\rho^2}{\left(\sigma_\omega^2 + \sigma_{\varepsilon_\omega}^2\right)\left(\sigma_\eta^2 + \sigma_{\varepsilon_\eta}^2\right) - \left(\rho\sigma_\omega\sigma_\eta\right)^2}$$

and

$$\beta^R = \sigma_{\varepsilon_\omega}^2 \frac{\sigma_\eta\sigma_\omega\rho}{\left(\sigma_\omega^2 + \sigma_{\varepsilon_\omega}^2\right)\left(\sigma_\eta^2 + \sigma_{\varepsilon_\eta}^2\right) - \left(\rho\sigma_\omega\sigma_\eta\right)^2}.$$

First, suppose $\sigma_{\varepsilon_\eta}^2$ and $\sigma_{\varepsilon_\omega}^2$ are both positive and finite. Equations $(A2.3)$ and $(A2.4)$ are identical for all s_ω and s_η if and only if

$$\sigma_{\varepsilon_\eta}^2 \rho\sigma_\omega\sigma_\eta = \sigma_\omega^2\left(\sigma_{\varepsilon_\eta}^2 + \sigma_\eta^2 - \sigma_\eta^2\rho^2\right)$$

and

$$\sigma_\eta^2\left(\sigma_{\varepsilon_\omega}^2 - \sigma_\omega^2\rho^2 + \sigma_\omega^2\right) = \sigma_\eta\sigma_\omega\rho\sigma_{\varepsilon_\omega}^2.$$

This requires that

$$\sigma_\eta^2\left(1-\rho^2\right) = \left(\frac{\rho\sigma_\eta}{\sigma_\omega} - 1\right)\sigma_{\varepsilon_\eta}^2$$

and

$$\sigma_\omega^2\left(1-\rho^2\right) = \left(\frac{\sigma_\omega\rho}{\sigma_\eta} - 1\right)\sigma_{\varepsilon_\omega}^2.$$

A necessary and sufficient condition for these two conditions to hold simultaneously is $\sigma_{\omega\eta} = \sigma_\eta^2 = \sigma_\omega^2$.

Consider now the limiting cases where one of the variances goes out of bounds. Applying l'Hôpital's rule to $(A2.3)$ and $(A2.4)$, we get in the limit as $\sigma_{\varepsilon_\eta}^2 \to \infty$

$$\mathbb{E}\left[\omega|s_\omega\right] = \frac{\sigma_\omega^2}{\sigma_\omega^2 + \sigma_{\varepsilon_\omega}^2} s_\omega \qquad \text{and} \qquad \mathbb{E}\left[\eta|s_\omega\right] = \frac{\rho_{\omega\eta}\sigma_\omega\sigma_\eta}{\sigma_\omega^2 + \sigma_{\varepsilon_\omega}^2} s_\omega,$$

so that

$$\mathbb{E}\left[\omega|s_\omega\right] \equiv \mathbb{E}\left[\eta|s_\omega\right] \qquad \Leftrightarrow \qquad \rho\sigma_\eta = \sigma_\omega.$$

Likewise, for the case where $\sigma_{\varepsilon_\omega}^2 \to \infty$, we get

$$\mathbb{E}\left[\omega|s_\eta\right] = \frac{\rho\sigma_\omega\sigma_\eta}{\sigma_\eta^2 + \sigma_{\varepsilon_\eta}^2} s_\eta \qquad \text{and} \qquad \mathbb{E}\left[\eta|s_\eta\right] = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{\varepsilon_\eta}^2} s_\eta,$$

so

$$\mathbb{E}\left[\omega|s_\eta\right] \equiv \mathbb{E}\left[\eta|s_\eta\right] \qquad \Leftrightarrow \qquad \rho\sigma_\omega = \sigma_\eta.$$

□

Proof of Lemma 2.2. Let $u \equiv u^R = u^S$. Recall from Lemmas 2.1 and 2.4 that $\theta = \mathbb{E}[\eta|s_\omega, s_\eta]$ and that the conditional distribution of η given s_ω, s_η is symmetric about θ. We first show that the sender's preferences over messages depend only on the distance between induced actions and θ. Let $x' - \mathbb{E}[\eta|\, s_\omega, s_\eta] = \mathbb{E}[\eta|\, s_\omega, s_\eta] - x'' \equiv z > 0$, then

$$\int u(x'-\eta)\, f(\eta|\, s_\omega, s_\eta)\, d\eta = \int u(z - (\eta - \mathbb{E}[\eta|\, s_\omega, s_\eta]))\, f(\eta|\, s_\omega, s_\eta)\, d\eta.$$

The random variable $\hat{\eta} \equiv \eta - \mathbb{E}[\eta|\, s_\omega, s_\eta]$ has mean zero and follows a symmetric distribution. Let $\hat{f}(\hat{\eta}|\, s_\omega, s_\eta)$ denote the standardized distribution (with mean zero). Then, we have

$$f(\eta|\, s_\omega, s_\eta) = \hat{f}(\eta - \mathbb{E}[\eta|\, s_\omega, s_\eta]|\, s_\omega, s_\eta) = \hat{f}(\hat{\eta}|\, s_\omega, s_\eta).$$

Take two realizations $\hat{\eta}'$ and $\hat{\eta}'' = -\hat{\eta}'$ of $\hat{\eta}$. By construction, we have $|z-\hat{\eta}'| = |-z-\hat{\eta}''|$ and hence by symmetry of u around 0, $u(z-\hat{\eta}') = u(-z-\hat{\eta}'')$. Symmetry of the distribution around zero is equivalent to $\hat{f}(\hat{\eta}'|\, s_\omega, s_\eta) = \hat{f}(\hat{\eta}''|\, s_\omega, s_\eta)$. Therefore, for all $\hat{\eta}'$ we have $u(z-\hat{\eta}')\, \hat{f}(\hat{\eta}'|\, s_\omega, s_\eta) = u(\hat{\eta}'-z)\, \hat{f}(\hat{\eta}'|\, s_\omega, s_\eta)$, implying that

$$\int u(z-\hat{\eta})\, \hat{f}(\hat{\eta}|\, s_\omega, s_\eta)\, d\hat{\eta} = \int u(\hat{\eta}-z)\, \hat{f}(\hat{\eta}|\, s_\omega, s_\eta)\, d\hat{\eta}.$$

By symmetry of the distribution, $\hat{\eta}$ and $-\hat{\eta}$ follow the exact same distribution, and we can write

$$\int u(\hat{\eta}-z)\, \hat{f}(\hat{\eta}|\, s_\omega, s_\eta)\, d\hat{\eta} = \int u(-z-\hat{\eta})\, \hat{f}(\hat{\eta}|\, s_\omega, s_\eta)\, d\hat{\eta}.$$

Hence,

$$\begin{aligned}\int u(z-\hat{\eta})\, \hat{f}(\hat{\eta}|\, s_\omega, s_\eta)\, d\hat{\eta} &= \int u(-z-\hat{\eta})\, \hat{f}(\hat{\eta}|\, s_\omega, s_\eta)\, d\hat{\eta} \\ &= \int u(-z-(\eta - \mathbb{E}[\eta|\, s_\omega, s_\eta]))\, f(\eta|\, s_\omega, s_\eta)\, d\eta \\ &= \int u(x''-\eta)\, f(\eta|\, s_\omega, s_\eta)\, d\eta,\end{aligned}$$

that is, the sender is indifferent between actions that are equidistant from θ. By concavity of the sender's payoff, the sender prefers action x' over x'' if and only if x' is closer to θ.

Suppose now that an equilibrium transmits more information to the receiver than θ. Then it must be that the sender is indifferent between the induced actions, x', x'', that is they must satisfy $\left|\theta - x'\right| = \left|x'' - \theta\right|$.

Suppose for sender type θ, the equilibrium induces two actions, equidistant from θ, with some distance $\varepsilon > 0$. Now take a type $\tilde{\theta} = \theta + \delta$ for some $\delta > 0$. We distinguish three cases. Suppose first type $\tilde{\theta}$ induces one action $\overline{x}(\tilde{\theta}) \geq \tilde{\theta}$. Then, to discourage any deviation, we need to have

$$\underline{x}(\theta) < \theta < \overline{x}(\theta) \leq \tilde{\theta} \leq \overline{x}(\tilde{\theta}). \tag{A2.5}$$

However, for all $\delta < \varepsilon$ we have $\overline{x}(\theta) = \theta + \varepsilon > \theta + \delta = \tilde{\theta}$ contradicting condition (A2.5) and implying that some types have a strict incentive to lie. If type $\tilde{\theta}$ induces one action $\underline{x}(\tilde{\theta}) \leq \tilde{\theta}$, or two actions that are equidistant from $\tilde{\theta}$ and satisfy $\underline{x}(\tilde{\theta}) \leq \tilde{\theta} \leq \overline{x}(\tilde{\theta})$, then condition (A2.5) needs to be amended to

$$\underline{x}(\theta) < \theta < \overline{x}(\theta) \leq \underline{x}(\tilde{\theta}) \leq \tilde{\theta} \left(\leq \overline{x}(\tilde{\theta})\right),$$

where the last inequality is absent if type $\tilde{\theta}$ induces only one action $\underline{x}(\tilde{\theta}) \leq \tilde{\theta}$. Since the new condition is even more difficult to satisfy than (A2.5), the same reasoning applies. □

Proof of Lemma 2.3. Letting $a \equiv \frac{\sigma^2_{\varepsilon\omega}}{\sigma^2_{\omega}}$ and $b \equiv \frac{\sigma^2_{\varepsilon\eta}}{\sigma^2_{\eta}}$ we can rewrite $Cov(\omega,\theta)$ and $Var(\theta)$ as

$$Cov(\omega,\theta) = \sigma_{\omega\eta}\frac{a+b+1-\rho^2}{(1+a)(1+b)-\rho^2},$$

and

$$Var(\theta) = \sigma^2_{\eta}\frac{a+b\rho^2+1-\rho^2}{(1+a)(1+b)-\rho^2}.$$

Consider first the set of feasible levels of $Cov(\omega,\theta) = C$. Note that for $a = 0$ or $b = 0$, the covariance is constant and equal to $\sigma_{\omega\eta}$. Moreover, the covariance is decreasing in a for given b and decreasing in b for given a. By l'Hôpital's rule, we have

$$\lim_{b\to\infty}\frac{a+b+1-\rho^2}{(1+a)(1+b)-\rho^2} = \frac{1}{1+a},$$

and

$$\lim_{a\to\infty}\frac{a+b+1-\rho^2}{(1+a)(1+b)-\rho^2} = \frac{1}{1+b}.$$

So, letting both a and b (in whatever order) go to infinity results in a covariance of zero. By continuity, any $C \in \left(0, \sigma_{\omega\eta}\right]$ can be generated by finite levels a, b. Including the case where no signal is observed at all, we can generate all $C \in \left[0, \sigma_{\omega\eta}\right]$.

Consider next the set of feasible $Var(\theta)$ for any given level $Cov(\omega,\theta) = C$. Distinguish two cases, i) $C = \sigma_{\omega\eta}$ and ii) $C \in \left[0, \sigma_{\omega\eta}\right)$.

Case i) requires that $a = 0$ or $b = 0$ or both. If $b = 0$, then $\frac{a+b\rho^2+1-\rho^2}{(1+a)(1+b)-\rho^2} = 1$ and thus $Var(\theta) = \sigma_\eta^2$ for all a. If $a = 0$, then

$$Var(\theta) = \sigma_\eta^2 \frac{b\rho^2 + 1 - \rho^2}{(1+b) - \rho^2}$$

is decreasing in b and attains value $Var(\theta) = \sigma_\eta^2$ for $b = 0$. Moreover,

$$\lim_{b\to\infty} \frac{b\rho^2 + 1 - \rho^2}{(1+b) - \rho^2} = \rho^2.$$

Hence, for $C = \sigma_{\omega\eta}$, $Var(\theta) \in \left[\rho^2\sigma_\eta^2, \sigma_\eta^2\right]$; the lower limit is included because we allow for the case where only one signal is observed.

Case ii) $C \in \left[0, \sigma_{\omega\eta}\right)$ requires that $a > 0$ and $b > 0$. Let $\gamma \equiv \frac{C}{\sigma_{\omega\eta}} \in [0,1)$. The combinations of a and b that generate C satisfy

$$\frac{a + b + 1 - \rho^2}{(1+a)(1+b) - \rho^2} = \gamma.$$

Solving for a as a function of b, we obtain

$$a\left(b;\gamma\right) = \frac{\left(1-\gamma\right)\left(1 + b - \rho^2\right)}{\gamma b - \left(1-\gamma\right)} = \frac{\left(1 + b - \rho^2\right)}{\frac{\gamma}{1-\gamma}b - 1}.$$

The function $a\left(b;\gamma\right)$ is decreasing in b and has the limit

$$\lim_{b\to\infty} \frac{1 + b - \rho^2}{\frac{\gamma}{1-\gamma}b - 1} = \frac{1-\gamma}{\gamma}.$$

In the limit as $b \to \frac{1-\gamma}{\gamma}$, we obtain $a \to \infty$. Hence, C can be generated for $b > \frac{1-\gamma}{\gamma}$ and $a = \frac{(1+b-\rho^2)}{\frac{\gamma}{1-\gamma}b-1}$. Substituting for $\frac{(1+b-\rho^2)}{\frac{\gamma}{1-\gamma}b-1}$ into $Var(\theta)$, we obtain

$$Var\left(\theta; b, a\left(b;\gamma\right), \gamma\right) = \sigma_\eta^2 \frac{\frac{(1+b-\rho^2)}{\frac{\gamma}{1-\gamma}b-1} + b\rho^2 + 1 - \rho^2}{\left(1 + \frac{(1+b-\rho^2)}{\frac{\gamma}{1-\gamma}b-1}\right)(1+b) - \rho^2} = \sigma_\eta^2 \frac{b\gamma\rho^2 + 1 - \rho^2}{1 + b - \rho^2}.$$

The derivative of this expression in b is $\frac{(\gamma\rho^2-1)(1-\rho^2)}{(1+b-\rho^2)^2} < 0$, so $Var\left(\theta; b, a\left(b;\gamma\right), \gamma\right)$ is continuous and monotone decreasing in b. In the limit as b tends to infinity, we obtain

$$\lim_{b\to\infty}\sigma_\eta^2\frac{b\gamma\rho^2+1-\rho^2}{1+b-\rho^2}=\sigma_\eta^2\gamma\rho^2=\sigma_\eta^2\frac{C}{\sigma_{\omega\eta}}\rho^2=\frac{\sigma_\eta}{\sigma_\omega}\rho C.$$

In the limit as $b\to\frac{1-\gamma}{\gamma}$, we obtain

$$\lim_{b\to\frac{1-\gamma}{\gamma}}\sigma_\eta^2\frac{b\gamma\rho^2+1-\rho^2}{1+b-\rho^2}=\sigma_\eta^2\frac{\frac{1-\gamma}{\gamma}\gamma\rho^2+1-\rho^2}{1+\frac{1-\gamma}{\gamma}-\rho^2}=\gamma\sigma_\eta^2=\frac{\sigma_\eta}{\sigma_\omega}\frac{1}{\rho}C.$$

Hence, we have shown that for any given $C\in\left[0,\sigma_{\omega\eta}\right)$, $Var(\theta)\in\left[\frac{\sigma_\eta}{\sigma_\omega}\rho C,\frac{\sigma_\eta}{\sigma_\omega}\frac{1}{\rho}C\right]$. We include the lower limit, because the case where $b\to\infty$ is equivalent to the case with one signal only. □

Proof of Theorem 2.1. Let $u\equiv u^R=u^S$, $C\equiv Cov(\omega,\theta)$, and $V\equiv Var(\theta)$. We prove the theorem in two steps. In step i) we derive the standardized distributions. In step ii) we solve the maximization problem.

i) Let $f_{\omega\theta}(\omega,\theta)=\int f(\omega,\eta,\theta)\,d\eta$ and let $f_{\eta\theta}(\omega,\theta)=\int f(\omega,\eta,\theta)\,d\omega$ denote the marginal joint densities of ω,θ and η,θ. Consider first the expected utility of the sender.

Let $\tau\equiv\frac{C}{V}\theta-\eta$ and let $g(\cdot)$ denote the density of τ. The expected utility of the sender satisfies

$$\begin{aligned}&\int\int u\left(\frac{C}{V}\theta-\eta\right)f_{\eta\theta}(\eta,\theta)\,d\eta d\theta=\int\int u(\tau)\,f_{\eta\theta}\left(\frac{C}{V}\theta-\tau,\theta\right)d\tau d\theta\\=&\int u(\tau)\int f_{\eta\theta}\left(\frac{C}{V}\theta-\tau,\theta\right)d\theta d\tau=\int u(\tau)\,g(\tau)\,d\tau=\int u(\sigma_\tau t)\,c\phi(t)\,dt.\end{aligned}$$

For the first equality, substitute τ and apply the switch of variables theorem. For the second, apply Fubini's theorem. For the third, note that $\Pr\left[\frac{C}{V}\theta-\eta\le\tau\right]=\Pr\left[\frac{C}{V}\theta-\tau\le\eta\right]$ and that by Leibniz's rule

$$g(\tau)=\frac{\partial}{\partial\tau}\Pr\left[\frac{C}{V}\theta-\eta\le\tau\right]=\frac{\partial}{\partial\tau}\int_{-\infty}^{\infty}\int_{\frac{C}{V}\theta-\tau}^{\infty}f_{\eta\theta}(\eta,\theta)\,d\eta d\theta=\int_{-\infty}^{\infty}f_{\eta\theta}\left(\frac{C}{V}\theta-\tau,\theta\right)d\theta.$$

Since τ is a linear function of θ and η, we can use Fang et al. (1990) Theorem 2.16 to conclude that $g(\tau)$ is the density of an elliptical distribution that has the same characteristic generator, $\phi(\cdot)$, as f has. The variance of τ is

$$\begin{aligned}\sigma_\tau^2 &= \frac{C^2}{V^2}Var(\theta) - 2\frac{C}{V}Cov(\theta,\eta) + Var(\eta)\\ &= \frac{C^2}{V} - 2C + \sigma_\eta^2.\end{aligned}$$

Standardizing to $t = \frac{\tau}{\sigma_\tau}$, we transform to a spherical (standardized elliptical) distribution with density $c\phi(\cdot)$.

To derive the receiver's expected utility, we let $\zeta \equiv \frac{C}{V}\theta - \omega$, let $h(\cdot)$ denote the density of ζ. Going through the exact same steps one finds that $h(\zeta) = \int f\left(\frac{C}{V}\theta - \zeta, \theta\right) d\theta$, again an elliptical density with the same characteristic generator. The variance of ζ is

$$\sigma_\zeta^2 = \sigma_\omega^2 - \frac{C^2}{V}.$$

Hence, with $z = \frac{\zeta}{\sigma_\zeta}$, we can write

$$\int\int u\left(\frac{C}{V}\theta - \omega\right) f_{\omega\theta}(\omega,\theta)\, d\omega d\theta = \int u\left(z\sigma_{\omega|\theta}\right) c\phi(z)\, dz.$$

ii) An optimal information structure solves:

$$\max_{C,V} \int u\left(z\left(\sigma_\omega^2 - \frac{C^2}{V}\right)^{\frac{1}{2}}\right) c\phi(z)\, dz + \int u\left(\left(\frac{C^2}{V} - 2C + \sigma_\eta^2\right)^{\frac{1}{2}} t\right) c\phi(t)\, dt.$$

We solve the problem by maximizing sequentially wrt C and V. For given C, the derivative wrt V is

$$\begin{aligned}&\frac{1}{2}\frac{C^2}{V^2}\int z\left(\sigma_\omega^2 - \frac{C^2}{V}\right)^{-\frac{1}{2}} u'\left(z\left(\sigma_\omega^2 - \frac{C^2}{V}\right)^{\frac{1}{2}}\right) c\phi(z)\, dz\\ &-\frac{1}{2}\frac{C^2}{V^2}\int \left(\frac{C^2}{V} - 2C + \sigma_\eta^2\right)^{-\frac{1}{2}} tu'\left(\left(\frac{C^2}{V} - 2C + \sigma_\eta^2\right)^{\frac{1}{2}} t\right) c\phi(t)\, dt.\end{aligned} \tag{A2.6}$$

Recall that $\sigma_\eta^2 = \sigma_\omega^2$. First, suppose $V = C$. Then, the derivative wrt V satisfies

$$\begin{aligned}&\int \frac{1}{2} z\left(\sigma_\omega^2 - V\right)^{-\frac{1}{2}} u'\left(z\left(\sigma_\omega^2 - V\right)^{\frac{1}{2}}\right) c\phi(z)\, dz\\ &-\int \frac{1}{2}\left(-V + \sigma_\eta^2\right)^{-\frac{1}{2}} tu'\left(\left(-V + \sigma_\eta^2\right)^{\frac{1}{2}} t\right) c\phi(t)\, dt\\ =\ & 0.\end{aligned}$$

Now suppose $V \neq C$. Note that both integrands in (A2.6) have the common representation

$$\int \frac{1}{a} k u'(ak)\, c\phi(k)\, dk. \tag{A2.7}$$

Differentiating wrt a, we observe that (A2.7) is monotone decreasing in a,

$$-\frac{1}{a^3}\int aku'(ak)\, c\phi(k)\, dk + \frac{1}{a^3}\int a^2k^2u''(ak)\, c\phi(k)\, dk \leq 0,$$

where the inequality follows from the curvature condition

$$q\frac{u''(q)}{u'(q)} = q\frac{\ell''(q)}{\ell'(q)} \geq 1. \tag{A2.8}$$

$V < C$ implies $\frac{C^2}{V} - 2C + \sigma_\eta^2 > \sigma_\omega^2 - \frac{C^2}{V}$. The curvature condition (A2.8) implies monotonicity and therefore

$$\begin{aligned} & \frac{1}{2}\frac{C^2}{V^2}\int z\left(\sigma_\omega^2 - \frac{C^2}{V}\right)^{-\frac{1}{2}} u'\left(z\left(\sigma_\omega^2 - \frac{C^2}{V}\right)^{\frac{1}{2}}\right) c\phi(z)\, dz \\ \geq\ & \frac{1}{2}\frac{C^2}{V^2}\int \left(\frac{C^2}{V} - 2C + \sigma_\eta^2\right)^{-\frac{1}{2}} tu'\left(\left(\frac{C^2}{V} - 2C + \sigma_\eta^2\right)^{\frac{1}{2}} t\right) c\phi(t)\, dt. \end{aligned}$$

Hence the derivative is non-negative for $V < C$. By symmetry, the derivative is non-positive for $V > C$. These inequalities become strict for functions that satisfy the curvature condition (A2.8) with strict inequality. It follows that the problem is maximized in V for $V = C$.

The second step is now to maximize over C, given that $V = C$.

$$\max_C \int u\left(z\left(\sigma_\omega^2 - C\right)^{\frac{1}{2}}\right) c\phi(z)\, dz + \int u\left(\left(\sigma_\eta^2 - C\right)^{\frac{1}{2}} t\right) c\phi(t)\, dt.$$

The derivative wrt C is given by

$$\begin{aligned} & -\int \frac{1}{2} z\left(\sigma_\omega^2 - V\right)^{-\frac{1}{2}} u'\left(z\left(\sigma_\omega^2 - V\right)^{\frac{1}{2}}\right) c\phi(z)\, dz \\ & -\int \frac{1}{2}\left(\sigma_\eta^2 - C\right)^{-\frac{1}{2}} tu'\left(\left(\sigma_\eta^2 - C\right)^{\frac{1}{2}} t\right) c\phi(t)\, dt \\ > \ & 0. \end{aligned}$$

The payoff is unambiguously increasing in C. The solution is thus $C = C^{\max}$. □

Proof of Proposition 2.1. By Theorem 2.1, for quadratic loss functions all information structures satisfying $Cov(\omega,\theta) = \sigma_{\omega\eta}$ are optimal for θ public. By Theorem 2.2, smooth communication is an equilibrium if and only if $Cov(\omega,\theta) = Var(\theta)$. By Lemma 2.3, the candidate solution $Var(\theta)^* = Cov(\omega,\theta)^* = \sigma_{\omega\eta}$ is feasible if $\frac{Cov(\omega,\theta)^*}{Var(\theta)^*} = 1 \in \left[\frac{\sigma_\eta}{\sigma_\omega}\rho, \frac{\sigma_\eta}{\sigma_\omega}\frac{1}{\rho}\right]$. This is guaranteed by the assumption $\min\left\{\sigma_\omega^2, \sigma_\eta^2\right\} \geq \sigma_{\omega\eta}$. □

3 A Bandit Model of Two-Dimensional Uncertainty – Rationalizing Mindsets

3.1 Introduction

Will I succeed if I work hard? Or, is ability more important than effort? Do I change my effort in the face of failure? These questions are often encountered, e.g. by kids at school and workers at the workplace. As it is well-known, people differ enormously in their response to failures. Our goal is to understand how failure impacts on the effort choice when the own ability as well as the production function for the task at hand are unknown.

In this chapter, we address these questions in a theoretical framework of optimal experimentation. We introduce a new type of bandit where the agent is confronted with two-dimensional uncertainty. There are two possible tasks; in task E high effort, whereas in task A high ability is necessary for a success. The agent repeatedly faces the same task but does not know whether the task is E or A, i.e. whether effort or ability is required to succeed. In addition, the agent is uncertain about her own type, whether she has high or low ability. In every period, the agent chooses whether to exert effort at known cost or not. After both decisions she observes a success or a failure, depending on the task and possibly on her choice of effort and her ability. Hence differently from the standard model, there are always outcomes and information in each period, regardless of effort. Nevertheless, the informational content of the outcome depends on the effort choice. Whether effort or no effort conveys more information depends on the agent's beliefs about her ability and the task. In our model, the agent updates her belief about the type of the bandit (task A or E) and her own type (ability L or H) at the same time and her belief about one influences her learning about the other.

We find four different patterns of behavior when analyzing our model with two periods. For low cost, the agent persistently exerts effort regardless of observing a success or a failure and despite the fact that with a positive probability effort is not needed to succeed. For medium cost and a low belief about her ability, the agent starts by exerting effort but gives up when facing a failure. By contrast, an agent with medium cost and a high belief about her ability initially does not exert effort, but starts to exert effort when facing a failure. Finally, for very high cost of effort the agent never exerts effort even though with a positive probability effort is necessary and a success would outweigh the cost. When considering an infinite-time horizon only the number but not the order of outcomes is relevant. After observing a success once, the agent will stick to her effort choice forever. An agent that only observes failures, tolerates only a maximal number of failures with effort before she stops exerting effort forever. Before this maximal number is reached, the agent may start and stop exerting effort repeatedly.

The intuition for our results relies on the fact that the choice of effort determines the kind of information for future periods. The outcome after each period is the pair of effort choice (yes or no) and result (success or failure). The agent observes different outcomes when exerting effort and not exerting effort, and both types of outcomes may be informative about the task and the agent's ability. Besides the agent's cost, also her beliefs determine the optimal effort choice. For example, we find that for the same effort cost, agents with low-ability belief rely more on effort, while agents with high ability belief initially avoid exerting effort in the hope that their ability will be sufficient. Moreover, a higher belief that the task is E results in a higher willingness to exert effort. Hence, in our model three factors determine the optimal effort strategy: the agent's cost of exerting effort, her beliefs about the production function as well as the beliefs about her ability.

In this chapter, we give a theoretical explanation for different reactions to failure. In applied research in educational psychology, in particular Dweck (2006) attributes diverging behavior in response to failure to different mindsets. Agents that have a "fixed mindset" believe that success is based on innate ability, whereas agents with "growth mindset" believe that success comes from hard work. Consequently, when facing a failure fixed types stop exerting effort whereas growth types start exerting more effort, as documented in this literature (e.g. Dweck (2000)). Given that it often is not clear whether effort or ability are needed for a success, at a first glance it seems intuitive that having a growth mindset and not giving up is desirable. But then, why do we find both types of behavior? Dweck (2006) emphasizes the role of education and feedback to establish the mindset beliefs. In our model, we show that observed responses to failure can also be explained as the result of different effort costs and beliefs about the own ability together with Bayesian updating about the sources of success.

In a standard bandit model, the agent has the option of pulling the arm of a bandit machine with an unknown output distribution. Pulling the arm is costly either because it requires costly effort or because of the opportunity cost of forgoing the known output of a safe arm. On the other hand, by pulling the risky arm the agent receives information about its output distribution which can be beneficial in the future. No output is generated when the risky arm is resting and no information is revealed. The first bandit problem in economics is developed in Rothschild (1974); a single firm has to determine the optimal price in a market with unknown demand. Weitzman (1979) studies where to allocate effort optimally when different opportunities with unknown rewards are available. Berry and Fristedt (1985) gives a summary of results for bandit problems. For a survey of the literature on multi-armed bandits see Bergemann and Valimaki (2006).

Similarly to a standard bandit model, in our model the output distribution is unknown and exerting effort is costly. By contrast to the literature, the agent receives an output and information with and without exerting effort in every period. Moreover, the agent faces a two-dimensional uncertainty, the production function as well as the own ability are unknown. We are not aware of any model that shares these features.

A bandit model where the inactive arm evolves over time, a "restless bandit", is first introduced in Whittle (1988). Fryer and Harms (2015) models human capital formation as a restless bandit. In their model the bandit is "bi-directional" since payoffs go up when the arm is used but they go down when the arm rests. The authors show that stopping rules are optimal. In our model, the bandit does not evolve over time but the agent obtains information in every period. The agent's beliefs change depending on the number of observed outcomes and the agent may repeatedly start and stop exerting effort.

In Heidhues et al. (2015) the agent's outcome depends on her action, her ability and some external factor. The agent is assumed to be overconfident about her ability and the authors analyze the impact of the overconfidence on the inferences the agent draws about other variables. We consider a rational agent that simultaneously learns about her ability and the production function and her belief about the one influences the learning about the other.

Bolton and Harris (1999) and Keller et al. (2005) analyze two-armed bandit models with many agents in continuous time. In Bolton and Harris (1999) the uncertainty of the risky arm is driven by a Brownian motion and both good and bad news arrive continuously. Keller et al. (2005) considers a model with exponential distributed uncertainty. An arm could be either good or bad, where the good arm has a certain arrival rate of a breakthrough, while at the bad arm a breakthrough never occurs. The analysis is complemented in Keller and Rady (2010), where high payoffs arrive in both states of the world, but the arrival rate is higher if the state is good. Since information is a public good in these contributions, a free-riding problem arises. However, the latter paper shows that the presence of other players encourages at least one of them to continue experimenting with the risky arm. By contrast, in our model time is discrete and the uncertainty does not follow a stochastic process. In the two-agents extension of our model we not only observe free-riding but also co-ordination as equilibrium behavior, some types of agents prefer to exert effort if and only if the other agent exerts effort as well.

In Lizzeri and Siniscalchi (2008) a child tries to learn the mean of his normally distributed task by choosing an action in every period. The parent of the child may manipulate the outcome; she faces the trade-off between allowing the child to learn from his mistakes and sheltering him from the consequences. By contrast, in our model the task is constant over time and a single player learns about the task and the own ability simultaneously.

In career concerns models (e.g. Holmström (1999), Dewatripont et al. (1999)) an agent works in a competitive labor market and is paid his expected output that depends on his private choice of effort and his unknown ability. In the earlier periods the agent has an incentives to exert effort to increase the beliefs about his ability. Bonatti and Hörner (2015) considers a continuous-time career concerns model, where ability and effort are complements. In addition, effort levels at different dates are strategic substitutes, increasing market expectation decreases incentives at earlier stages. In a career-concerns version of our model, we find that indeed an agent may exert effort in the first period to induce higher beliefs about her ability for the second period.

Instead of inducing the agent to exert more effort, in Manso (2011) the optimal incentive scheme motivates the agent to innovate; early failure is tolerated and long-term success is rewarded. In our model future payoffs are discounted but do not change over time.

Similarly to our trade-off between effort and ability, in Piketty (1995) a continuum of agents tries to learn from their experienced income mobility whether predetermined (social) factors or individual effort are more important for high income. The author explains different preferences over redistribution by showing that in any long-run steady state in his model, types learn the importance of predetermined factors on average but stay with different estimates about the influence of individual effort. In our model some types of agents learn the production function and the role of effort perfectly, while the other agents stay uninformed. Depending on effort cost and beliefs some of the latter continue exerting effort forever, while others never exert effort.

We proceed as follows. We state our model in Section 3.2. In Section 3.3 we analyze the case where outcomes are certain. By contrast, in Section 3.4 we allow for uncertainty of the outcome. We shortly discuss possible extensions in Section 3.5 and conclude in Section 3.6. All proofs are gathered in the Appendix.

3.2 Model

There is one agent who has unknown ability $a \in \{L, H\}$. Her prior belief that her ability is high is $\alpha = Pr(a = H) \in (0,1)$. Time is discrete, there are $0 < t \leq T$ periods, for $T \in \mathbb{N} \cup \{\infty\}$, and we discount the output of future periods with discount factor $\delta \in (0,1)$. In each period, the agent chooses a level of effort $e \in \{Y, N\}$. Exerting effort, $e = Y$, has cost $c \in (0, 1-\varepsilon)$, while not exerting effort is costless.[1] The parameter $\varepsilon \in [0,1)$ denotes the probability of having "bad luck", i.e. with probability ε the agent observes a failure instead of a merited success. There are two possible tasks, E and A, as shown in Table 3.1. In task E, the agent is successful with probability $1-\varepsilon$ only if she exerts effort, ability is irrelevant. By contrast, in task A high ability determines success and effort is irrelevant.[2] The agent's prior belief is that with probability $q \in (0,1)$ the task is E and with probability $1-q$ the task is A. Hence, there are three distinguishable states, E, AH and AL that occur with probabilities q, $(1-q)\alpha$, and $(1-q)(1-\alpha)$, respectively. The state is fixed for all periods of the game.

E	$a = L$	$a = H$
$e = N$	0	0
$e = Y$	$1-\varepsilon$	$1-\varepsilon$

A	$a = L$	$a = H$
$e = N$	0	$1-\varepsilon$
$e = Y$	0	$1-\varepsilon$

Table 3.1. Probabilities of success in tasks E and A.

After each period, the agent observes either failure F or success S. For $e = Y$ the possible outcomes are FY (failure, effort) or SY (success, effort). For $e = N$ the possible outcomes are FN (failure, no effort) or SN (success, no effort). The agent's payoff

[1] The assumption that c is bounded by $1-\varepsilon$ ensures that exerting effort may be valuable for the agent. For effort cost above this bound the agent never exerts effort.

[2] We analyze the agent's behavior when effort and ability are perfect substitutes and perfect complements, and state a continuous version as possible benchmarks for $\varepsilon = 0$ in Appendix B.

of success is 1, her payoff of failure is 0. However, due to the possibility of bad luck the expected payoff of success is $1-\varepsilon$. After observing the outcome, the agent updates her beliefs about the task and her ability and decides about her level of effort in the next period.

We split our analysis in two cases. In Section 3.3, the agent is not confronted with the possibility of bad luck, i.e. $\varepsilon = 0$, and we characterize equilibrium behavior for an arbitrary number of periods. In Section 3.4, we allow for uncertainty, i.e. $\varepsilon > 0$, and examine equilibrium behavior for two as well as for infinitely many periods.

3.3 Certainty

In this section we analyze our model for $\varepsilon = 0$ and explain different patterns of behavior depending on the agent's cost and beliefs.

In each period, after observing the outcome the agent updates her beliefs about the task and her ability. In particular, after observing SN the agent knows that the task is A and her ability is H, exerting effort is unnecessary. Similarly, the agent knows after FY that the state is AL and exerting effort is futile. By contrast, after FN and SY the agent remains uncertain about the state and whether effort is necessary or not. After FN she can rule out that the state is AH, while after SY the state cannot be AL. In both cases she updates in favor of state E.

$$\begin{aligned}
Pr\,(E|SN) &= 0,\\
Pr\,(E|FY) &= 0,\\
Pr\,(E|FN) &= \frac{q}{q+(1-q)(1-\alpha)} \geq q,\\
Pr\,(E|SY) &= \frac{q}{q+\left(1-q\right)\alpha} \geq q.
\end{aligned}$$

Note that for large α the updated probability that effort is necessary is higher after FN than after SY,

$$\alpha \geq \frac{1}{2} \Leftrightarrow Pr\,(E|FN) \geq Pr\,(E|SY)\,.$$

Moreover, certainty implies that for the same level of effort the outcome will necessarily always be the same, thus there is no further updating when facing the same outcome repeatedly. Note further that the agent is always fully informed about the state after changing the effort choice once and observing the outcomes.[3] However,

[3] E.g., if the agent exerts no effort after SY, and the result is FN she will know that the task is E. If instead the result is SN she will know that she has high ability and faces task A.

as we will see in Proposition 3.1 some types of agents in equilibrium never learn the actual state.

As a reference it is useful to consider the optimal effort choice when the game is only played once, $T = 1$. In this case the choice of effort does not affect the information in later periods and hence it pays off to exert effort for $q \geq c$. The cost of effort c is weighed against the probability q that the state is E, i.e. that effort is necessary for success. Note that for any finite-time horizon the last period is analogous to the single-period game, since the agent will exert effort if and only if the cost is below the updated probability that the state is E.

Now consider $T > 1$. In period $t = 1$ the choice of effort will not only affect the immediate payoff but also the information available in all future periods. Recall that we discount future periods with δ and that the agent obtains information with as well as without exerting effort. The expected payoff will depend on how the agent expects to behave after receiving the information. We can solve the finite-horizon game by backward induction. Our analysis smoothly extends to an infinite-time horizon, since only the effort choices of the first two periods are relevant.

Remark 1: After the second period, the agent will not change her effort choice until the end of the game.

The reason is twofold. Either the uncertainty about the state is resolved and the agent knows whether effort is required or not, or the agent is still uncertain but nonetheless sticks to her choice of effort. The latter can be explained as follows. If the agent observed FN or SY in the first period, she can get either fully informed by changing the effort choice in the second period or remain uninformed. For FN in $t = 1$ the decision to stay uninformed in $t = 2$ implies that the trade-off between immediate cost and informational gain given the beliefs did resolve in favor of not exerting effort. But since the updated beliefs stay constant given the same outcome and the gain from information (weakly) declines over time, the same result of the trade-off will hold in all future periods and the agent will never exert effort.[4] On the other hand, the outcome SY shows that it was valuable to exert effort in the first period despite the outcome being uncertain. But after SY success after effort is certain. Therefore, the agent will continue to exert effort forever.

The remark immediately implies that the agent will never change her choice of effort after a success. The agent's updated beliefs define cost thresholds for the decision of effort in the second period. For costs identical to or below these beliefs the

[4] For $T = \infty$ the informational gain stays constant and the same argument holds.

agent exerts effort.[5] For $T = 2$, after observing FY or SN the agent has learned that effort is irrelevant for success, therefore her cost thresholds are $c_{FY} \equiv 0$ and $c_{SN} \equiv 0$. By contrast, after observing SY or FN the agent updates her belief about the task being E upward and would therefore exert effort for higher cost than in the one-shot game, $c_{SY} \equiv \frac{q}{q+(1-q)\alpha} \geq q$ and $c_{FN} \equiv \frac{q}{q+(1-q)(1-\alpha)} \geq q$. For $T > 2$, we additionally have to incorporate the gain from information for all future periods. The following cost thresholds depend on T and are useful to characterize the agent's equilibrium behavior. For any $T \geq 2$, the cut-off c_{FNT} (c_{SYT}) determines the agent's upper-cost limit for exerting effort in the second period after observing FN (SY) in $t = 1$.

$$c_{FNT} \equiv \frac{c_{FN}(1-\delta^{T-1})}{1-\delta+c_{FN}\left(\delta-\delta^{T-1}\right)},$$
$$c_{SYT} \equiv \frac{c_{SY}(1-\delta)}{1-(1-c_{SY})\delta^{T-1}-\delta c_{SY}}.$$

In the first period, the cut-offs c_{1T} and c_{2T} define the agent's effort decision for low and medium costs, respectively. The derivation can be found in the proof of Lemma 3.1.

$$c_{1T} \equiv \frac{q}{1+\delta(1-q)\left(\frac{\alpha(1-\delta^{T-1})}{1-\delta}-(1-\alpha)\right)},$$
$$c_{2T} \equiv \frac{q(1-\delta^{T})}{(1-\delta)+(q(1-\alpha)+\alpha)\delta(1-\delta^{T-1})}.$$

The following lemma shows three different patterns of equilibrium behavior.

Lemma 3.1. *Let $\varepsilon = 0$ and $T \leq \infty$. There exist three different types of equilibrium behavior:*

1. *Inert: If the cost is high, $c > \max\{c_{2T}, c_{FNT}\}$, the agent never exerts effort.*
2. *Growth: For medium cost, $c \in (c_{1T}, c_{FNT}]$, the agent does not exert effort in $t = 1$, but does in $t = 2$ after failure. Thereafter, the agent is fully informed and exerts effort if and only if the task is E.*
3. *Fixed: For low cost, $c \leq \min\{c_{1T}, c_{2T}\}$, the agent exerts effort in all periods unless she observes a failure after which she stops exerting effort.*

The agent's effort choices are characterized by different cost thresholds. To get an intuition for the result, recall that after the second period the effort choice and

[5] Note that in our model in general the thresholds are not equal to the agent's beliefs. The only exception is the case of certainty and two periods, $\varepsilon = 0$ and $T = 2$.

the corresponding outcome stay constant. In $t = 2$, the agent compares the continuation values from exerting and not exerting effort given the outcome of the first period. This implies the cost thresholds c_{SYT} and c_{FNT} that are depicted in Figure 3.3.1. These thresholds define four different areas in the cost-ability space. In $t = 1$, the agent anticipates her second-period behavior that depends on the respective area.[6] Therefore, her choice of effort in the first period does not only determine the immediate outcome but also influences the information available in the second period; e.g. after $e = Y$ only SY and FY are possible outcomes. Taking also into account the expected resulting continuation values, the comparison of exerting and not exerting effort in $t = 1$ results in the thresholds c_{1T} and c_{2T} that determine the behavior in $t = 1$ for the respective area.

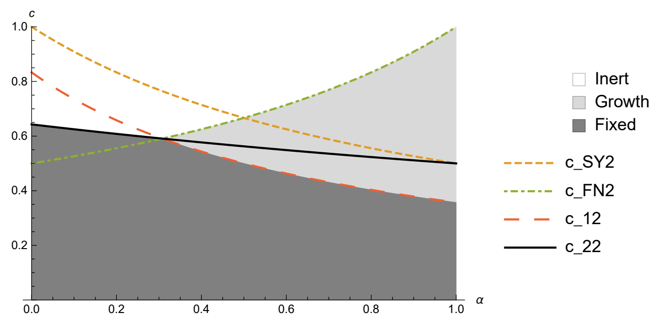

Figure 3.3.1. $T = 2$. Thresholds for effort in terms of c and α for $q = 0.5, \delta = 0.8$.

Figure 3.3.1 illustrates the different types of behavior. In the white area, an agent never exerts effort because her cost is too high. The agent remains inert regardless of the outcome. In the light grey area, an agent does not exert effort in period $t = 1$ but starts exerting effort in case of a failure (FN). After FN the probability that the task is E is increased. Since the agent's ability belief is relatively high and her cost is low enough, she tries to succeed with effort in $t = 2$. In this case, a failure motivates the agent to exert more effort than before. In the grey area, an agent exerts effort in period $t = 1$, since her cost as well as her ability belief are relatively low. She only continues after a success (SY) in $t = 2$, when facing a failure the agent gives up exerting effort.

Given our analysis we obtain the following result.

Proposition 3.1. *For all $T \geq 1$, there exist types of agents that never learn the actual state. An inert-type agent never exerts effort even if the task is E. A fixed-type agent with high ability exerts effort forever even if the task is A.*

6 For example, for $c_{SYT} < c \leq c_{FNT}$ the agent exerts effort in $t = 2$ after observing FN but not after SY.

The proposition follows from Lemma 3.1 and the observation that the agent will never change her choice of effort after observing a success.

As q increases all thresholds rise. The agent becomes more willing to exert effort in any period when the likelihood that effort leads to success increases.

Concerning patience, as δ increases both c_{1T} and c_{2T} become steeper in α. Therefore, for low beliefs in ability, $\alpha \leq \frac{\delta}{1+2\delta}$, the agent exerts effort in $t = 1$ also for relatively high cost. On the other hand, for high beliefs in ability, $\alpha \geq \frac{\delta}{1+2\delta}$, the agent becomes less willing to exert effort in $t = 1$ as she becomes more patient. Instead, it becomes more attractive for the agent to test in the first period whether success is also possible without effort. This would potentially allow her to save the cost of effort in the future. However, after a failure, she will start exerting effort.

The threshold c_{FNT} is increasing in T, while c_{SYT} is decreasing in T. This implies that for large T (and δ not too small) the white area where the agent never exerts effort and the grey area where the agent gives up exerting effort after a failure shrink, whereas the light grey area where the agent does not exert effort in the first period but exerts effort in the second period after FN grows. The longer the time horizon the lower the level of beliefs in ability necessary to motivate the agent to experiment whether effort is needed for success. Figure 3.3.2 illustrates equilibrium behavior for $T = \infty$.

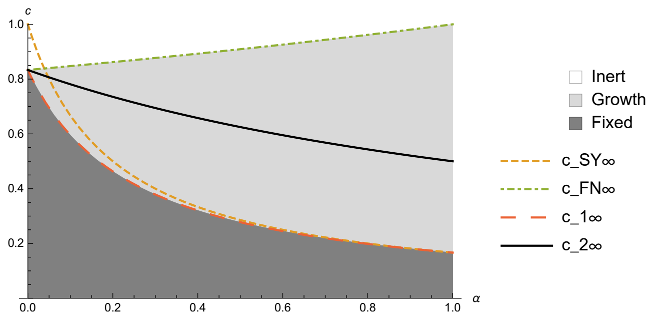

Figure 3.3.2. $T = \infty$. Thresholds for effort in terms of c and α for $q = 0.5, \delta = 0.8$.

Going back to the motivation given in the introduction, we relate the agent's response to failure to the two mindsets. Agents with growth mindset (light grey) increase their effort level, while agents with fixed mindset (grey) stop exerting effort after observing a failure.

3.4 Uncertainty

In this section we allow for uncertainty of the outcome in the sense that with probability $\varepsilon \in (0, 1)$ the agent has bad luck and observes a failure instead of a merited success.[7] We first analyze a finite-time horizon with two periods, then we consider an infinite-time horizon.

Consider $T = 2$. In the second period, the agent updates her beliefs about the task and her ability. In particular, after observing SN the agent knows that the task is A and her ability is H, so exerting effort is unnecessary. After observing SY, the agent knows that the state cannot be AL and she updates in favor of states E and AH. In contrast, after a failure the agent cannot rule out any state with certainty. After FY she becomes less confident while after FN she becomes more confident that the state is E.

$$\begin{aligned}
Pr\,(E|SN) &= 0,\\
Pr\,(E|SY) &= \frac{q}{q+\left(1-q\right)\alpha} \geq q,\\
Pr\,(E|FY) &= \frac{\varepsilon q}{\varepsilon\left[q+\left(1-q\right)\alpha\right]+(1-q)(1-\alpha)} \leq q,\\
Pr\,(E|FN) &= \frac{q}{q+\varepsilon(1-q)\alpha+(1-q)(1-\alpha)} \geq q.
\end{aligned}$$

The analysis is analogous to the case of certainty in Section 3.3. The agent will exert effort in the second period for low enough costs. The threshold values for exerting effort incorporate the probability of bad luck and are given by

$$c_{h\varepsilon} \equiv Pr\,(E|h)\cdot(1-\varepsilon) \quad \text{for } h \in \{SY, FY, SN, FN\}.$$

In the first period, the agent's effort choice affects the immediate cost as well as the expected informational gains for the second period. The cost thresholds for the first period are derived in the proof of Lemma 3.2, the relevant thresholds are

$$\begin{aligned}
c_{1\varepsilon} &\equiv \frac{q(1-\varepsilon)}{1+(1-q)\alpha\delta(1-\varepsilon)},\\
c_{2\varepsilon} &\equiv \frac{q(1-\varepsilon)(1-\delta\varepsilon)}{1+\delta(-1+2\alpha+q(1-2\alpha)(1-\varepsilon)-2\alpha\varepsilon)},\\
c_{3\varepsilon} &\equiv \frac{q(1+\delta(1-\varepsilon))(1-\varepsilon)}{1+\delta(q(1-\alpha)+\alpha)(1-\varepsilon)}.
\end{aligned}$$

The following lemma is the analog to Lemma 3.1 in Section 3.3.

Lemma 3.2. *Let $\varepsilon \in (0,1)$ and $T = 2$. There are four different types of equilibrium*

[7] For the case of two periods, our results qualitatively do not change when we additionally allow for the possibility that the agent has "good luck", i.e. she observes a success instead of a merited failure. We expect the same for the case of infinitely many periods.

behavior:

1. *Inert: If the cost is high,* $c > \max\{c_{3\varepsilon}, c_{FN\varepsilon}\}$*, the agent does not exert effort in either period.*

2. *Growth: For medium high cost,* $c \in (\min\{c_{2\varepsilon}, c_{1\varepsilon}\}, c_{FN\varepsilon}]$*, the agent does not exert effort in* $t = 1$*, but does exert effort in* $t = 2$ *after failure.*

3. *Fixed: For medium low cost,* $c \in (c_{FY\varepsilon}, \min\{c_{3\varepsilon}, c_{2\varepsilon}\}]$*, the agent exerts effort in* $t = 1$ *and stops exerting effort after observing a failure.*

4. *Persistent: For low cost,* $c \leq \min\{c_{FY\varepsilon}, c_{1\varepsilon}\}$ *the agent exerts effort in both periods regardless of the outcome.*

Similarly to the case of certainty, we characterize the agent's effort choices by different cost thresholds. Taking bad luck into account, in $t = 2$ the agent compares her cost with her immediate expected outcome given the observation of the first period. This implies the cost thresholds $c_{SY\varepsilon}, c_{FY\varepsilon}$, and $c_{FN\varepsilon}$ for $t = 2$. The five resulting areas in the cost-ability space influence the agent's effort choice in $t = 1$. The comparison of the sums of the expected outcomes in both periods when exerting and not exerting effort in $t = 1$ results in the thresholds $c_{1\varepsilon}, c_{2\varepsilon}$, and $c_{3\varepsilon}$.

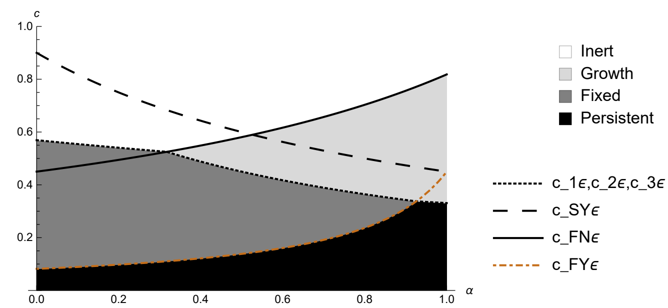

Figure 3.4.1. $T = 2$. Thresholds for effort for $q = 0.5, \delta = 0.8$ and $\varepsilon = 0.1$.

Figure 3.4.1 illustrates the four different behavior patterns that depend on cost c and beliefs α and q. We call them persistent, fixed, growth, and inert. In addition to the patterns that exist under certainty already, a new "persistent" type arises (black area). The agent persistently exerts effort in both periods regardless of the outcome. The reason is that FY does not rule out task E perfectly, with probability ε the agent has had bad luck in $t = 1$. Therefore, types with very low cost try for a second time to succeed with effort.

We now expand the time horizon to $T = \infty$ periods. A history h_t is a sequence of outcomes from periods 1 to t. Recall that after the outcome SN the agent is certain that the state is AH, while after SY the agent is certain that the state is not AL. The updated beliefs of the agent about the state E after a history h_t are given by

$$Pr\left[E|h_t\right] = \mathbb{1}_{\{SN \notin h_t\}} \frac{q\varepsilon^{n_{FY}}}{q\varepsilon^{n_{FY}} + (1-q)\alpha \cdot \varepsilon^{n_{FN}+n_{FY}} + (1-q)(1-\alpha) \cdot \mathbb{1}_{\{SY \notin h_t\}}}, \qquad (3.4.1)$$

where the integers n_{FY} and n_{FN} count the number of experienced failures with and without effort, respectively. Since the updated beliefs only depend on the numbers of the outcomes, the order of outcomes in h_t is irrelevant. Given equation (3.4.1), the outcomes SY and FN increase the updated belief that the task is E, while the outcomes SN and FY decrease it.

As before, at every point in time the agent trades off the immediate cost of exerting effort, which only depends on her belief about the likelihood of the state being E, and the informational returns of exerting effort or not, which in addition depends on her belief about her ability. Differently from the finite-time case, this trade-off is constant over time for $T = \infty$.

We split the analysis in two cases. First, we look at the agent's behavior after observing a success. Second, we focus on the case where the agent has observed failures only.

Remark 2: The agent will not change her effort level after observing a success; she will never exert effort after the outcome SN and she will continue to exert effort after observing SY.

To see this, suppose first that the agent observes the outcome SN in period t. Since this immediately implies that the state is AH, the agent has no incentive to exert effort in any future period. Suppose now that the agent observes the outcome SY in period t. This implies that the state cannot be AL. Furthermore, the agent's updated belief in $t+1$ about the likelihood that the state is E increases. But given that it was profitable for the agent to exert effort in period t, this implies that it also must be profitable for her exert effort in period $t+1$. The same argument holds for period $t+1$ and so forth, and the agent continues to exert effort forever.

We can now calculate the agent's continuation value after a success, i.e. the sums of the expected payoffs of all future periods,

$$V_{SN} = \sum_{s=0}^{\infty} \delta^s (1-\varepsilon) = \frac{1-\varepsilon}{1-\delta},$$
$$V_{SY} = \sum_{s=0}^{\infty} \delta^s (1-\varepsilon-c) = \frac{1-\varepsilon-c}{1-\delta}.$$

We continue to analyze the agent's behavior in the class of histories where the agent has only experienced failures (FY or FN). We find that an agent is willing tolerate only a maximum number $\bar{n}_{FY}$ of outcomes FY until she stops exerting effort forever. Similarly, we can define the minimum number $\underline{n}_{FN}$ of outcomes FN before the agent exerts effort for the $\bar{n}_{FY}$th time. Moreover, the number $\bar{n}_{FY}$ corresponds to the following cost threshold depicted in Figure 3.4.2

$$c_{\infty FN}(\bar{n}_{FY}) \equiv \frac{q(1-\varepsilon)(1-\delta\varepsilon)\varepsilon^{\bar{n}_{FY}}}{(1-q)(1-\alpha)(1-\delta)+q\varepsilon^{\bar{n}_{FY}}(1-\delta\varepsilon)}.$$

Before $\bar{n}_{FY}$ is reached, the agent may repeatedly start and stop exerting effort.

Lemma 3.3. *There exists a threshold $c_{\infty FN}(\bar{n}_{FY})$ such that for effort cost $c \leq c_{\infty FN}(\bar{n}_{FY})$ the agent is willing to observe at most $\bar{n}_{FY}$ times the outcome FY before she stops exerting effort forever. For any $n_{FY} < \bar{n}_{FY}$ the agent exerts effort at the earliest period where postponing effort is disadvantageous, this may occur before or after $\underline{n}_{FN}$ is reached.*

The intuition for the result relies on the fact that the likelihood of the state being E is decreasing in FY and increasing in FN. Indifference between starting to exert effort and continuing not to exert effort after infinitely many FN outcomes defines the cost threshold $c_{\infty FN}(\bar{n}_{FY})$ for any maximal number of FY outcomes. The key step in the proof is the comparison of the expected utilities of exerting effort "today" or "tomorrow". We define the difference as a function $\beta(n_{FY}, n_{FN})$ that is decreasing in n_{FY} and increasing in n_{FN}. The agent exerts effort whenever this comparison is in favor of exerting effort today. Monotonicity implies that the agent may stop after observing FY and start after some number of FN until the maximum $\bar{n}_{FY}$ is reached.

Figure 3.4.2 illustrates how the cost thresholds decline for an increase in $\bar{n}_{FY}$.

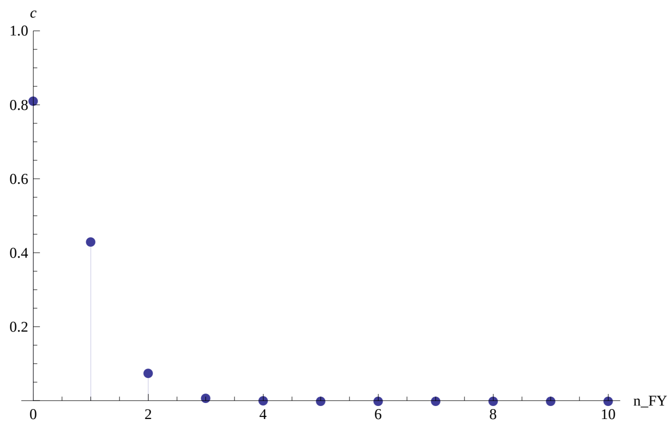

Figure 3.4.2. Cost thresholds $c_{\infty FN}(\bar{n}_{FY})$ for $q = 0.5, \delta = 0.8$, and $\varepsilon = 0.1$.

3.5 Extensions

We shortly illustrate two possible extensions of our model with certainty.[8] First, we consider a competitive labor market where the agent is confronted with a principal. Due to competition between employers and a lack of commitment power on either side, in every period the agent is paid her expected output. Is the agent willing to exert effort to increase the principal's belief about her ability? In other words, do career concerns arise? We find that under the assumption that the task is known to the principal or to the agent, the agent does not exert effort due to career concerns. However, if the task is unknown the agent may have incentives to exert effort in order to protect or inflate her reputation. In the career-concerns version of our model three different equilibria are present. For low cost, the principal correctly anticipates the agent to exert effort, while for high cost, the principal correctly anticipates the agent to not exert effort. For medium cost, the agent's incentives are opposed to the principal's expectations. If the principal expects the agent to exert effort, she would prefer not to and if the principal expects the agent not to exert effort, the agent instead has an incentive to exert effort. Therefore, there is no equilibrium in pure strategies in this case.[9] Pursuing the principal-agent approach allows for other interesting questions, for example how a principal would assign agents to tasks or tasks to agents. We leave this for future research.

Second, we investigate on the following question. How does equilibrium behavior change in the presence of other agents? We add a second agent with identical

---

8 A formal analysis of the findings can be obtained from the authors upon request.

9 More precisely, for medium cost, $\delta \frac{q(1-q)\alpha}{q+(1-q)\alpha} < c \leq \delta q$, the agent will follow a mixed strategy and exert effort with probability $p = \alpha \frac{1-q}{q}\left(\frac{\delta q}{c} - 1\right)$.

beliefs about task and ability to our model. Both agents face the same task and perfectly observe each others effort choices and outcomes. We can show that the equilibrium behavior patterns of the single-agent model continue to exist. For most cost and ability combinations, the agents play dominant strategies in equilibrium. However, for medium cost there are two areas where the optimal strategy depends on the other agent's effort choice. For medium-ability beliefs the agent tries to match the other agent's effort decision (co-ordination), while for low and high beliefs the agent tries to oppose the other agent's effort choice (free-riding). As a next step in a setup with more agents we plan to add competition between the agents.

3.6 Conclusions

We introduce a new type of bandit model where an agent is confronted with two-dimensional uncertainty. The agent does not know whether effort or ability is required to succeed. Moreover, her own ability is unknown and exerting effort is costly. In each period after deciding whether to exert effort or not, the agent observes a success or a failure and updates her beliefs about the task and her ability accordingly. Importantly, the agent gains information with as well as without exerting effort.

We show that depending on cost and beliefs some agents stick with their initial effort decision forever, thereby potentially foregoing a benefit. By contrast, other types of agents experiment with their effort level to gather information that is valuable for future periods. Amongst the latter, agents with low-ability beliefs see an advantage in starting by exerting effort while agents who believe their ability to be high start by not exerting effort. However, when facing a failure both types adjust their effort choice. For the case of uncertainty of the outcome in a infinite-time horizon, the agent may start and stop exerting effort repeatedly.

Our theoretical analysis gives similar results as the observations in Dweck (2006), where agents with growth mindset increase their effort level, while agents with fixed mindset stop exerting effort when observing a failure. In contrast to this literature, in our setup the ability level is rigid. The analysis of our model with ability as an increasing function of effort is an interesting task for future research.

3.7 Appendix

Proof of Lemma 3.1. After FY in $t = 1$, the expected payoff of all future periods is 0.

After SN in $t = 1$, the expected payoff in all future periods is 1 and amounts to $U^2_{SNT}(N) = \sum_{t=0}^{T-2} \delta^t$.

After SY in $t=1$, the expected payoff from exerting effort in period $t=2$ and all future periods is

$$U^2_{SYT}(Y)=\sum_{t=0}^{T-2}\delta^t\cdot(1-c).$$

If the agent does not exert effort in $t=2$ the uncertainty is resolved and the agent's payoff is

$$U^2_{SYT}(N)=Pr\left(E|SY\right)\cdot\sum_{t=1}^{T-2}\delta^t\cdot(1-c)+\left(1-Pr\left(E|SY\right)\right)\cdot\sum_{t=0}^{T-2}\delta^t.$$

Recall that $c_{SY}=Pr\left(E|SY\right)$. It pays off for the agent to exert effort in $t=2$ after observing SY in $t=1$ if and only if

$$U^2_{SYT}(Y)\geq U^2_{SYT}(N)\Leftrightarrow c\leq\frac{c_{SY}(1-\delta)}{1-(1-c_{SY})\delta^{T-1}-\delta c_{SY}}\equiv c_{SYT}.$$

After FN in $t=1$, the agent can resolve the uncertainty by exerting effort in period $t=2$. The expected payoff is

$$U^2_{FNT}(Y)=Pr\left(E|FN\right)\cdot\sum_{t=0}^{T-2}\delta^t\cdot(1-c)-\left(1-Pr\left(E|FN\right)\right)\cdot c.$$

Recall that $c_{FN}=Pr\left(E|FN\right)$. Since the agent receives a payoff of zero for sure without effort, $U^2_{FNT}(N)=0$, it pays off for her after observing FN to exert effort in $t=2$ if and only if

$$U^2_{FNT}(N)\leq U^2_{FNT}(Y)\Leftrightarrow c\leq c_{FN}\cdot\frac{1-\delta^{T-1}}{1-\delta+c_{FN}\left(\delta-\delta^{T-1}\right)}\equiv c_{FNT}.$$

For the first period we distinguish four cases. For low cost, $c\leq\min\{c_{SYT},c_{FNT}\}$, the agent will choose effort after both SY and FN in $t=2$. In $t=1$ we have

$$\begin{aligned}U^1_T(Y)=&q\left[(1-c)+\delta\cdot U^2_{SYT}(Y)\right]+(1-q)\alpha\left[(1-c)+\delta\cdot U^2_{SYT}(Y)\right]+(1-q)(1-\alpha)\left[-c+\delta\cdot 0\right],\\U^1_T(N)=&q\left[0+\delta\cdot U^2_{FNT}(Y)\right]+(1-q)\alpha\left[1+\delta\cdot U^2_{SNT}(N)\right]+(1-q)(1-\alpha)\left[0+\delta\cdot U^2_{FNT}(Y)\right],\end{aligned}$$

and

$$U^1_T(Y)\geq U^1_T(N)\Leftrightarrow c\leq\frac{q}{1+\delta(1-q)\left(\frac{\alpha(1-\delta^{T-1})}{1-\delta}-(1-\alpha)\right)}\equiv c_{1T}.$$

Similarly, for $c_{FNT}\leq c\leq c_{SYT}$ the agent will choose effort after SY but not after

FN in $t = 2$. In $t = 1$ we have

$$U_T^1(Y) \geq U_T^1(N) \Leftrightarrow c \leq \frac{q\frac{1-\delta^T}{1-\delta}}{1 + (q(1-\alpha)+\alpha)\delta\frac{1-\delta^{T-1}}{1-\delta}} \equiv c_{2T}.$$

For $c_{SYT} \leq c \leq c_{FNT}$ the agent will choose effort after FN but not SY in $t = 2$. In $t = 1$ we have $U_T^1(Y) \geq U_T^1(N) \Leftrightarrow c \leq \frac{q(1-\delta)}{1-\delta(q+(1-q)(1-\alpha))}$. This constraint is not binding.

For $c \geq \max\{c_{SY}, c_{FN}\}$ the agent will not choose effort after neither SY nor FN in $t = 2$. In $t = 1$ we have $U_T^1(Y) \geq U_T^1(N) \Leftrightarrow c \leq q$. This constraint is not binding.

Finally, by taking the limit $T \longrightarrow \infty$ in the above formulas, our results extend to $T = \infty$. □

Proof of Lemma 3.2. The agent's expected payoff from exerting effort in period $t = 1$ is given by

$$\begin{aligned} U(Y) = -c + q &\left[(1-\varepsilon)\left(1+\delta\begin{cases}1-\varepsilon-c, & \text{if } c \leq c_{SY\varepsilon}\\ 0, & \text{otherwise}\end{cases}\right) + \varepsilon\delta\begin{cases}1-\varepsilon-c, & \text{if } c \leq c_{FY\varepsilon}\\ 0, & \text{otherwise}\end{cases}\right] \\ &+ (1-q)\alpha\left[(1-\varepsilon)\left(1+\delta\begin{cases}1-\varepsilon-c, & \text{if } c \leq c_{SY\varepsilon}\\ 1-\varepsilon, & \text{otherwise}\end{cases}\right) + \varepsilon\delta\begin{cases}1-\varepsilon-c, & \text{if } c \leq c_{FY\varepsilon}\\ 1-\varepsilon, & \text{otherwise}\end{cases}\right] \\ &+ (1-q)(1-\alpha)\delta\begin{cases}-c, & \text{if } c \leq c_{FY\varepsilon}\\ 0, & \text{otherwise}\end{cases}. \end{aligned}$$

Meanwhile, her expected payoff from not exerting effort in period $t = 1$ is given by

$$\begin{aligned} U(N) = q &\left[(1-\varepsilon)\delta\begin{cases}1-\varepsilon-c, & \text{if } c \leq c_{FN\varepsilon}\\ 0, & \text{otherwise}\end{cases}\right] \\ &+ (1-q)\alpha\left[(1-\varepsilon)(1+\delta(1-\varepsilon)) + \varepsilon\delta\begin{cases}1-\varepsilon-c, & \text{if } c \leq c_{FN\varepsilon}\\ 1-\varepsilon, & \text{otherwise}\end{cases}\right] \\ &+ (1-q)(1-\alpha)\delta\begin{cases}-c, & \text{if } c \leq c_{FN\varepsilon}\\ 0, & \text{otherwise}\end{cases}. \end{aligned}$$

When comparing $U(Y)$ with $U(N)$, there are five different cases to consider depending on the agent's cost and her beliefs about the task and her ability.

For low cost, $c \leq \min\{c_{SY\varepsilon}, c_{FN\varepsilon}, c_{FY\varepsilon}\}$, the agent exerts effort in period $t = 1$ if her cost is below $c_{1\varepsilon} \equiv \frac{q(1-\varepsilon)}{1+(1-q)\alpha\delta(1-\varepsilon)}$.

For medium cost, $c_{FY\varepsilon} \leq c \leq \min\{c_{FN\varepsilon}, c_{SY\varepsilon}\}$, the agent exerts effort in period

$t = 1$ if her cost is below $c_{2\varepsilon} \equiv \frac{q(1-\varepsilon)(1-\delta\varepsilon)}{1+\delta[(-1+2\alpha+q(-1+2\alpha)(-1+\varepsilon)-2\alpha\varepsilon)]}$.

For highish cost and low ability $c_{FN\varepsilon} \leq c \leq c_{SY\varepsilon}$, the agent exerts effort in period $t = 1$ if her cost is below $c_{3\varepsilon} \equiv \frac{q(1+\delta(1-\varepsilon))(1-\varepsilon)}{1+(q(1-\alpha)+\alpha)\delta(1-\varepsilon)}$.

For highish cost and high ability, $c_{SY\varepsilon} \leq c \leq c_{FN\varepsilon}$, the agent does not exert effort in period $t = 1$.

For high cost, $c \geq \max\{c_{SY\varepsilon}, c_{FN\varepsilon}\}$, the agent does not exert effort in period $t = 1$. □

Proof of Lemma 3.3. To simplify notation for histories that only include failures we define the updated beliefs about the states E and AH as follows

$$p_E \equiv \frac{q\varepsilon^{n_{FY}}}{q\varepsilon^{n_{FY}} + (1-q)\alpha \cdot \varepsilon^{n_{FN}+n_{FY}} + (1-q)(1-\alpha)},$$
$$p_{AH} \equiv \frac{(1-q)\alpha \cdot \varepsilon^{n_{FN}+n_{FY}}}{q\varepsilon^{n_{FY}} + (1-q)\alpha \cdot \varepsilon^{n_{FN}+n_{FY}} + (1-q)(1-\alpha)}.$$

The probability of state AL is $p_{AL} \equiv 1 - p_E - p_{AH}$. The probabilities p_E and p_{AL} increase in n_{FN}, while p_{AH} decreases. The limits of observing repeatedly the outcome FN, are given by

$$\lim_{n_{FN}\to\infty} p_E = \frac{q \cdot \varepsilon^{n_{FY}}}{q \cdot \varepsilon^{n_{FY}} + \left(1-q\right)(1-\alpha)},$$
$$\lim_{n_{FN}\to\infty} p_{AH} = 0,$$
$$\lim_{n_{FN}\to\infty} p_{AL} = \frac{(1-q)(1-\alpha)}{q \cdot \varepsilon^{n_{FY}} + \left(1-q\right)(1-\alpha)}.$$

Furthermore, p_{AL} increases in n_{FY}, while p_E and p_{AH} decrease. The limits of observing repeatedly the outcome FY, are given by

$$\lim_{n_{FY}\to\infty} p_E = 0,$$
$$\lim_{n_{FY}\to\infty} p_{AH} = 0,$$
$$\lim_{n_{FY}\to\infty} p_{AL} = 1.$$

We next calculate the expected payoffs in period t of exerting effort in t ("today") but not in $t+1$ ("tomorrow") and vice versa (i.e. (Y,N) and (N,Y)). Since for the updating process only the number but not the order of the outcomes plays a role, the continuation values after observing FY and FN can be denoted as $V_{FF} \equiv V_{FNFY} = V_{FYFN}$. The expected utility of effort choices (Y,N) and outcome

FY is given by

$$\begin{aligned} U_{h_t}(Y,N) =& -c + p_E\left\{(1-\varepsilon)\left[1+\delta\cdot\frac{1-\varepsilon-c}{1-\delta}\right]+\varepsilon\cdot 0\right\} \\ &+ p_{AH}\left\{(1-\varepsilon)\left[1+\delta\cdot\frac{1-\varepsilon-c}{1-\delta}\right]+\varepsilon\cdot\left[0+\delta\cdot(1-\varepsilon)\left(1+\delta\cdot\frac{1-\varepsilon}{1-\delta}\right)\right]\right\} \\ &+ (p_E\cdot\varepsilon + p_{AH}\cdot\varepsilon^2 + p_{AL})\cdot V_{FF}. \end{aligned}$$

The expected utility of effort choices (N,Y) and outcome FN is given by

$$\begin{aligned} U_{h_t}(N,Y) =& p_E\left\{0+\delta\left[-c+(1-\varepsilon)\left(1+\delta\cdot\frac{1-\varepsilon-c}{1-\delta}\right)\right]\right\} \\ &+ p_{AH}\left\{(1-\varepsilon)\left[1+\delta\cdot\frac{1-\varepsilon}{1-\delta}\right]+\varepsilon\cdot\left(0+\delta\left[-c+(1-\varepsilon)\left(1+\delta\cdot\frac{1-\varepsilon-c}{1-\delta}\right)\right]\right)\right\} \\ &+ (p_E\cdot\varepsilon + p_{AH}\cdot\varepsilon^2 + p_{AL})\cdot V_{FF}. \end{aligned}$$

Define

$$\begin{aligned} \beta(n_{FY},n_{FN}) \equiv& U_{h_t}(Y,N) - U_{h_t}(N,Y) \\ =& p_E(1-\varepsilon-c)(1-\delta\varepsilon) + p_{AH}(-c)\frac{(1-\delta\varepsilon)^2}{1-\delta} + p_{AL}(-c)(1-\delta) \\ =& \frac{-c(1-q)(1-\alpha)(1-\delta)^2 + \varepsilon^{n_{FY}}(1-\delta\varepsilon)(q(1-\delta)(1-c-\varepsilon) - c(1-q)\alpha\varepsilon^{n_{FN}}(1-\delta\varepsilon))}{(1-\delta)(1-\alpha+(1-q)\alpha\varepsilon^{n_{FN}+n_{FY}} - q(1-\alpha-\varepsilon^{n_{FY}}))}. \end{aligned}$$

In the second line, the first term of β is positive while the second and third term are negative. The function β is monotonically increasing in n_{FN} and monotonically decreasing in n_{FY}. The intuition is that p_{AH} is decreasing in n_{FN} while the ratio $\frac{p_E}{p_{AL}}$ remains constant, and p_{AL} is increasing in n_{FY} while the ratio $\frac{p_E}{p_{AH}}$ remains constant. However, for large enough n_{FY} the function β stays negative for all n_{FN}.

The agent prefers to exert effort today rather than tomorrow, $U_{h_t}(Y,N) \geq U_{h_t}(N,Y)$ if and only if $\beta(n_{FY},n_{FN}) \geq 0$. Moreover, $0 \leq \beta(n_{FY},n_{FN})$ implies $0 \leq \beta(n_{FY},n_{FN}+m)$ for $m \geq 0$, provided that no success occurs; if the agent prefers effort today rather than tomorrow in period t, she will also prefer effort today rather than tomorrow in any later period. Monotonicity then implies that for a given n_{FY}, an agent that prefers effort today rather than tomorrow also prefers effort today rather than in any future period.

Consider the limit $n_{FN} \to \infty$. We define $\bar{n}_{FY}$ as the maximum number of FY outcomes an agent is willing to experience before stopping effort forever

$$\bar{n}_{FY} \equiv \max\left\{n_{FY} | \lim_{n_{FN}\to\infty} \beta(n_{FY},n_{FN}) \geq 0\right\}.$$

For a given $\bar{n}_{FY}$ we define $\underline{n}_{FN}$ as the minimum number of periods with no effort before the agent attempts effort for the $\bar{n}_{FY}$th time (provided she experiences no success),

$$\underline{n}_{FN} \equiv \min\left\{n_{FN} | \beta(\bar{n}_{FY}, n_{FN}) \geq 0\right\}.$$

Recall that β is increasing in n_{FN} and decreasing in n_{FY}. For $n_{FY} \leq \bar{n}_{FY}$ we have $\beta(n_{FY}, \underline{n}_{FN}) \geq 0$; after not exerting effort for $\underline{n}_{FN}$ periods, the agent exerts effort until $\bar{n}_{FY}$, the maximum number of FY, is reached. On the other hand, for $n_{FY} < \bar{n}_{FY}$ it is possible that $\beta(n_{FY}, n_{FN}) \geq 0$ also for $n_{FN} < \underline{n}_{FN}$; before reaching the maximum $\bar{n}_{FY}$, the agent prefers to not further delay effort before $\underline{n}_{FN}$ periods of inactivity have elapsed.

We can conclude that the optimal behavior of the agent is as follows: The agent exerts effort whenever $\beta(n_{FY}, n_{FN}) \geq 0$. The agent starts exerting effort as soon as $\beta(0, n_{FN}) \geq 0$. After experiencing FY, she stops exerting effort until again $\beta(1, n_{FN}) \geq 0$. This process continues until $\bar{n}_{FY}$ times the outcome FY is observed and the agent stops exerting effort forever, $\beta(\bar{n}_{FY}, n_{FN}) < 0$ for any n_{FN}.

The cost threshold $c_{\infty FN}(\bar{n}_{FY})$ corresponds to an agent that chooses $e = N$ after observing an infinite number of FN outcomes

$$
\begin{aligned}
0 &> -c + \frac{q \cdot \varepsilon^{n_{FY}}}{q \cdot \varepsilon^{n_{FY}} + \left(1-q\right)(1-\alpha)}(1-\varepsilon)(1+\delta\frac{1-\varepsilon-c}{1-\delta}) \\
\Leftrightarrow c &> \frac{q(1-\varepsilon)(1-\delta\varepsilon)\varepsilon^{n_{FY}}}{(1-q)(1-\alpha)(1-\delta) + q\varepsilon^{n_{FY}}(1-\delta\varepsilon)} \equiv c_{\infty FN}(n_{FY}).
\end{aligned}
$$

□

Bibliography

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

Alonso, R., Dessein, W. and Matouschek, N. (2008). When does coordination require centralization?, *American Economic Review* **98**(1): 145–179.

Angeletos, G.-M. and Pavan, A. (2007). Efficient use of information and social value of information, *Econometrica* **75**(4): 1103–1142.

Argenziano, R., Severinov, S. and Squintani, F. (2013). Strategic information acquisition and transmission. Working paper.

Austen-Smith, D. and Banks, J. S. (1996). Information aggregation, rationality, and the condorcet jury theorem, *American Political Science Review* **90**(1): 34–45.

Austen-Smith, D. and Feddersen, T. J. (2006). Deliberation, preference uncertainty and voting rules, *American Political Science Review* **100**(2): 209–218.

Battaglini, M. (2002). Multiple referrals and multidimensional cheap talk, *Econometrica* **70**(4): 1379–1401.

Bergemann, D. and Valimaki, J. (2006). Bandit problems, *Cowles Foundation Discussion Paper No. 1551* .

Berry, D. A. and Fristedt, B. (1985). *Bandit problems: sequential allocation of experiments*, Springer.

Bhattacharya, S. (2013). Preference monotonicity and information aggregation in elections, *Econometrica* **81**(3): 1229–1247.

Blume, A., Board, O. J. and Kawamura, K. (2007). Noisy talk, *Theoretical Economics* **2**(4): 395–440.

Bolton, P. and Harris, C. (1999). Strategic experimentation, *Econometrica* **67**(2): 349–374.

Bonatti, A. and Hörner, J. (2015). Career concerns with exponential learning. Working paper.

Chakraborty, A. and Harbaugh, R. (2007). Comparative cheap talk, *Journal of Economic Theory* **132**(1): 70–94.

Chakraborty, A. and Harbaugh, R. (2010). Persuasion by cheap talk, *American Economic Review* **100**(5): 2361–2382.

Coughlan, P. J. (2000). In defense of unanimous jury verdicts: Mistrials, communication, and strategic voting, *American Political Science Review* **94**(2): 375–393.

Crawford, V. and Sobel, J. (1982). Strategic information transmission, *Econometrica* **50**(6): 1431–1451.

Cyert, R. M., Dill, W. R. and March, J. G. (1958). The role of expectations in business decision making, *Administrative Science Quarterly* **3**(3): 307–340.

Cyert, R. M. and March, J. G. (1963). A behavioral theory of the firm, *Englewood Cliffs, NJ* **2**.

Deimen, I., Ketelaar, F. and Le Quement, M. T. (2015). Consistency and communication in committees, *Journal of Economic Theory* **160**: 24 –35.

Deimen, I. and Szalay, D. (2015). Information, authority, and smooth communication in organizations, *CEPR Discussion Paper10969* .

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

Dewatripont, M., Jewitt, I. and Tirole, J. (1999). The economics of career concerns, part i: Comparing information structures, *Review of Economic Studies* **66**(1): 183–198.

Dickson, D., Hafer, C. and Landa, D. (2008). Cognition and strategy: A deliberation experiment, *Journal of Politics* **70**(4): 974–989.

Dimitrakas, V. and Sarafidis, Y. (2005). Advice from an expert with unknown motives, *SSRN working paper 797844* .

Dweck, C. (2000). *Self-theories: Their role in motivation, personality, and development*, Psychology Press.

Dweck, C. (2006). *Mindset: The new psychology of success*, Random House.

Eső, P. and Szalay, D. (2015). Incomplete language as an incentive device. Working paper.

Fang, K.-T., Kotz, S. and Ng, K. W. (1990). *Symmetric multivariate and related distributions*, Chapman and Hall.

Feddersen, T. and Pesendorfer, W. (1998). Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting, *American Political Science Review* **92**(1): 23–35.

Fryer, R. G. J. and Harms, P. (2015). Two-armed restless bandits with imperfect information: Stochastic control and indexability. Working paper.

Gerardi, D., McLean, R. and Postlewaite, A. (2009). Aggregation of expert opinions, *Games and Economic Behavior* **65**(2): 339–371.

Gerardi, D. and Yariv, L. (2007). Deliberative voting, *Journal of Economic Theory* **134**(1): 317–338.

Goeree, J. K. and Yariv, L. (2011). An experimental study of collective deliberation, *Econometrica* **79**(3): 893–921.

Goltsman, M., Hörner, J., Pavlov, G. and Squintani, F. (2009). Mediation, arbitration and negotiation, *Journal of Economic Theory* **144**(4): 1397–1420.

Gordon, S. (2010). On infinite cheap talk equilibria. Working paper.

Gordon, S. and Nöldeke, G. (2013). Figures of speech and informative communication. Working paper.

Heidhues, P., Koszegi, B. and Strack, P. (2015). Unrealistic expectations and misguided learning. Working paper.

Holmström, B. (1999). Managerial incentive problems: A dynamic perspective, *Review of Economic Studies* **66**(1): 169–182.

Hummel, P. (2012). Deliberation in large juries with diverse preferences, *Public Choice* **150**(3): 595–608.

Ivanov, M. (2010). Informational control and organizational design, *Journal of Economic Theory* **145**(2): 721–751.

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

Kawamura, K. (2015). Confidence and competence in communication, *Theory and Decision* **78**(2): 233–259.

Keller, G. and Rady, S. (2010). Strategic experimentation with poisson bandits, *Theoretical Economics* **5**(2): 275–311.

Keller, G., Rady, S. and Cripps, M. (2005). Strategic experimentation with exponential bandits, *Econometrica* **73**(1): 39–68.

Le Quement, M. T. (2013). Communication compatible voting rules, *Theory and Decision* **74**(4): 479–507.

Le Quement, M. T. and Yokeeswaran, V. (2015). Subgroup deliberation and voting, *Social Choice and Welfare* **45**(1): 155–186.

Levy, G. and Razin, R. (2007). On the limits of communication in multidimensional cheap talk: a comment, *Econometrica* **75**(3): 885–893.

Li, M. and Madarász, K. (2008). When mandatory disclosure hurts: Expert advice and conflicting interests, *Journal of Economic Theory* **139**(1): 47–74.

Lizzeri, A. and Siniscalchi, M. (2008). Parental guidance and supervised learning, *The Quarterly Journal of Economics* **123**(3): 1161–1195.

Manso, G. (2011). Motivating innovation, *Journal of Finance* **66**(5): 1823–1860.

March, J. G. and Simon, H. A. (1958). *Organizations*, Wiley.

Meirowitz, A. (2007). In defense of exclusionary deliberation: Communication and voting with private beliefs and values, *Journal of Theoretical Politics* **19**(3): 329–360.

Meyer, M., de Barreda, I. M. and Nafziger, J. (2013). Robustness of full revelation in multisender cheap talk. Working paper.

Morgan, J. and Stocken, P. C. (2003). An analysis of stock recommendations, *RAND Journal of Economics* **34**(1): 183–203.

Moscarini, G. (2007). Competence implies credibility, *American Economic Review* **97**(1): 37–63.

Nöldeke, G. and Tröger, T. (2006). A characterization of the distributions that imply existence of linear equilibria in the kyle-model, *Annals of Finance* **2**(1): 73–85.

Ottaviani, M. and Sørensen, P. N. (2006a). Professional advice, *Journal of Economic Theory* **126**(1): 120–142.

Ottaviani, M. and Sørensen, P. N. (2006b). The strategy of professional forecasting, *Journal of Financial Economics* **81**(2): 441–466.

Pavan, A. and Vives, X. (2015). Information, coordination, and market frictions: An introduction, *Journal of Economic Theory* **158,**(Part B): 407 – 426.

Piketty, T. (1995). Social mobility and redistributive politics, *Quarterly Journal of Economics* **110**(3): 551–584.

Rothschild, M. (1974). A two-armed bandit theory of market pricing, *Journal of Economic Theory* **9**(2): 185–202.

Stein, J. (1989). Cheap talk and the fed: A theory of imprecise policy announcements, *American Economic Review* **79**(1): 32–42.

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

Van Weelden, R. (2008). Deliberation rules and voting, *Quarterly Journal of Political Science* **3**(1): 83–88.

Vives, X. (1999). *Oligopoly Pricing: Old Ideas and New Tools*, MIT Press.

Vives, X. (2008). *Information and Learning in Markets*, Princeton University Press.

Weitzman, M. L. (1979). Optimal search for the best alternative, *Econometrica* **47**(3): 641–654.

Whittle, P. (1988). Restless bandits: Activity allocation in a changing world, *Journal of applied probability* **25**: 287–298.

Wolinsky, A. (2002). Eliciting information from multiple experts, *Games and Economic Behavior* **41**(1): 141–160.