



HAL
open science

Certifiable information : revelation and persuasion

Mohamed Mehdi Ayouni Dit Houimel

► **To cite this version:**

Mohamed Mehdi Ayouni Dit Houimel. Certifiable information : revelation and persuasion. Economics and Finance. Université de Cergy Pontoise, 2016. English. NNT : 2016CERG0875 . tel-01644804

HAL Id: tel-01644804

<https://theses.hal.science/tel-01644804>

Submitted on 22 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIVERSITÉ DE CERGY-PONTOISE

THÈSE

pour le **Doctorat de Sciences Économiques**

Présentée et soutenue publiquement par

Mehdi Ayouni

le 14 décembre 2016

Information certifiable : révélation et persuasion

Directeurs de thèse :

M. Frédéric Koessler

Directeur de recherche au CNRS, Ecole d'Economie de Paris

M. Régis Renault

Professeur à l'Université de Cergy-Pontoise

Jury :

M. Gabriel Desgranges,	Professeur à l'Université de Cergy-Pontoise (<i>Président</i>)
M. Sidartha Gordon,	Professeur à l'Université Paris-Dauphine
M. Frédéric Koessler,	Directeur de recherche au CNRS, Ecole d'Economie de Paris
Mme Lucie Ménager,	Professeur à l'Université Paris II
M. Régis Renault,	Professeur à l'Université de Cergy-Pontoise
Mme Vasiliki Skreta,	Professeur à University College London

Remerciements

En tout premier lieu, je tiens à remercier mes directeurs de thèse Frédéric Koessler et Régis Renault pour leur encadrement sans faille et leur disponibilité pendant la réalisation de ce travail. Leurs conseils et nombreuses suggestions m'ont été d'une valeur inestimable. Je les remercie pour tout ce qu'ils m'ont appris tant en termes de connaissances que sur le plan méthodologique et scientifique.

Je remercie Gabriel Desgranges, Sidartha Gordon, Lucie Ménager et Vasiliki Skreta d'avoir accepté de faire partie de mon jury de thèse. Leurs commentaires pertinents et remarques intéressantes me seront d'une grande utilité pour la publication de mes travaux et la poursuite de mes recherches.

En rejoignant le Thema, j'ai pu d'intégrer une équipe de recherche aux intérêts divers et variés ce qui a permis d'élargir mon horizon. Je remercie tout particulièrement les membres du groupe de théorie du Thema pour les discussions intéressantes à l'occasion des multiples présentations que j'ai eu l'occasion de réaliser tout au long de mon doctorat. J'ai également eu la chance de faire partie d'un groupe de doctorants avec qui j'ai partagé une expérience très enrichissante tant sur le plan humain que sur le plan professionnel. Je tiens donc à les remercier pour tous les moments agréables que nous avons vécus ensemble au cours de ces quatre années.

Je remercie également l'équipe administrative du Thema et de l'Ecole doctorale EM2P pour leur disponibilité et le Labex MME-DII pour le soutien financier.

J'exprime mon immense gratitude envers tous mes amis qui m'ont toujours soutenu et encouragé. Enfin, je remercie tous les membres de ma famille pour leurs encouragements, leur confiance sans faille et leur indéfectible soutien.

À mes parents,
À mes petites sœurs,

Introduction

A decision maker is required to be well informed in order to choose the best action or policy. If some relevant information is missing, bad decisions can be made. For that reason, she might benefit from acquiring additional information prior to selecting an action. In some situations, this information is publicly available and the decision maker can learn it by reading a book or scientific publication. However, the expected benefit might not justify the required time and effort to understand it. In other situations, the missing information is privately held and it is either impossible or prohibitively costly to acquire it directly. For example, firms have private information about their activities and it is difficult for regulators to access it if they do not willingly reveal it. In both types of situations, the decision maker may choose communication as a way to learn more information instead of direct acquisition.

Depending on the problem at hand, the informed party can send one of two types of messages (or a combination of both): *cheap talk* and *hard evidence*. Cheap talk messages are statements that can be sent regardless of the true state of the world. As such, their meaning results from the interaction itself and the preferences of the parties involved. By contrast, hard evidence carries an intrinsic meaning because every piece of evidence can be produced in some states and not in others. In that sense, it *certifies* that the true state is in a certain subset of states.

In the present thesis, I study frameworks where at least some information can be certified through hard evidence. Many economic settings satisfy this condition. For instance, taxpayers can be audited and asked to provide official documents regarding their income, employer in-

interview job applicants and may test their abilities before hiring them and food producers are required to prove their products are harmless for consumers.

It is natural to assume that some statements cannot be certified. Usually it is possible for an agent to prove that he masters a skill by performing a task. However, it is not possible to prove the opposite. For instance, one can perform card tricks to show his sleight-of-hand abilities. But in a card game, participants cannot be sure that players are not capable of cheating using the same techniques because there is no test to prove it.

More generally, one can hide evidence if it is detrimental to reveal it. In some settings, agents are required by law to disclose all information but some information is inevitably concealed in the absence of proper incentives to reveal it. In fact, enforcing disclosure in this manner requires knowledge about information storage and legal authority to seize it. Although this might be the case in some court mandated investigations, it is generally not feasible in ordinary situations and economic interactions. For that reason, the decision maker needs to focus on providing incentives for revelation and cannot expect the informed party to abide by such legal requirements.

In the present thesis, I consider two frameworks satisfying the aforementioned conditions. The first is a setting where a decision maker has to evaluate a claim or respond to a request made by an individual or an institution. Information in this context is private and difficult to access if the party sending the request is not willing to reveal it.

In some cases, information is certified by a trusted third party. In other cases, certification takes place during the interaction: the decision maker needs to *verify* the presented evidence, possibly with the help of the informed party. I focus on situations where at least some information falls in the latter category. In real life applications, the need for verification induces a limitation on the amount of certifiable information that can be exchanged during the interaction between the decision maker and the informed agent. For instance, verification can be costly or time consuming. If in addition, the decision maker is overloaded with activities or has to

quickly select an action, she would not be able to verify all information the agent is willing to present.

These constraints determine the amount of information that can be verified before the decision is made. For example, a financial auditor cannot examine every single transaction of a firm in the limited time he can allocate to this task. Assuming there are N transactions and the auditor verifies n transactions before making a decision, one can define the set of *certifiable events* as the set of all possible choices of n among N transactions. In general, a certifiable event is a statement that can be verified before the final decision given the constraints of the problem at hand. A useful benchmark is obtained by relaxing these constraints so that all certifiable events can be verified at once. In the example above, this is equivalent to having as much time as necessary to examine all N transactions. The impact of these constraints can be studied by comparing implementable outcome functions in the original setting and this benchmark. In the first two chapters, I consider situations where implementation is effectively restricted by such constraints and identify ways in which the decision maker can undo this restriction or at least enlarge the set of implementable outcome functions in order to implement better outcomes. An outcome function is a mapping from states to actions and it is implemented if the action it specifies for every state is chosen when the state is realized. Once a set of certifiable events is defined, we can interchangeably say that the agent certifies an event or that the decision maker verifies it because in both cases the same constraints apply.

In the first chapter, I study the basic framework where the informed agent's preferences over the decision maker's actions are independent of the state. It offers a good description of settings where the decision maker chooses a reward, a budget allocation or a wage for the agent who wants the action to be as high as possible. I compare two types of communication mechanisms, namely *unilateral* and *bilateral* communication, under the condition that the same amount of information can be verified in both cases. In unilateral communication, only the agent sends a message to the decision maker. In bilateral communication, the decision maker can also actively

communicate with the agent. The additional exchange of messages allows the decision maker to influence the choice of the event that needs to be certified. I argue that the decision maker can benefit from bilateral communication even though she does not verify more information than in unilateral communication.

In the canonical bilateral communication mechanism, the agent announces the state of the world, the decision maker asks him to certify a certain event among those that can be certified in the announced state. Then she bases her decision on his ability to do so. In addition, the agent truthfully reveals the state in the first stage. This mechanism is canonical in the sense that it achieves any result that can be achieved using any other bilateral communication mechanism.

In order to show that decision maker can benefit from active communication, I first show that bilateral and unilateral achieve the same results if and only if there are no effective limitations on information verification. Therefore, as long as time and cost constraints apply as explained above, bilateral communication allows the decision maker to implement more outcome functions than unilateral communication. However, these additional outcome functions do not necessarily constitute an improvement.

The main result of this chapter essentially states that if information certification is limited so as to prevent the decision maker from achieving her first-best in unilateral communication then she strictly benefits from bilateral communication.

A first-best outcome function is an outcome function that maps every state to an action that maximizes the decision maker's payoff in that state. Limited certification prevents the decision maker from achieving her first-best in unilateral communication if three conditions are satisfied. First, the first-best needs to be well-defined. In other words, in every state there is at least one action that gives the decision maker her maximal payoff in that state. Second, no first-best outcome function is implementable in unilateral communication. Otherwise, it would not be possible to make any improvements. Finally, there needs to be at least one first-best outcome function that is implementable in the benchmark where certification is unlimited: if

unlimited time and resources can be allocated to information verification, the decision maker can achieve her maximal payoff in every state.

Under these conditions, bilateral communication allows the decision maker to improve the outcome in comparison with unilateral communication thereby reducing the impact of certification limitation while verifying the same amount of information.

In the second chapter, which results from a joint work with my advisor Frédéric Koessler (CNRS, Paris School of Economics), we consider a similar framework but we do not require agent's preferences to be state independent. We focus on the difference between limited and unlimited certification and we study the way *ambiguous* mechanisms can exploit agent's *ambiguity aversion* to enlarge the set of implementable outcome functions.

In the aforementioned bilateral communication mechanism, after the agent announces his type the decision maker selects an event according to a certain probability distribution and asks him to certify it. If the decision maker uses a set of probability distributions instead of a single distribution, the communication mechanism is called ambiguous. The use of such a mechanism generates ambiguous beliefs for the agent in a similar way to the use of an Ellsberg urn where the proportion of balls of each color is unknown.

We focus on settings where the agent is ambiguity averse in the sense of maxmin expected utility. It describes an extreme case of ambiguity aversion where the agent anticipates the worst case scenario. If the decision maker uses the whole set of possible probability distributions in the ambiguous mechanism, the agent acts as if certification is unlimited. Under unlimited certification, an agent who announces a state but cannot prove some certifiable events of that state knows that the decision maker will discover that he misreported it. Under limited certification and an ambiguous mechanism that uses all probability distributions, the worst case scenario for the same agent is to be asked to certify an event that he cannot certify. Therefore, he anticipates that the decision maker will also discover that he misreported the state.

The main result of this chapter is that any allocation rule that is implementable with unlim-

ited certification is also implementable with limited certification and ambiguous mechanisms if the agent is ambiguity averse in the sense of maxmin expected utility. Ambiguous mechanisms completely eliminate the effect of limited certification in that case. We also show that the result holds if there are multiple agents.

In addition, we show that the converse is true if there is only one agent and a *worst outcome*. A worst outcome is an action that gives the agent his minimal payoff in all states. Such an outcome does not exist in general because agent's preferences may depend on the state. But if it exists, it can be used if certification is unlimited to implement all outcome functions that can be implemented using ambiguous mechanisms.

In the last chapter, I consider a different framework where the decision maker does not respond to a request but rather seeks advice by sequentially consulting informed agents before choosing an action that affects her payoff and the payoff of all informed agents.

In many economic settings, decision makers rely on the advice of informed agents. If they have the same objectives, revelation is guaranteed and there is no strategic interaction. If she consults an agent who is biased but his bias is known it may well be enough to extract all information. But if agents' preferences are unknown, I argue that sequential consultation of more than one agent can be beneficial to the decision maker even though she incurs a cost due to waiting time, delayed actions or information processing.

Unlike the first two chapters, the decision maker does not design a mechanism but simply plays a game where she consults informed agents until she decides to stop and choose an action. I consider two settings where the decision maker has to choose a scalar action and every agent has one of two types: his payoff is state independent and is either strictly increasing or strictly decreasing in the action.

In the first setting, states and actions are binary. They can take values 0 and 1. The decision maker wants to choose the action that matches the state while every agent wants her to choose his favorite action regardless of the state. As in the first two chapters, information is certifiable

so that the agent can credibly reveal the state if he is consulted or remain silent if he chooses to do so.

The decision maker wants to reduce her uncertainty about the state before choosing an action and she benefits from sequential consultation if prior uncertainty is high enough. As she consults informed agents, she either learns the state because an agent reveals it or she acquires information from silent agents if types are not equiprobable. In that case, a silent agent is more likely to be in the state where the majority remains silent so that the decision maker updates her belief by increasing the probability of that state.

As a consequence, a bad decision can be made if only agents of the minority are consulted in the state where they prefer to withhold information. It is therefore possible for agents of the minority to persuade the decision maker to choose their favorite action when she should not by remaining silent, as long as she does not encounter an agent of the majority before her uncertainty about the state becomes low enough.

In the second setting, I consider a richer environment where the state and action can take any value between 0 and 1. The decision maker wants to match the state and faces a quadratic loss function if her action is different from the state. A consulted agent can certify any interval containing the true state so that he chooses the precision of the information he reveals.

In the case where sequential consultation is not allowed, there exists a unique equilibrium and it turns out that the message space richness does not affect the result. This equilibrium has the same outcome as the one that obtains when the agents can only reveal the state or remain silent.

But if the decision maker can consult more than one agent, it is possible to have an equilibrium where she consults only one agent and uses her ability to keep seeking advice to make him reveal more precise information. The agent does so only to prevent her from consulting other agents. He sends a message that is precise enough to make her stop and choose an action thereby guaranteeing that he is the only one to influence her decision. I show the existence of

such an equilibrium for a low enough consultation cost and equiprobable types. The analysis in general is tricky due to the additional information that the decision maker can acquire when types are not equiprobable as in the first setting.

Contents

1	Bilateral communication and hard evidence	21
1.1	Introduction	21
1.2	The model	25
1.2.1	The environment	25
1.2.2	The canonical form	27
1.3	Implementable outcome functions	29
1.4	The value of bilateral communication	36
1.5	Extensions	42
1.5.1	Tightness of Proposition 1.5	42
1.5.2	Optimal solutions: Randomization over actions and Commitment	45
1.6	Appendix	48
2	Hard evidence and ambiguity aversion	59
2.1	Introduction	59
2.2	Model	61
2.3	Implementation without Ambiguity	63
2.4	Implementation with Ambiguous Mechanisms	66
2.5	Discussion	72

3	Seeking advice	77
3.1	Introduction	77
3.2	Model	81
3.3	Equilibrium analysis	83
3.3.1	Agents prefer immediate decisions	84
3.3.2	Agents prefer delayed bad outcomes	85
3.4	Extensions	94
3.4.1	Commitment	94
3.4.2	Cost of waiting for information	95
3.4.3	Processing cost avoidance	100
3.5	Alternative model: information precision	101
3.5.1	Game without sequential consultation	102
3.5.2	Game with sequential consultation	106
3.6	Appendix	113
	Bibliography	123

Chapter 1

Bilateral communication and hard evidence

1.1 Introduction

In a persuasion problem, an agent wishes to influence a principal who has to implement an outcome. The agent privately knows the state of the world, also called his *type*, and has *hard evidence* about it. Any certified message that proves a non trivial statement is considered hard evidence. Formally, a piece of evidence is a message certifying a certain *event*. Namely, that the agent belongs to a certain subset of types. Not all events are necessarily certifiable and the set of certifiable events depends on the problem at hand.

The principal ignores the state of the world which is relevant to her decision, but she can interact with the agent before implementing an action. The standard setting to model such an interaction is the sender-receiver game: the agent (sender) presents information to the principal (receiver) by sending a message containing certifiable information before the principal chooses an outcome. We call this setting the *unilateral communication* framework. In contrast, the *bilateral communication* setting is one where both the agent and the principal are active in the

communication phase, exchanging messages sequentially.

As an illustration, consider the example of a hiring process. The agent is the applicant who knows his skills and abilities which define his type. The principal is the employer who does not observe that information so she interviews him before making a decision. If the employer can learn all information during the interview, she would not gain from being active in the communication phase. However, in some cases only some information can be certified: for example it might not be possible to test all the skills the applicant claims to master due to cost or time constraints. In such cases, bilateral communication might allow the employer to improve the outcome by choosing what the applicant has to certify based on his claims instead of letting him choose the information he presents as in unilateral communication.

In the unilateral communication framework, the principal has to choose an *implementation rule* that assigns an outcome to every possible message the agent can send. In the bilateral communication framework, the principal has to design the *communication mechanism* in addition to the implementation rule. The communication mechanism specifies the active player and the set of available messages at each node. The implementation rule, in this case, specifies the outcome possibly as a function of the history of exchanged messages.

Our goal is to study and compare both frameworks. To that end, we impose the restriction that the same amount of information can be certified in both settings. This guarantees that any difference of implementable outcomes is only the result of the bilateral exchange of non-certifiable information between the principal and the agent.

We show that the *canonical mechanism* in bilateral communication has the following simple structure: a three-stage communication mechanism where (i) the agent announces his type, (ii) the principal asks him to certify a specific event, (iii) he certifies an event of his choice, and an *implementation rule* that selects the outcome based on the announced type and whether the requested event was certified. Namely, for the implementation of an outcome function f , if the agent announces type t and certifies the requested event then $f(t)$ is implemented,

otherwise a punishment action is implemented. By applying Theorem 6 of [Bull and Watson \[2007\]](#), who study the introduction of hard evidence to mechanism design, we obtain a partial characterization of the canonical mechanism. In order to explicitly determine the principal's message in step (ii) and the implementation rule, we use the fact that there is only one agent with type-independent preferences.

Having identified the canonical mechanism, we establish the necessary and sufficient conditions for the implementation of any outcome function in both settings and we show that the sets of implementable outcome functions coincide if and only if the *normality* condition is satisfied. This condition states that every type can certify a maximal evidence event, i.e. an event that is equivalent to certifying all information about that type. In other words, unilateral and bilateral communication are outcome equivalent only when there are no effective limitations on the amount of information that can be certified. Bilateral communication is potentially beneficial to the principal only in settings where normality is not satisfied. The hiring process is an example of such a setting if it is not possible to certify all (available) events (at least for some types).

Our main result gives sufficient conditions for bilateral communication to improve the outcome for the principal in comparison with unilateral communication. It is essentially shown that if the principal's first-best is well defined (not necessarily by a unique outcome function) and is not achievable in unilateral communication but would be achieved if any amount of information can be certified then bilateral communication strictly increases the principal's expected payoff. In other words, the principal gains from being active in the communication phase if she is unable to achieve her first-best in unilateral communication because of the cost or time constraints that limit information certification.

As an extension, we examine whether the results in the literature about commitment and outcome randomization hold in our framework. [Sher \[2011\]](#) studies the unilateral communication setting and shows that under a concavity assumption, namely that the principal's utility

function is a type-dependent concave transformation of the agent’s utility function, the principal needs neither commitment over the implementation rule nor randomization of outcome. These results are in fact generalizations of the findings of [Glazer and Rubinstein \[2006\]](#) who considered only binary action spaces, for which the concavity assumption is always satisfied. [Hart et al. \[2016\]](#) show that commitment is unnecessary for a class of certifiability structures (which satisfy normality) and strongly single-peaked preferences. We show that, in bilateral communication under the conditions of our main result and the concavity assumption stated above, randomization is not necessary if the action space is continuous but we give an example with a discrete action space where it is needed. We also give an example where commitment is necessary under the same conditions.

Related Literature Certifiable information has been extensively studied in both sender-receiver games and mechanism design by authors such as [Green and Laffont \[1986\]](#), [Glazer and Rubinstein \[2001, 2004\]](#), [Forges and Koessler \[2005\]](#), [Bull \[2008\]](#), [Deneckere and Severinov \[2008\]](#), [Ben-Porath and Lipman \[2012\]](#), [Kartik and Tercieux \[2012\]](#), [Koessler and Perez-Richet \[2014\]](#), [Sher \[2011\]](#), [Sher and Vohra \[2015\]](#), [Singh and Wittman \[2001\]](#), [Strausz \[2016\]](#). These papers, among others, give rise to important results about implementable allocation rules and some of them establish a revelation principle for settings with certifiable information.

[Bull and Watson \[2007\]](#) study hard evidence in a general mechanism design setting (with multiple agents) and characterize a three-stage communication mechanism in Theorem 6 which we use to determine the canonical bilateral communication mechanism as explained above. We also apply the revelation principle in unilateral communication given in their Theorem 2. Moreover, we show that unilateral and bilateral communication are outcome equivalent if and only if normality is satisfied. The if part holds in general and is established in their Theorem 5. The only if part relies on the agent’s type-independent preferences.

Sher [2014] is closely related to our work but focuses on a framework where the decision space of the principal is binary. It is shown for instance, that unilateral communication is optimal under *foresight* which is a condition related to, but weaker than normality. We note that our main result does not apply in that framework because it requires punishment to be non optimal which is impossible with a binary action space.

1.2 The model

1.2.1 The environment

Consider a setting where a principal faces an agent who is privately informed about his type t in a finite set of agent types T . The principal ignores t but knows the probability distribution of types. We assume the existence of a *certifiability structure* $\mathcal{C} \subseteq 2^T$, where for every $t \in T$ there exists $C \in \mathcal{C}$ such that $t \in C$. We denote by $\mathcal{C}(t) = \{C \in \mathcal{C} : t \in C\}$ the set of events the agent can certify when his type is t .

The principal has to implement an action a in A . Prior to her decision, she can communicate with the agent. The principal's goal is to design the *communication mechanism* along with an *implementation rule*. There are two types of communication mechanisms:

- **Unilateral Communication:** Only the agent is active in the communication. He can certify an event C in \mathcal{C} (which has to be in $\mathcal{C}(t)$ if his type is t) and (possibly) send a message m in some predetermined set (independent of the true type).
- **Bilateral Communication:** Both the agent and the principal partake in sequential communication. The mechanism has to specify the active player at each node and the set of available messages at that node. The only constraint is that, along every possible path, the agent must have exactly one node at which he can certify an event C in \mathcal{C} . At every node, the active player chooses a message from a predetermined set of messages.

With these definitions, we can analyze the benefit of active communication for the principal insofar as the same amount of information is certified in both mechanisms: if bilateral communication increases the principal's expected payoff in comparison with unilateral communication then the difference is due only to the non-certifiable information that is exchanged back and forth between the agent and the principal.

The requirement that the agent does not certify more than one event in \mathcal{C} corresponds to given constraints on the amount of information that can be verified during an interaction between the agent and the principal¹. For example, such constraints apply if the agent has limited time to present this information or the principal has limited time to check it. Hiring processes generally fall in this category when it is impossible to verify whether the applicant fits all the requirements of the job. Recruiters have to decide which aspects to verify and which aspects to ignore.

The implementation rule specifies the principal's action for every possible history in the communication mechanism. In the case of unilateral communication, the history contains exactly one node so that the principal's action is simply a function of the information that the agent presents.

An outcome function $f : T \rightarrow \Delta(A)$ is a mapping from types to lotteries over actions. The agent has a utility function $u : A \rightarrow \mathbb{R}$ which is independent of his type. Let a_0 denote an action such that $u(a_0) = \min_{a \in A} u(a)$ whenever the minimum exists². Throughout the chapter, a_0 will be called the punishment action and the value of $u(a_0)$ will be set to 0 w.l.o.g. We also assume that u is not constant over A (otherwise all outcome functions would be implementable). The principal has a utility function $v : T \times A \rightarrow \mathbb{R}$ which not only depends on the action she chooses

¹This is essentially without loss of generality because if we want to model a limitation to N events instead of one, we would have to replace \mathcal{C} with the set of events that combine up to N elements of \mathcal{C} , i.e. $\{\cap_{i=1}^n C_i \text{ s.t. for all } i, C_i \in \mathcal{C} \text{ and } n \leq N\}$.

²It is implicitly assumed that the minimum exists in all the results stated in this chapter. But essentially, the results still hold with minor modifications if this assumption is not satisfied. See Appendix for a study of the case where the infimum is not reached.

to implement, but also on the type of the agent.

1.2.2 The canonical form

Consider the following communication mechanism:

Definition 1.1. A three-stage communication mechanism is a bilateral communication mechanism with the following timing:

- Stage 1 : The agent reports a type.
- Stage 2 : The principal asks the agent to certify a particular event.
- Stage 3 : The agent certifies an event of his choice.

In stage 1, the agent makes a claim by reporting a type $t \in T$. Then the principal asks him to certify a particular event in \mathcal{C} . Her choice at stage 2, is given by $\sigma : T \rightarrow \Delta(\mathcal{C})$ with $\sigma(t; C)$ denoting the probability of asking the agent to certify C given that he announced type t . In stage 3, the agent certifies C' (either the requested C or a different event).

Definition 1.2. For given $f : T \rightarrow \Delta(A)$ and $\sigma : T \rightarrow \Delta(\mathcal{C})$, the (σ, f) -mechanism is a three-stage communication mechanism along with an implementation rule such that:

- σ is used in stage 2.
- If the agent certifies the requested event the outcome $f(t)$ is implemented, otherwise the punishment action a_0 is implemented.

As we show in the last part of this section, (σ, f) -mechanisms are *canonical* in the sense that we can restrict attention to such mechanisms when studying the implementation of a given outcome function f . Furthermore, this implementation is achieved with truthful reporting in the first stage: if σ is such that for every type t , an agent of type t has no incentive to report

a different type and is able to certify any C that is requested with positive probability $\sigma(t; C)$ (i.e. the support of $\sigma(t)$ is in $\mathcal{C}(t)$) then f is implementable in the (σ, f) -mechanism.

Proposition 1.1. *If f is implemented using a general bilateral communication mechanism and a general implementation rule then there exists $\sigma : T \rightarrow \Delta(\mathcal{C})$ such that it is also implemented in the (σ, f) -mechanism with truthful reporting in stage 1.*

Proof. See Appendix. □

The argument of this proof is split in two steps. First, note that a bilateral communication mechanism is an extensive form game with three types of nodes :

- Message nodes : one player (principal or agent) sends a message to the other.
- Certification nodes : the agent certifies an event.
- Terminal nodes : the principal implements an outcome.

such that along every path in the game tree, there is exactly one certification node. Theorem 6 of [Bull and Watson \[2007\]](#) guarantees that if an outcome function is implementable using such a general mechanism then it is also implementable using a three-stage mechanism with truthful reporting in stage 1:

- Stage 1 : the agent reports his type to the principal.
- Stage 2 : the principal sends a message to the agent.
- Stage 3 : the agent certifies an event.

This mechanism is similar to our three-stage communication mechanism except that instead of directly asking for evidence, the message of stage 2 identifies an information set for the agent in the original extensive form game, more specifically, the one where he has to present evidence.

In the second step, we use the fact that, in our framework, there is only one agent whose preferences are the same across types in order to show that we can restrict attention even further and focus only on (σ, f) -mechanisms.

1.3 Implementable outcome functions

Proposition 1.1 asserts that an outcome function f is implementable in bilateral communication if and only if there exists σ that implements it (in the (σ, f) -mechanism). In this section, we determine the necessary and sufficient conditions for an outcome function to be implementable in bilateral communication. Then, we characterize implementable outcome functions in unilateral communication and identify the link between the two types of mechanisms.

Lemma 1.1. *σ implements f with truthful reporting if and only if*³

$$\begin{aligned} \forall t, \quad \sigma_{tt} &= 1 \\ \forall t, t' \quad \sigma_{t't} &\leq \frac{u(f(t'))}{u(f(t))} \end{aligned}$$

where $\sigma_{t't} = \sum_{C \in \mathcal{C}(t')} \sigma(t; C)$ is the probability for an agent of type t' to successfully persuade the principal that he is of type t .

Proof. See Appendix. □

The first set of conditions say that the principal asks only for events that an agent of the announced type can certify. This guarantees that if the agent reports truthfully then the

³ $u(f(t))$ denotes in general the expectation of agent's utility given the lottery over actions $f(t)$.

principal will certainly implement the right outcome. The second set of conditions are in fact the incentive compatibility constraints of the agent, which ensure that he reports his type truthfully in the first stage. Truthful reporting in stage 1 is necessary to implement the outcome function, and these conditions make sure that the agent has incentive to tell the truth and that the principal does not make the mistake of asking an agent who reported his true type for evidence he cannot present, which in turn, would induce punishment erroneously.

Using Lemma 1.1 we can determine the necessary and sufficient conditions for an outcome function f to be implementable in bilateral communication. We focus on the strategies satisfying the first set of conditions, i.e. strategies such that the support of $\sigma(t)$ is contained in $\mathcal{C}(t)$ for all types t , and we study the existence of an incentive compatible strategy among them. Consider an indexing of types in T from 1 to n : $T = \{t_1, \dots, t_n\}$. Let q_k be the number of events that are certifiable by type t_k : $q_k = \text{card}(\mathcal{C}(t_k))$. $\mathcal{C}(t_k)$ may then be written as $\mathcal{C}(t_k) = \{C_k^1, \dots, C_k^{q_k}\}$. The vector $\sigma(t_k, C)|_{C \in \mathcal{C}(t_k)}$ describes a point M_k in \mathbb{R}^{q_k} . Using this definition, the second set of conditions of Lemma 1.1 can be interpreted as linear inequalities satisfied by the coordinates of the M_k 's for $k \in \{1, \dots, n\}$. From this formulation, we can derive the following result about the implementability of an outcome function f :

Proposition 1.2. *An outcome function f is implementable in bilateral communication if and only if for all $k \in \{1, \dots, n\}$, the following linear program P_k has a value greater than or equal to 1:*

$$\begin{aligned} \text{Max} \quad & c \cdot x \\ \text{s.t.} \quad & Ax \leq b \\ & x \geq 0, \end{aligned}$$

where $x, c \in \mathbb{R}^{q_k}$, $b \in \mathbb{R}^{n-1}$ and A a matrix $(n-1) \times q_k$. $\forall l \in \{1, \dots, q_k\}, \forall j \in \{1, \dots, k-1, k+1, \dots, n\}$, $c_l = 1$, $b_j = \frac{u(f(t_j))}{u(f(t_k))}$ and $A_{jl} = \mathbf{1}_{\{t_j \in C_k^l\}}$.

Proof. From Lemma 1.1, we know that f is implementable (in bilateral communication) if and only if there exists a strategy σ such that

$$\begin{aligned} \forall k, \quad \sigma_{kk} &= 1 \\ \forall k, \forall j, \quad \sigma_{jk} &\leq \frac{u(f(t_j))}{u(f(t_k))} \end{aligned}$$

For a given $k \in \{1, \dots, n\}$, let $x \in \mathbb{R}^{q_k}$ denote the vector $\sigma(t_k, C)|_{C \in \mathcal{C}(t_k)}$, i.e. $x_l = \sigma(t_k, C_k^l)$. The condition $\sigma_{kk} = 1$ is then equivalent to the condition $\sum_{l \in \{1, \dots, q_k\}} x_l = c \cdot x = 1$, where $c \in \mathbb{R}^{q_k}$ and $\forall l, c_l = 1$. Consider the matrix A such that, $\forall l \in \{1, \dots, q_k\}, \forall j \in \{1, \dots, k-1, k+1, \dots, n\}$, $A_{jl} = \mathbb{1}_{\{t_j \in C_k^l\}}$. We can then write $\sigma_{jk} = (Ax)_j$ for every j . By defining the vector $b \in \mathbb{R}^{n-1}$ such that $b_j = \frac{u(f(t_j))}{u(f(t_k))}$, we conclude that the set of conditions on σ_{jk} for $j \in \{1, \dots, k-1, k+1, \dots, n\}$ is equivalent to $Ax \leq b$. So far, we have shown that f is implementable if and only if for every k there exists a vector $x \in \mathbb{R}^{q_k}$, with positive coordinates, such that

$$\begin{aligned} c \cdot x &= 1 \\ Ax &\leq b \end{aligned}$$

If such a vector exists, then the value of P_k is at least 1. Conversely, if x^* is the solution of P_k , with $v = c \cdot x^*$ greater than 1, then the vector $x = \frac{1}{v}x^*$ satisfies the conditions above. □

The implementability of an outcome function f is therefore equivalent to conditions on the values of n linear programs. Moreover, if these conditions are satisfied then we obtain a σ that implements f : σ such that $\sigma(t_k, C)|_{C \in \mathcal{C}(t_k)}$ is a solution of P_k divided by its value.

In the second part of this section, we focus on implementation in unilateral communication. The standard revelation principle (see Theorem 1 of [Bull and Watson \[2007\]](#) or Proposition 2 of [Myerson \[1982\]](#)) applies in this context: if an outcome function f is implementable in unilateral

communication then it is implementable in a unilateral communication mechanism where the agent reports a type and certifies an event in \mathcal{C} with truthful type reporting. Using this fact, we can characterize implementable outcome functions in unilateral communication.

In bilateral communication, σ is called *deterministic* if for every type t , there exists an event C that is requested with certainty if type t is announced in stage 1, i.e. $\sigma(t; C) = 1$.

Definition 1.3. An outcome function f is *implementable in deterministic bilateral communication* if there exists a deterministic σ such that f is implemented in the (σ, f) -mechanism.

The fact that a deterministic σ maps every type to one event with certainty makes it possible to reduce the communication game to a single stage as in the models of [Glazer and Rubinstein \[2006\]](#) and [Sher \[2011\]](#). Consider an outcome function f and a deterministic σ that implements it. In the (σ, f) -mechanism, if the agent announces a type t then the principal asks him for some C with certainty (which can be denoted $\sigma(t)$), and if he certifies it the outcome $f(t)$ is implemented, otherwise the outcome a_0 is implemented. It becomes clear that f is implementable in unilateral communication as follows: if an agent of type t reports t and certifies $\sigma(t)$ then $f(t)$ is implemented, otherwise a_0 is implemented.

Notice that if an agent wants to get the outcome $f(t)$ for some type t , he just has to be able to certify $\sigma(t)$. Therefore if the agent strictly prefers $f(t)$ to $f(t')$, then the incentive compatibility constraint implies that t' is not in $\sigma(t)$. This property is formalized in the following definition:

Definition 1.4. An outcome function f is *evidence compatible* if for every type t there exists C in $\mathcal{C}(t)$ such that:

$$\forall t', \text{ if } u(f(t')) < u(f(t)) \text{ then } t' \notin C.$$

The evidence compatibility of an outcome function f means that every type t can certify an event that no type with a worse outcome than $f(t)$ can certify. The previous analysis shows that if an outcome function is implementable in deterministic bilateral communication then it is evidence compatible.

We conclude this analysis with the following equivalence result:

Proposition 1.3. *Let f be an outcome function. The three following statements are equivalent:*

(i) *f is implementable in unilateral communication.*

(ii) *f is evidence compatible.*

(iii) *f is implementable in deterministic bilateral communication.*

Proof. See Appendix. □

Propositions 1.2 and 1.3 characterize the sets of implementable outcome functions in bilateral and unilateral communication. In the remainder of this section, we identify the necessary and sufficient condition for these sets to coincide. This analysis is interesting insofar as it allows us to determine when the principal can potentially benefit from being active in the communication phase. Let $c^*(t)$ denote the intersection of all events that type t can certify:

$$c^*(t) = \bigcap_{C \in \mathcal{C}(t)} C,$$

Definition 1.5. The certifiability structure \mathcal{C} is called *normal*⁴ if for every type t there exists a certifiable event providing maximal evidence about t , that is:

$$\forall t \in T, c^*(t) \in \mathcal{C}$$

Proposition 1.4. *The sets of implementable outcome functions in unilateral and bilateral communication coincide if and only if the certifiability structure is normal.*

Proof. First, note that outcome functions that are implementable in unilateral communication are also implementable in bilateral communication (see Proposition 1.3).

⁴This condition is called normality by [Bull and Watson \[2007\]](#). It has also been called the full reports condition by [Lipman and Seppi \[1995\]](#) and the minimal closure condition by [Forges and Koessler \[2005\]](#).

If the certifiability structure is normal, Theorem 5 of [Bull and Watson \[2007\]](#) implies that outcome functions that are implementable in bilateral communication are also implementable in unilateral communication. More specifically, assume that the certifiability structure is normal and consider an outcome function f that is implementable in bilateral communication. From normality and Lemma 1.1 we get that f is evidence compatible and therefore implementable in unilateral communication (by Proposition 1.3). Thus the two sets of implementable outcome functions coincide under normality.

To prove the converse, we assume that the certifiability structure is not normal and we construct an outcome function that is implementable in bilateral communication but not implementable in unilateral communication. Under non-normality there exists a type \bar{t} such that $c^*(\bar{t}) \notin \mathcal{C}$. $c^*(\bar{t})$ is not empty (it contains at least \bar{t}) and does not contain all types: if we had $c^*(\bar{t}) = T$ then $\mathcal{C}(\bar{t}) = \{T\}$ and as a consequence $c^*(\bar{t})$ would be in \mathcal{C} .

Consider an action a such that $u(a) > u(a_0) = 0$ and the outcome function f_λ defined by:

$$f_\lambda(t) = \begin{cases} a & \text{if } t \in c^*(\bar{t}) \\ (a_0; a) \text{ with proba. } (\lambda; 1 - \lambda) & \text{if } t \notin c^*(\bar{t}) \end{cases}$$

Types in $c^*(\bar{t})$ can certify any event that \bar{t} can certify. Any type that is not in $c^*(\bar{t})$ is unable to certify at least one event in $\mathcal{C}(\bar{t})$. The outcome function f_λ separates types in two sets and gives a higher payoff to the set of types that can certify any event in $\mathcal{C}(\bar{t})$. Let $\bar{\lambda} = \frac{1}{\text{card}(\mathcal{C}(\bar{t}))}$, and consider σ defined as follows:

$$\sigma(t; C) = \begin{cases} \bar{\lambda} & \text{if } t \in c^*(\bar{t}) \text{ and } C \in \mathcal{C}(\bar{t}) \\ 1 & \text{if } t \notin c^*(\bar{t}) \text{ and } C = T \\ 0 & \text{otherwise} \end{cases}$$

In the (σ, f_λ) -mechanism, if the agent reports a type t in $c^*(\bar{t})$ (i.e., he wants to get the payoff

$u(a)$ with certainty), the principal selects an element in $\mathcal{C}(\bar{t})$ randomly (with equal probability) and asks him to certify it. If the agent reports any other type, he is not required to certify any event and he gets $u(a)$ with probability $1 - \lambda$ and 0 with probability λ . It is readily verifiable that this mechanism implements f_λ for any λ smaller than $\bar{\lambda}$.

We now show that f_λ is not implementable in unilateral communication by proving that it is not evidence compatible (see Proposition 1.3). Indeed, the evidence compatibility condition of f_λ would imply that there exists C in $\mathcal{C}(\bar{t})$ that does not contain any type t outside of $c^*(\bar{t})$. Such an event can only be $c^*(\bar{t})$ which is not in \mathcal{C} . Thus, f_λ is not evidence compatible. □

Example 1.1. Let $T = \{t_1, t_2, t_3\}$ and $\mathcal{C} = \{\{t_1, t_3\}, \{t_2, t_3\}\}$. \mathcal{C} does not satisfy normality: $c^*(t_3) = \{t_3\}$ is not certifiable. Implementable outcome functions in unilateral communication are the evidence compatible outcome functions. If f is evidence compatible, it follows that $u(f(t_3)) \geq \max\{u(f(t_1)), u(f(t_2))\}$. The reason is that t_1 (respectively, t_2) cannot certify an event that does not contain t_3 . Moreover, $u(f(t_3))$ cannot be strictly greater than $\max\{u(f(t_1)), u(f(t_2))\}$, otherwise t_3 would have to certify an event that contains neither t_1 nor t_2 , i.e. the event $\{t_3\}$ which is not certifiable. If $u(f(t_3)) = \max\{u(f(t_1)), u(f(t_2))\}$, it is easy to check that f is evidence compatible. In conclusion, f is implementable in unilateral communication if and only if:

$$u(f(t_3)) = \max\{u(f(t_1)), u(f(t_2))\}.$$

Implementable outcome functions in bilateral communication are those that satisfy the conditions of Proposition 1.2. For type t_1 (respectively, t_2), we only need to have $u(f(t_1))$ (respectively, $u(f(t_2))$) smaller or equal to $u(f(t_3))$. For type t_3 , the value of the following linear program has to be greater or equal to 1:

$$\begin{aligned}
& \text{Max} && x_1 + x_2 \\
& \text{s.t.} && x_1 \leq \frac{u(f(t_1))}{u(f(t_3))} \\
& && x_2 \leq \frac{u(f(t_2))}{u(f(t_3))} \\
& && x_1 \geq 0, x_2 \geq 0.
\end{aligned}$$

That is equivalent to the following condition: $u(f(t_1)) + u(f(t_2)) \geq u(f(t_3))$. In conclusion, f is implementable in bilateral communication if and only if:

$$\max\{u(f(t_1)), u(f(t_2))\} \leq u(f(t_3)) \leq u(f(t_1)) + u(f(t_2)).$$

Because \mathcal{C} does not satisfy normality, bilateral communication allows the implementation of more outcome functions than unilateral communication. If the certification structure is *normalized*, i.e. if the event $c^*(t_3) = \{t_3\}$ is added to \mathcal{C} , it is easy to check that f is implementable in unilateral (respectively, bilateral) communication if and only if $u(f(t_3)) \geq \max\{u(f(t_1)), u(f(t_2))\}$.

1.4 The value of bilateral communication

We know that bilateral communication enlarges the set of implementable outcome functions if and only if the certifiability structure \mathcal{C} does not satisfy normality (see Proposition 1.4). In this section, we establish sufficient conditions for bilateral communication to (strictly) increase the principal's expected payoff. Assume the action space A is a subset of \mathbb{R} (with $a_0 = \min A$ and $A \neq \{a_0\}$), the agent's utility function u is increasing, and both u and v are continuous (on any interval $I \subseteq A$). Frameworks where the principal chooses a reward, a salary or a budget allocation for the agent fit this description.

Definition 1.6. An outcome function f is *weakly evidence compatible* if

$$\forall t, \forall t', \text{ if } u(f(t')) < u(f(t)) \text{ then } t' \notin c^*(t).$$

Recall that an outcome function f is evidence compatible if every type t can certify an event that no type with an outcome worse than $f(t)$ can certify. Weak evidence compatibility only requires that for every type t , no type with an outcome worse than $f(t)$ can certify all events in $\mathcal{C}(t)$. Note that if \mathcal{C} satisfies normality, both notions are equivalent.

Remark 1.1. If an outcome function f is implementable in bilateral communication then it is weakly evidence compatible. Indeed, if f is not weakly evidence compatible then there exist two types t and t' such that t' is in $c^*(t)$ and $u(f(t')) < u(f(t))$ and therefore the incentive compatibility constraint of t' is violated because he can perfectly mimic t . However, weak evidence compatibility is not sufficient for implementation (see Example 1.3).

As a consequence of this observation and Proposition 1.4, there exist outcome functions that are weakly evidence compatible but not evidence compatible if \mathcal{C} does not satisfy normality.

Definition 1.7. The principal's utility function v is *single-plateau* if for every type t there exists \underline{a}_t and \bar{a}_t such that $v(t, \cdot)$ is strictly increasing before \underline{a}_t , constant between \underline{a}_t and \bar{a}_t , and strictly decreasing after \bar{a}_t , i.e. for any action a in $[\underline{a}_t, \bar{a}_t]$, $v(t, a) = v(t, \underline{a}_t) = v(t, \bar{a}_t)$ and for all actions a' and a'' in A :

$$\text{if } a'' < a' \leq \underline{a}_t \text{ or } \bar{a}_t \leq a' < a'' \text{ then } v(t, a'') < v(t, a').$$

If in addition $\underline{a}_t = \bar{a}_t = a_t^*$ for every type t , then v is *single-peaked* at a^* .

Example 1.2. In the hiring process example, assume the agent has (regardless of t) a quadratic disutility of work: if he works h hours, his disutility is $\frac{h^2}{2}$. If the wage he obtains is a then his surplus is $h(a - \frac{h}{2})$. Therefore the optimal number of hours for the agent is $h = a$. Let s_t be

the gross hourly surplus that an agent of type t generates. The principal's payoff if she hires type t at an hourly wage a is therefore $v(t, a) = a(s_t - a)$. It is single-peaked at $\frac{s_t}{2}$.

Let $A_t^* = \arg \max_{a \in A} v(t, a)$ and if it is nonempty for all t , i.e. if $v(t, \cdot)$ reaches its maximum in A for all t , let $F^*(v)$ denote the set of first-best outcome functions:

$$F^*(v) = \{f : T \rightarrow \Delta(A) \mid f(t) \in \Delta(A_t^*)\},$$

If v is single-plateau, $A_t^* = [\underline{a}_t, \bar{a}_t]$ and the principal wants his action to be in A_t^* if the agent's type is t , or as close as possible to this interval.

Proposition 1.5. *Bilateral communication strictly increases the principal's expected payoff if*

(i) *For all t , A_t^* nonempty and $a_0 \notin A_t^*$.*

(ii) *No first-best outcome f^* in $F^*(v)$ is evidence compatible.*

(iii) *There exists a weakly evidence compatible first-best outcome f^* in $F^*(v)$.*

If randomization over actions is not allowed the result holds if in addition, v is single-plateau and A is an interval.

Proof. See Appendix. □

Proposition 1.5 gives sufficient conditions for bilateral communication to improve the outcome for the principal in comparison with unilateral communication. Such an improvement results only from the principal being active in the communication phase given that the same constraints on information certification apply.

Condition (i) guarantees the existence of at least one first-best outcome function, which would be implemented if the principal can observe the agent's type. It also states that, regardless of the agent's type, punishment is not optimal. Under this condition, if f is an optimal

outcome function in unilateral communication then $u(f(t)) > 0$ for all t . As a consequence, the threat of punishment can be used to increase the principal's expected payoff through bilateral communication (regardless of the specifics of utility functions and type distribution).

Condition (ii) guarantees that no first-best outcome function f^* is implementable in unilateral communication (by Proposition 1.3). This condition is necessary for bilateral communication to improve the outcome for the principal.

Condition (iii) states that there exists a weakly evidence compatible first-best outcome function f^* . This implies that f^* would be implementable in unilateral communication if the certification structure is *normalized*, i.e. if the maximal evidence events $\{c^*(t)\}_{t \in T}$ are added to \mathcal{C} . In other words, f^* can be implemented if the principal can ask the agent to certify all events in $\mathcal{C}(t)$ when he reports type t . Under this condition, the constraint on the amount of evidence that can be certified is the reason that the principal is unable to achieve her first-best in unilateral communication. Note that under conditions (ii) and (iii), \mathcal{C} does not satisfy normality (see Definition 1.6).

In general, the result depends on the possibility of randomization over actions. But if v is single-plateau and A is an interval, it holds even if randomization is not allowed: instead of improving the outcome by finding an implementable function with a higher probability of choosing an optimal action, we can simply choose a closer action to the interval of optimal actions.

Example 1.3. Consider an employer (principal) who wants to design a hiring process for a job at her firm. There are different profiles (types in the set T) of applicants that fit the description of this job. However, these profiles are not equally valued by the employer due to differences in productivity. The action space is \mathbb{R}_+ : she chooses the wage at which she is willing to hire an applicant (agent). The punishment action is to reject the application, i.e. to choose a wage equal to 0. The applicant wants the highest possible wage. The hiring process is subject to a time limit which implies that a limited amount of information (about the applicant's skills) can

be verified. Therefore, a certifiability structure \mathcal{C} can be defined.

If the principal wants to implement an allocation f , where $f(t)$ is the wage for type t , she can use canonical form bilateral communication mechanism. The hiring process starts when the agent reports a type t by sending his curriculum vitae (which describes his profile). The principal then asks him to certify an element of \mathcal{C} by testing his abilities in certain tasks and/or by asking for third party certifications (such as diplomas). If the applicant passes the test and/or provides the required certificates, he is hired at wage $f(t)$. Otherwise, he is not hired (punishment action). For simplicity, we choose to preclude randomization over actions because they represent wages. A similar analysis can be conducted if randomization is allowed.

Let $T = \{t_1, t_2, t_3\}$ and $\mathcal{C} = \{\{t_1, t_3\}, \{t_2, t_3\}\}$: there are two skills and three possible types with the possibility to verify only one skill during the hiring process. The first type masters the first skill, the second masters the other, while the third masters both skills. Note that \mathcal{C} does not satisfy normality: $c^*(t_3) = \{t_3\}$ is not certifiable.

Let $u(a) = a$: the agent's utility is equal to his wage. Assume that the employer's utility v single-peaked at a^* . We give an instance where this condition is satisfied in Example 1.2 (a change of variable would allow us to have $u(a) = a$ and keep v single-peaked).

As established in Example 1.1, f is implementable in unilateral communication if and only if $f(t_3) = \max\{f(t_1), f(t_2)\}$ and implementable in bilateral communication if and only if $\max\{f(t_1), f(t_2)\} \leq f(t_3) \leq f(t_1) + f(t_2)$. We can easily check that f is weakly evidence compatible if and only if $f(t_3) \geq \max\{f(t_1), f(t_2)\}$.

Let a_3^* be strictly larger than a_1^* and a_2^* : the first-best wage for t_3 is strictly higher than the first-best wages for t_1 and t_2 . That means a^* is weakly evidence compatible but not evidence compatible. If in addition, $a_k^* > 0$ for all k then all conditions of Proposition 1.5 are satisfied. Therefore, bilateral communication strictly increases the principal's payoff in comparison with unilateral communication. In the remainder, we examine how the payoff increase is achieved. Assume w.l.o.g that $a_1^* = \min_k a_k^*$.

If a^* is such that $a_3^* \leq a_1^* + a_2^*$ then it is implementable in bilateral communication. Otherwise, consider f optimal in unilateral communication. We have $f(t_3) = \max\{f(t_1), f(t_2)\}$. Moreover, $f(t_k)$ is in $[a_1^*, a_3^*]$ for all k because f is optimal (in unilateral communication) and v is single-peaked. We show how to construct a function \hat{f} that gives the principal a higher payoff than f .

If f is such that $f(t_1) \leq f(t_2) = f(t_3) < a_3^*$, define \hat{f} such that $\hat{f}(t_k) = f(t_k)$ for k in $\{1, 2\}$ and

$$\hat{f}(t_3) = \min\{f(t_1) + f(t_2), a_3^*\}.$$

\hat{f} is implementable in bilateral communication. Given that $f(t_1) \geq a_1^* > 0$, it follows that $f(t_3) < \hat{f}(t_3) \leq a_3^*$. Consequently, \hat{f} gives a strictly higher payoff to the principal than f (because v is single-peaked).

If f is such that $f(t_1) \leq f(t_2) = f(t_3) = a_3^*$, define \hat{f} such that $\hat{f}(t_k) = f(t_k)$ for k in $\{1, 3\}$ and

$$\hat{f}(t_2) = \max\{f(t_3) - f(t_1), a_2^*\}.$$

\hat{f} is implementable in bilateral communication. We have $a_2^* \leq \hat{f}(t_2) < f(t_3)$. Therefore, \hat{f} gives a strictly larger payoff to the principal than f (because v is single-peaked). Similar arguments apply if $f(t_2) < f(t_1) = f(t_3)$.

1.5 Extensions

1.5.1 Tightness of Proposition 1.5

In this section, we focus on the conditions of the main result. Note that if A_t^* is empty for a given type t , these conditions are not well defined. For that reason, we assume that A_t^* is nonempty for all t and we show that if any of the other conditions of Proposition 1.5 is dropped we can construct an example where the result does not hold, i.e. where the outcome is not improved by bilateral communication. In all examples we have $u(a) = a$, $T = \{t_1, t_2, t_3\}$ and $\mathcal{C} = \{\{t_1, t_3\}, \{t_2, t_3\}\}$ as in Example 1.3. In the general case, i.e. when randomization is allowed, we choose $A = \mathbb{R}_+$ and single-peaked piecewise linear v which ensures that randomization over actions does not improve the outcome for the principal.

Non-optimal punishment condition

Let a^* be such that $0 = a_1^* < a_2^* < a_3^*$. a^* is weakly evidence compatible but not evidence compatible. Let $v(t_1, a) = -2a$ and $v(t_k, a) = -|a - a_k^*|$ for k in $\{2, 3\}$ and assume types are uniformly distributed.

Let f be an optimal outcome function in bilateral communication. Implementation conditions (see Example 1.3) ensure that $\max\{f(t_1), f(t_2)\} \leq f(t_3) \leq f(t_1) + f(t_2)$. If $f(t_1) = 0$, then we necessarily have $f(t_2) = f(t_3)$ and it is in $[a_2^*, a_3^*]$ (because v is single-peaked at a^*). Such a function is also implementable in unilateral communication. If $f(t_1) > 0$ then $f(t_3)$ must be equal to $f(t_1) + f(t_2)$ and below a_3^* : otherwise we would increase the principal's payoff by lowering $f(t_1)$. It follows that the principal's optimization problem is

$$\begin{aligned}
& \text{Max} && v(t_1, a_1) + v(t_2, a_2) + v(t_3, a_1 + a_2) \\
& \text{s.t.} && a_1 \geq 0, a_2 \geq 0, \\
& && a_1 + a_2 \leq a_3^*.
\end{aligned}$$

Under these constraints, we have

$$v(t_1, a_1) + v(t_2, a_2) + v(t_3, a_1 + a_2) = -a_1 + a_2 + v(t_2, a_2) - a_3^*.$$

It follows that $a_1 = 0$ at the optimum. Therefore, $f(t_1)$ cannot be strictly positive. In conclusion, the optimal f is implementable in unilateral communication and the outcome is not improved by bilateral communication.

Evidence compatibility conditions

It is obvious that if a first-best outcome function is evidence compatible then the result does not hold because it is implementable in unilateral communication. Condition (ii) is necessary for the result to hold.

Let $v(t_k, a) = -|a - a_k^*|$ for all k with $0 < a_1^* < a_3^* < a_2^*$. Condition (iii) is not satisfied in this case because a^* is the unique first-best and is not weakly evidence compatible. If f is optimal in bilateral communication then we necessarily have $f(t_1) = a_1^*$ and $f(t_2) = f(t_3) \in [a_3^*, a_2^*]$, thus it is implementable in unilateral communication. Bilateral communication cannot improve the outcome.

Without randomization over actions

Now we assume randomization over actions is not allowed and we show that if A is not an interval or v is not single-plateau, we can construct an example where the result does not hold.

The action space is not an interval: Let $v(t_k, a) = -|a - a_k^*|$ for all k with $0 < a_1^* < a_2^* < a_3^*$ and $a_1^* + a_2^* < a_3^*$. If $A = [0, a_2^*] \cup [a_3^*, +\infty)$, it is easy to see that if f is optimal in bilateral communication then $f(t_1) = a_1^*$ and $f(t_2) = f(t_3) \in \{a_2^*, a_3^*\}$. Therefore, it is also implementable in unilateral communication.

The principal's utility is not single-plateau: Assume types are uniformly distributed. Let $A = \mathbb{R}_+$ and $v(t_k, a) = -2|a - a_k^*|$ for k in $\{1, 2\}$ with $0 < a_1^* < a_2^*$. However, let $v(t_k, \cdot)$ be increasing strictly increasing. Let there be a_3^* strictly higher than $a_1^* + a_2^*$. Let $v(t_3, \cdot)$ be strictly increasing for $a < a_2^*$, constant between a_2^* and $a_1^* + a_2^*$, strictly increasing between $a_1^* + a_2^*$ and a_3^* , and strictly decreasing for $a > a_3^*$. Assume in addition that $v(t_3, \cdot)$ is piecewise linear with a slope equal to 1 (in absolute value) outside the interval $[a_2^*, a_1^* + a_2^*]$. a^* is the unique first-best and it satisfies the conditions of Proposition 1.5, i.e. it is weakly evidence compatible but not evidence compatible and $a_k^* > 0$ for all k . But v is not single-plateau.

The optimal outcome function in unilateral communication is f such that $f(t_1) = a_1^*$ and $f(t_2) = f(t_3) = a_2^*$. Any outcome function \hat{f} such that $\hat{f}(t_1) = a_1^*$, $\hat{f}(t_2) = a_2^*$ and $\hat{f}(t_3) \in [a_2^*, a_1^* + a_2^*]$ is implementable in bilateral communication and gives the same expected payoff to the principal as f . To see that such \hat{f} is optimal, we need to observe that in order to have $v(t_3, \tilde{f}(t_3)) > v(t_3, \hat{f}(t_3))$, we need to have $\tilde{f}(t_3) > a_1^* + a_2^*$. If \tilde{f} is implementable, $\tilde{f}(t_3) \leq \tilde{f}(t_1) + \tilde{f}(t_2)$ and therefore:

$$\tilde{f}(t_3) - (a_1^* + a_2^*) \leq (\tilde{f}(t_1) - a_1^*) + (\tilde{f}(t_2) - a_2^*).$$

Given the chosen function v , namely the fact that $v(t_1, \cdot)$ and $v(t_2, \cdot)$ have a twice larger slope than $v(t_3, \cdot)$, it follows that \tilde{f} decreases the principal's payoff in comparison with f .

1.5.2 Optimal solutions: Randomization over actions and Commitment

An outcome function f is optimal if it maximizes the principal's expected payoff among the set of implementable outcome functions. In unilateral communication, Sher [2011] shows that there exists an optimal function that does not involve randomization over actions and can be implemented without the principal's commitment if the following concavity assumption is satisfied: the principal's utility function is a type-dependent concave transformation of the agent's utility function. Hart et al. [2016] show that commitment is unnecessary for a class of certifiability structures satisfying a stronger condition than normality if the principal's preferences are strongly single-peaked, that is every convex combination of elements of $\{v(t, \cdot)\}_{t \in T}$ is single-peaked.

In this section, we show that under the conditions of Proposition 1.5 and the concavity assumption stated above, randomization over actions is not needed if the action space is an interval and we give an example with a discrete action space where it is necessary. We also give an example where commitment is necessary at the optimum under the same conditions.

We focus on settings where the conditions of Proposition 1.5 are satisfied because it guarantees that bilateral communication is beneficial and it is then interesting to study the properties of optimal solutions in comparison with unilateral communication.

Randomization over actions

The result of Sher [2011] about randomization holds whether the principal's actions space A is continuous or discrete. We show that it holds in bilateral communication if A is an interval

and we give an example that illustrates the need for randomization for discrete A .

Proposition 1.6. *Assume the conditions of Proposition 1.5, A is an interval and the following concavity assumption are satisfied: for all t , there exists a concave function c_t such that $v(t, \cdot) = c_t(u(\cdot))$. Then there exists an optimal outcome function in bilateral communication f such that for all t , $f(t) \in A$.*

Proof. Consider an outcome function f . For a given type t , the outcome $f(t)$ is a distribution over actions. Let $\mathbb{E}_{f(t)}(u)$ be the expected utility of an agent under the lottery $f(t)$. Because u is continuous over A , there exists an action $\widehat{f}(t) \in A$ such that

$$u(\widehat{f}(t)) = \mathbb{E}_{f(t)}(u)$$

This defines a deterministic outcome function \widehat{f} . If f is implementable then \widehat{f} is also implementable (because the agent's expected utilities are identical for both outcome functions). Now we compare the principal's utilities under f and \widehat{f} when she faces an agent of type t .

$$\begin{aligned} \mathbb{E}_{f(t)}(v(t, \cdot)) &= \mathbb{E}_{f(t)}(c_t(u)) \\ &\leq c_t(\mathbb{E}_{f(t)}(u)) \quad (\text{concavity of } c_t) \\ &\leq c_t(u(\widehat{f}(t))) = v(t, \widehat{f}(t)). \end{aligned}$$

The principal is therefore (weakly) better off not randomizing over actions. The conclusion follows. □

As we can see in the proof above, the fact that A is an interval plays an essential role in the argument. In the following example, we consider a discrete actions space and we find that randomization is necessary at the optimum.

Example 1.4. Consider a setting where $u(a) = a$, $T = \{t_1, t_2, t_3\}$ and $\mathcal{C} = \{\{t_1, t_3\}, \{t_2, t_3\}\}$ as in Example 1.3. Assume types are uniformly distributed, $A = \{0, 1, 3\}$ and let the principal's utility function be given by the following table:

$v(t, a)$	0	1	3
t_1	0	1	0
t_2	0	1	0
t_3	0	1	3

Note that the concavity assumption is satisfied. The best deterministic outcome function in this case is one such that type t_3 receives the action 3 along with one of the other two types, while the remaining type receives the outcome 1. The utility of the principal for such a function is $V = 4$. This function is not optimal though. The optimal solution assigns action 3 to type t_3 and the same randomized outcome to types t_1 and t_2 such that action 1 has a probability $\frac{3}{4}$ and action 3 has a probability $\frac{1}{4}$. The optimal payoff of the principal is $V = \frac{9}{2}$.

Commitment

We now give an example where commitment is necessary at the optimum in bilateral communication under the conditions of Proposition 1.5 and the concavity assumption.

Example 1.5. Consider the same framework as Example 1.1 with the utility functions of Example 1.2. Namely, $T = \{t_1, t_2, t_3\}$, $\mathcal{C} = \{\{t_1, t_3\}, \{t_2, t_3\}\}$, $A = [0, +\infty)$, $u(a) = \frac{a^2}{2}$ and $v(t, a) = a(s_t - a)$ with s_t the surplus that type t generates. Assume in addition that $s_{t_1} = s_{t_2} = 1$ and $s_{t_3} = 2$. v is single peaked at $a^* = (\frac{1}{2}, \frac{1}{2}, 1)$ which is weakly evidence compatible but not evidence compatible. Also, punishment is not optimal regardless of the agent type. All conditions of Proposition 1.5 are satisfied. In addition the concavity assumption is satisfied: for every type t , $v(t, \cdot) = c_t(u(\cdot))$, where $c_t(x) = s_t\sqrt{2x} - 2x$. In fact, the strict concavity ensures that there can be no randomization over actions at the optimum.

It follows from the characterization given in Example 1.1 that f , such that $f(t_k) = a_k$ for all k , is implementable if and only if

$$\max\{a_1^2, a_2^2\} \leq a_3^2 \leq a_1^2 + a_2^2.$$

The optimal solution is such that $a_1 = a_2 = \frac{1+\sqrt{2}}{4}$ and $a_3 = \frac{2+\sqrt{2}}{4}$. The principal needs commitment in order to implement this outcome function because these actions are not rational given her beliefs at the time of implementation: for example, when choosing action $a_3 = \frac{2+\sqrt{2}}{4}$, she knows that the agent's type is t_3 and her rational decision would be $a = 1$.

1.6 Appendix

Proof of Proposition 1.1. Bull and Watson [2007] show that if f is implemented in a general mechanism, then it is also implemented in a special three-stage mechanism characterized by $g : T \times M \times \mathcal{C} \rightarrow \Delta(A)$ and $\sigma : T \rightarrow \Delta(M)$ with truthful reporting at stage 1. $g(t, m, C)$ is the outcome when the agent reports t , the principal sends message m and the agent certifies C . $\sigma(t; m)$ is the probability that the principal sends the message m if the agent reports t . Therefore, for every type t and every message m , there must exist an event $C_{t,m}$ in $\mathcal{C}(t)$ such that the outcome $f(t)$ is implemented whenever the agent announces t , the principal sends m and the agent shows C . Formally:

$$\forall t, \forall m, \exists C_{t,m} \in \mathcal{C}(t) \text{ such that } g(t, m, C_{t,m}) = f(t)$$

For every type t , let ϕ_t be a mapping from messages m to events $C_{t,m}$:

$$\forall m, \phi_t(m) \in \mathcal{C}(t) \text{ and } g(t, m, \phi_t(m)) = f(t)$$

Incentive compatibility constraints are given by:

$$\forall t, \forall t', \sum_m \sigma(t', m) \max_{C \in \mathcal{C}(t)} u(g(t', m, C)) \leq u(f(t)).$$

Consider the mechanism $\hat{\sigma}$ and \hat{g} defined by:

- $\forall t, \forall C, \hat{\sigma}(t, C) = \sum_{m \in \phi_t^{-1}(C)} \sigma(t, m)$.
- $\forall t, \forall C, \hat{g}(t, C, C) = f(t)$.
- $\forall t, \forall C' \neq C, \hat{g}(t, C, C') = a_0$.

We can easily check that $\forall t, \sum_{C \in \mathcal{C}(t)} \hat{\sigma}(t, C) = 1$. Note that this is a description of the three-stage $(\hat{\sigma}, f)$ -mechanism. In order to prove that $\hat{\sigma}$ and \hat{g} implement f , we check that incentive compatibility constraints are satisfied. First, using the definition of $\hat{\sigma}$, we have:

$$\forall t, \forall t', \sum_C \hat{\sigma}(t', C) \max_{C' \in \mathcal{C}(t)} u(\hat{g}(t', C, C')) = \sum_C \sum_{m \in \phi_{t'}^{-1}(C)} \sigma(t', m) \max_{C' \in \mathcal{C}(t)} u(\hat{g}(t', C, C'))$$

By definition, if $m \in \phi_{t'}^{-1}(C)$ then $g(t', m, C) = g(t') = \hat{g}(t', C, C)$, and for $C' \neq C$, we have $u(g(t', m, C')) \geq u(a_0) = u(\hat{g}(t', C, C'))$. Therefore

$$\forall t, \forall t', \sum_C \hat{\sigma}(t', C) \max_{C' \in \mathcal{C}(t)} u(\hat{g}(t', C, C')) \leq \sum_C \sum_{m \in \phi_{t'}^{-1}(C)} \sigma(t', m) \max_{C' \in \mathcal{C}(t)} u(g(t', m, C'))$$

The r.h.s term is equal to $\sum_m \sigma(t', m) \max_{C \in \mathcal{C}(t)} u(g(t', m, C))$ and we finally get:

$$\forall t, \forall t', \sum_C \hat{\sigma}(t', C) \max_{C' \in \mathcal{C}(t)} u(\hat{g}(t', C, C')) \leq u(f(t)).$$

Which proves that $\hat{\sigma}$ implements f . □

Proof of Lemma 1.1. Given the structure of the (σ, f) -mechanism, σ implements f with truthful reporting if and only if (i) the support of $\sigma(t)$ is in $\mathcal{C}(t)$ and (ii) no type has an incentive to misreport.

Let $\sigma_{t't} = \sum_{C \in \mathcal{C}(t')} \sigma(t; C)$ denote the probability for an agent of type t' to successfully persuade the principal that he is of type t . Using this notation, (i) states that for all t , $\sigma_{tt} = 1$. (ii) describes the incentive compatibility constraints and states that for all t and t' , $\sigma_{t't} u(f(t)) \leq u(f(t'))$. Note that for t such that $u(f(t)) = 0$, these conditions are satisfied and do not constrain the choice of $\sigma(t)$. Therefore, we can write (ii) as follows:

$$\forall t, t' \sigma_{t't} \leq \frac{u(f(t'))}{u(f(t))},$$

with the right hand side equal to $+\infty$ if $u(f(t)) = 0$. □

Proof of Proposition 1.3. Let f be an outcome function. Recall the three statements:

- (i) f is implementable in unilateral communication.
- (ii) f is evidence compatible.
- (iii) f is implementable in deterministic bilateral communication.

In order to prove the equivalence, we will show the following implications: (i) \Rightarrow (ii) \Rightarrow (iii) \Rightarrow (i).

(i) \Rightarrow (ii) Consider an outcome function f implementable in unilateral communication. Using the revelation principle, we know that there must exist a unilateral communication mechanism that implements it with truthful type reporting. For every type t , there must exist at least C in $\mathcal{C}(t)$ such that if the agent reports t and certifies C , the principal implements $f(t)$. Consequently, if t' is such that $u(f(t')) < u(f(t))$ then $t' \notin C$: otherwise t' benefits from deviating by reporting t and certifying C . Therefore, f is evidence compatible.

(ii) \Rightarrow (iii) If f is evidence compatible, define the deterministic σ such that for every t , $\sigma(t)$ is an element of $\mathcal{C}(t)$ that contains no type t' with $u(f(t')) < u(f(t))$. It is guaranteed to exist by evidence compatibility. It is readily verified that σ implements f .

(iii) \Rightarrow (i) If there exists a deterministic σ that implements f , consider the following implementation rule in unilateral communication: if the agent reports type t and certifies $\sigma(t)$ implement $f(t)$, otherwise implement a_0 . This mechanism implements f .

□

Proof of Proposition 1.5. Let f be an optimal outcome function in unilateral communication. Proposition 1.3 guarantees the existence of deterministic σ that implements f in bilateral communication. Our goal is to slightly modify the (σ, f) -mechanism so that we obtain an implementable outcome function \hat{f} with a strictly higher expected payoff for the principal than f .

Let f^* be a weakly evidence compatible element of $F^*(v)$ (it is guaranteed to exist by condition (iii)). Consider an indexing of types in T from 1 to n : $T = \{t_1, \dots, t_n\}$. Given that $A_{t_k}^*$ is nonempty for all k , let \bar{a}_k (respectively, \underline{a}_k) denote its largest (respectively, smallest) element. Let the indexing be such that $u(f(t_1)) \leq u(f(t_2)) \leq \dots \leq u(f(t_n))$ and if there exist k and l such that $u(f(t_{k-1})) < u(f(t_k)) = u(f(t_{k+1})) = \dots = u(f(t_{k+l})) < u(f(t_{k+l+1}))$, rearrange the indexing so that $u(f^*(t_k)) \leq u(f^*(t_{k+1})) \leq \dots \leq u(f^*(t_{k+l}))$.

For any two lotteries over actions μ, μ' and any $\alpha \in [0, 1]$ let $L(\mu, \mu', \alpha) = (1 - \alpha)\mu + \alpha\mu'$. If $u(f(t_k))$ is in $[u(\underline{a}_k), u(\bar{a}_k)]$ then $v(t_k, f(t_k))$ is necessarily maximal because f is optimal. Otherwise $f(t_k)$ can be replaced with $L(\bar{a}_k, \underline{a}_k, \alpha)$ where α is such that $u(f(t_k)) = \alpha u(\underline{a}_k) + (1 - \alpha)u(\bar{a}_k)$. This would make $v(t_k, f(t_k))$ maximal without affecting the evidence compatibility constraints of f .

Condition (ii) implies that f is not in $F^*(v)$ (because f is evidence compatible). Therefore, using the previous observation, there must exist \tilde{k} such that $u(f(t_{\tilde{k}}))$ is not in $[u(\underline{a}_{\tilde{k}}), u(\bar{a}_{\tilde{k}})]$,

and as a consequence $v(t_{\tilde{k}}, f(t_{\tilde{k}})) < v(t_{\tilde{k}}, f^*(t_{\tilde{k}}))$. Let $J = \{j : u(f(t_j)) = u(f(t_{\tilde{k}}))\}$. It follows that there exist l, l', m and m' such that

$$J = \{\tilde{k} - l, \dots, \tilde{k} - l', \dots, \tilde{k}, \dots, \tilde{k} + m', \dots, \tilde{k} + m\}$$

with:

$$\begin{cases} u(f^*(t_j)) < u(f^*(t_{\tilde{k}})) & \text{if } \tilde{k} - l \leq j < \tilde{k} - l' \\ u(f^*(t_j)) = u(f^*(t_{\tilde{k}})) & \text{if } \tilde{k} - l' \leq j \leq \tilde{k} + m' \\ u(f^*(t_j)) > u(f^*(t_{\tilde{k}})) & \text{if } \tilde{k} + m' < j \leq \tilde{k} + m \end{cases}$$

It follows from condition (i) that $\underline{a}_k > a_0$ for all k . Therefore, $u(f(t_1)) > 0$: otherwise $f(t_1) = a_0$ and it would be possible to improve the outcome for the principal (and the agent incidentally) by replacing $f(t_1)$ with $L(a_0, a_1, \alpha)$ where $\alpha > 0$ is small enough for f to remain evidence compatible (the same argument applies if more than one type receive the punishment action). Let ε be such that $0 < \varepsilon < \min\{\frac{u(f(t_1))}{u(f(t_{\tilde{k}}))}, \frac{1}{2}\}$. We now construct an outcome function \hat{f} that gives the principal a strictly higher expected payoff than f by having $v(t, \hat{f}(t)) \geq v(t, f(t))$ for all t and $v(t_{\tilde{k}}, \hat{f}(t_{\tilde{k}})) > v(t, f(t_{\tilde{k}}))$.

- (I) If $u(f(t_{\tilde{k}})) > u(\bar{a}_{\tilde{k}})$ then we necessarily have $m > m'$: if m is equal to m' , we can replace $f(t_j)$ with $L(f(t_j), f^*(t_j), \alpha_j)$ for every j in J with $\alpha_j > 0$ such that this lottery gives the agent a payoff equal to $\max\{u(f^*(t_{\tilde{k}})), u(f(t_{\tilde{k}-l-1}))\}$. This change would increase the principal's expected payoff without affecting the evidence compatibility of f , which contradicts the fact that f is optimal in unilateral communication.

Let $\underline{J} = \{\tilde{k} - l, \dots, \tilde{k} + m'\}$ and $\bar{J} = \{\tilde{k} + m' + 1, \dots, \tilde{k} + m\}$. \underline{J} contains \tilde{k} and \bar{J} is nonempty because $m > m'$. We also have $u(f^*(j)) > u(f^*(j'))$ for all $j \in \bar{J}$ and $j' \in \underline{J}$.

For every j in \bar{J} , note that $t_{j'} \notin c^*(t_j)$ for all $j' \in \underline{J}$ because f^* is weakly evidence compatible. Let C_j denote the event $\sigma(t_j)$. C_j may contain types $t_{j'}$ with $j' \in \underline{J}$ but

cannot be certified by any type t_k with $k < \tilde{k} - l$ (because σ implements f).

Let $\hat{\sigma}$ be identical to σ except for types t_j with j in \bar{J} . For each of these types and every j' such that $j' \in \underline{J}$ and $t_{j'} \in C_j$ choose $C_{jj'}$ in $\mathcal{C}(t_j)$ that does not contain $t_{j'}$. Let l_j be the number of $C_{jj'}$'s. If $l_j = 0$, let $\hat{\sigma}(t_j, C_j) = 1$. Otherwise, let $\hat{\sigma}(t_j, C_j) = 1 - \varepsilon$ and $\hat{\sigma}(t_j, C_{jj'}) = \frac{\varepsilon}{l_j}$ for each $C_{jj'}$. Let $\tilde{l} = \max_{j \in \bar{J}} l_j$.

Let \hat{f} be identical to f except for types $t_{j'}$ such that $j' \in \underline{J}$. For each of these types, let $\hat{f}(t_{j'}) = L(f(t_{j'}), f^*(t_{j'}), \alpha_{j'})$ with $\alpha_{j'} > 0$ such that this lottery gives the agent an expected payoff equal to $u(f(t_{\tilde{k}})) - \eta$ for some η satisfying the following condition:

$$0 < \eta \leq u(f(t_{\tilde{k}})) - \max\{u(f^*(t_{\tilde{k}})), u(f(t_{\tilde{k}-l-1}))\}.$$

This outcome function gives the principal a strictly higher expected payoff than f . If $\tilde{l} = 0$, \hat{f} would be evidence compatible and f would not be optimal in unilateral communication. Therefore $\tilde{l} > 0$ and $\hat{\sigma}$ is not deterministic. Moreover, $\hat{\sigma}$ implements \hat{f} if

$$\begin{aligned} (1 - \frac{\varepsilon}{\tilde{l}})u(f(t_{\tilde{k}})) &\leq u(f(t_{\tilde{k}})) - \eta \\ \varepsilon u(f(t_{\tilde{k}})) &\leq u(f(t_1)) \end{aligned}$$

The first condition guarantees that types $t_{j'}$ such that $j' \in \underline{J}$ have no incentive to deviate and is satisfied for $\eta > 0$ small enough. The second condition guarantees that types below $t_{\tilde{k}-l}$ have no incentive to deviate and is satisfied by definition of ε .

- (II) If $u(f(t_{\tilde{k}})) < u(a_{\tilde{k}})$ then $l > l'$: if l is equal to l' , we can replace $f(t_j)$ with $L(f(t_j), f^*(t_j), \alpha_j)$ for every j in J with $\alpha_j > 0$ such that this lottery gives the agent a payoff equal to

$\min\{u(f^*(t_{\tilde{k}})), u(f(t_{\tilde{k}+m+1}))\}$. This change would increase the principal's expected payoff without affecting the evidence compatibility of f , which contradicts the fact that f is optimal in unilateral communication.

Let $\underline{J} = \{\tilde{k} - l, \dots, \tilde{k} - l' - 1\}$ and $\overline{J} = \{\tilde{k} - l', \dots, \tilde{k} + m\}$. \underline{J} is nonempty because $l > l'$ and \overline{J} contains \tilde{k} . We also have $u(f^*(j)) > u(f^*(j'))$ for all $j \in \overline{J}$ and $j' \in \underline{J}$ so that we construct $\hat{\sigma}$ in the same way as in (I).

Let \hat{f} be identical to f except for types t_j with j in \overline{J} . For each of these types, let $\hat{f}(t_j) = L(f(t_j), f^*(t_j), \alpha_j)$ with $\alpha_j > 0$ such that this lottery gives the agent an expected payoff equal to $u(f(t_{\tilde{k}})) + \eta$ for some η satisfying the following condition:

$$0 < \eta \leq \min\{u(f^*(t_{\tilde{k}})), u(f(t_{\tilde{k}+l+1}))\} - u(f(t_{\tilde{k}})).$$

Similarly to (I), this outcome function gives the principal a strictly higher expected payoff than f and we have $\tilde{l} > 0$ and $\hat{\sigma}$ non-deterministic. Moreover, $\hat{\sigma}$ implements \hat{f} if

$$\begin{aligned} (1 - \frac{\varepsilon}{\tilde{l}})(u(f(t_{\tilde{k}})) + \eta) &\leq u(f(t_{\tilde{k}})) \\ \varepsilon(u(f(t_{\tilde{k}})) + \eta) &\leq u(f(t_1)) \end{aligned}$$

These conditions are analogous to those of (I) and are satisfied for $\eta > 0$ small enough.

This concludes the proof in the general case.

In order to prove the result when v is single-plateau, A is an interval and randomization over actions is not allowed, we simply have to replace $L(a, a', \alpha)$ with the action $(1 - \alpha)a + \alpha a'$ for any actions a and a' .

□

Study of the case where a_0 does not exist. This happens when $\inf_{a \in A} u(a)$ is not attained. If $\inf_{a \in A} u(a) = -\infty$ then the punishment can be as big as the principal wants. Formally, f is implementable if and only if there exists $\alpha \in \mathbb{R}$ and σ such that

$$\forall t, t' \quad \sigma_{t't} u(f(t)) - (1 - \sigma_{t't}) \alpha \leq u(f(t'))$$

First, note that if $t' \in c^*(t)$ then $\sigma_{t't}$ is necessarily equal to 1 which implies $u(f(t)) \leq u(f(t'))$. Consider the following mechanism : if the agent reports type t , ask for all events in $\mathcal{C}(t)$ with the same probability. Then $\forall t, t'$, if $t' \notin c^*(t)$, the above inequality is satisfied for α large enough. Because we have a finite number of such inequalities, we can take the largest α to satisfy all of them. We conclude that if $\inf_{a \in A} u(a) = -\infty$, f is implementable if and only if

$$\forall t, t', \text{ if } u(f(t')) < u(f(t)) \text{ then } t' \notin c^*(t),$$

that is, f weakly evidence compatible. If the punishment can be as large as we want, all weakly evidence compatible outcome functions are implementable and the limitation on information certification has no effect.

If on the other hand $\inf_{a \in A} u(a)$ is finite, we can set it to 0 w.l.o.g and denote by a_ϵ an action such that $u(a_\epsilon) = \epsilon$ for all $\epsilon > 0$. By continuity of u , such action always exists. In this case, f is implementable if and only if there exists σ such that

$$\begin{aligned} \exists \epsilon > 0 \text{ s.t } \forall t, t' \quad & \sigma_{t't} u(f(t)) + (1 - \sigma_{t't}) \epsilon \leq u(f(t')) \\ \Leftrightarrow \forall t, t' \quad & \text{if } u(f(t')) < u(f(t)) \text{ then } \sigma_{t't} u(f(t)) < u(f(t')) \end{aligned}$$

We conclude that σ implements f iff

$$\forall t, t', \text{ if } u(f(t')) < u(f(t)) \text{ then } \sigma_{t't} < \frac{u(f(t'))}{u(f(t))}$$

Implementation results follow from Lemma 1.1, where, in this context, certain inequalities are replaced with strict inequalities. The subsequent results still hold but have to be modified accordingly.

□

Chapter 2

Hard evidence and ambiguity aversion

2.1 Introduction

Consider a situation in which a decision-maker wants to implement different actions in different states of the world, but does not observe the true state while an informed agent does. Because the agent's interest might not be aligned with the decision-maker's preferences, he can be asked to certify—either partially or totally—the information he reports. When certain constraints (like time, complexity, cost, or technological constraints) preclude the option of unlimited information certification, the decision-maker has to choose which information has to be certified by the agent. An easy way to illustrate this framework is to consider a situation with multidimensional states (representing, for example, all transactions and activities of a firm) where the decision-maker (for example, a financial auditor) has to choose one dimension of the state to be certified by the firm.

There are several ways to choose which activities to investigate. The first and most basic is to define rules that associate an activity to investigate to each possible state: for example, investigate the largest activity. The problem with such deterministic rules is that they let the agent know with certainty which dimension is going to be investigated, which implies

that he could easily manipulate the decision; hence, the set of implementable allocations will usually be rather small. The second way is for the decision-maker to use mixed strategies: for example, choose randomly between all activities with equal probability. Using such a rule, it could be more difficult for the agent to manipulate the decision, thereby enlarging the set of implementable allocations.

Although requesting information certification with mixed strategies could improve the situation for the decision-maker, it does not eliminate the effect of the limitation on information certification. The original idea of this article is to show that if the agent is ambiguity averse in the sense of maxmin expected utility [Gilboa and Schmeidler, 1989], then requesting information certification with an ambiguous (instead of a pure or mixed) strategy completely eliminates the effect of limited certification. The reason behind this is that an agent who is ambiguity averse in the sense of maxmin expected utility always anticipates the worst case scenario. In the audit example, if the auditor chooses an activity using an ambiguous strategy, an agent who lies about one or more activities gets investigated on one of these activities in the worst case scenario. Therefore, he would act exactly as if all activities were going to be investigated. In this sense, ambiguity saves time on information investigation.

We prove this result in a general model: any allocation rule that is implementable with unlimited certification (with or without ambiguity) is also implementable with limited certification and ambiguity if agents are averse to ambiguity in the sense of maxmin expected utility. We give examples where the converse assertion does not hold, but show that the equivalence holds if there is a single agent and a worst outcome.

Related Literature The topic of mechanisms with certifiable information has been actively investigated in the literature such as in Ben-Porath and Lipman [2012], Bull [2008], Bull and Watson [2007], Deneckere and Severinov [2008], Glazer and Rubinstein [2001, 2004], Green and Laffont [1986], Forges and Koessler [2005], Kartik and Tercieux [2012], Koessler and Perez-

Richet [2014], Sher [2011], Sher and Vohra [2015], Singh and Wittman [2001], Strausz [2016]. These papers, among others, give rise to important results about implementable allocation rules and some of them establish a revelation principle for settings with certifiable information.

Ambiguity has recently been incorporated into mechanism design mainly in two different ways: either through incomplete preferences as in Lopomo et al. [2013] or through the strategic use of ambiguity as in Bose and Renou [2014] and Di Tillio et al. [2012]. In a model of insurance under moral hazard, Lang and Wambach [2013] have shown how uncertainty about an insurer's cost of an audit leads to ambiguity about the probability of an audit, which in turn induces ambiguity-averse agents to undertake less fraud. In this chapter, we allow the designer to use ambiguous communication devices, but unlike the cited papers we consider mechanism design problems with certifiable information.

In section 2.2, we present the model. In section 2.3, we give a characterization of the implementable allocation rules under unlimited information certification. In section 2.4, we show that ambiguity allows to implement all those allocation rules in settings of limited information certification. We also show that if there exists a worst outcome, then any allocation rule that is implementable with ambiguity is also implementable under unlimited information certification. In section 2.5, we discuss the implications of our results for communication games with certifiable information, and their limits in settings with multiple agents or other types of ambiguity averse preferences.

2.2 Model

Consider a mechanism designer with a set of alternatives denoted by A , and an agent¹ who is privately informed about his type $t \in T$. The utility of the agent when his type is $t \in T$ and alternative $a \in A$ is implemented is $u(a; t)$. In this section, the agent evaluates lotteries

¹Section 2.5 discusses which of our results extend to multiple agents.

according to the expected utility.

We assume that the agent can certify some of his information to the designer. This ability to disclose hard evidence is represented by a *certifiability structure* $\mathcal{C} \subseteq 2^T$, where for every $t \in T$ there exists $C \in \mathcal{C}$ such that $t \in C$. We denote by $\mathcal{C}(t) = \{C \in \mathcal{C} : t \in C\}$ the set of events the agent can certify when his actual type is t , and

$$c^*(t) = \bigcap_{C \in \mathcal{C}(t)} C,$$

the intersection of all these events.²

Definition 2.1. The certifiability structure \mathcal{C} satisfies *normality*³ if for every t there exists a certifiable event providing maximal evidence about type t . That is, $c^*(t) \in \mathcal{C}(t)$ for every $t \in T$. The *normalized certifiability structure* of \mathcal{C} is the normal certifiability structure \mathcal{C}^* defined by

$$\mathcal{C}^* \equiv \mathcal{C} \cup \{c^*(t) : t \in T\}.$$
⁴

In the environment $\langle T, \mathcal{C}, A, u \rangle$ the designer would like to implement an *allocation rule* $f : T \rightarrow A$. A *dynamic mechanism* (with unlimited certification) is an extensive form game for the agent with moves of Nature and with perfect recall, that satisfies the following properties: the information structure is compatible with the assumption that the agent privately knows his type; every decision node is either a cheap talk node with an arbitrary number of cheap talk messages (from the agent or from Nature), or a certification node, where the agent of type t

²The certifiability structure defined above could be deduced from any arbitrary message correspondence $\mathcal{M}(t)$, $t \in T$, by letting $\mathcal{C}(t) \equiv \{\mathcal{M}^{-1}(m) : m \in \mathcal{M}(t)\}$. The number of messages in $\mathcal{M}(t)$ may be larger than in $\mathcal{C}(t)$ because several messages in $\mathcal{M}(t)$ may certify the same event in $\mathcal{C}(t)$, but since we consider mechanisms with arbitrary sets of additional cheap talk messages, taking the certifiability structure as a primitive of the model is without loss of generality.

³This terminology is due to [Bull and Watson \[2007\]](#). The normality condition has also been called the full reports condition by [Lipman and Seppi \[1995\]](#) and the minimal closure condition by [Forges and Koessler \[2005\]](#).

⁴We could also define the normalized certifiability structure of \mathcal{C} by the smallest set of events including \mathcal{C} which is closed under intersection without affecting any of the results below.

must certify an event in $\mathcal{C}(t)$; finally, each terminal node is associated to an alternative in A . A *dynamic mechanism with limited certification* is a dynamic mechanism in which there is one and only one certification node in every path through the tree.

The allocation rule $f : T \rightarrow A$ is *implementable* (with unlimited certification) if there exists a dynamic mechanism and an optimal strategy for the agent inducing the allocation rule f . It is *implementable with limited certification* if there exists a dynamic mechanism with limited certification and an optimal strategy for the agent inducing the allocation rule f .

A *direct mechanism* is a mechanism in which each type t simultaneously sends a cheap talk message from the set T and certifies an event in $\mathcal{C}(t)$ to the designer. The strategy of the agent in a direct mechanism is truthful if he sends the cheap talk message t when his type is t .

The allocation rule $f : T \rightarrow A$ is *directly and truthfully implementable* if there exists a direct mechanism and a strategy that is both truthful and optimal for the agent inducing the allocation rule f .

2.3 Implementation without Ambiguity

Forges and Koessler [2005] have shown that an allocation rule is implementable (with unlimited certification) given a certifiability structure \mathcal{C} if and only if it is directly and truthfully implementable given the normalized certifiability structure \mathcal{C}^* , with a mechanism in which each type t is asked to certify the event $c^*(t) \in \mathcal{C}^*(t)$. In particular, if \mathcal{C} is normal, implementation with limited or unlimited certification is equivalent to direct and truthful implementation.

Let

$$M(t) \equiv \{s \in T : c^*(s) \in \mathcal{C}^*(t)\}.$$

The set $M(t)$ represents the set of types of the agent that type t is able to mimic without being observed by the designer if the latter is expecting that each type s certifies $c^*(s)$. Notice that $c^*(s) \in \mathcal{C}^*(t) \Leftrightarrow \mathcal{C}^*(s) \subseteq \mathcal{C}^*(t) \Leftrightarrow \mathcal{C}(s) \subseteq \mathcal{C}(t)$.

For every $S \subseteq T$, let

$$\text{INTIR}(S) = \{x \in \mathbb{R}^S : \exists \sigma(S) \in \Delta(A), x_t \geq u(\sigma(S); t) \forall t \in S\},$$

be the individually rational payoffs for the agent given S . That is, a payoff profile is individually rational given S if there exists a strategy $\sigma(S)$ for the designer that gives a lower payoff to every type in S . Notice that when $S = T$ individual rationality as defined above corresponds to the definition of individual rationality of [Myerson \[1991, p. 282\]](#).

Proposition 2.1. The allocation rule $f : T \rightarrow A$ is implementable (with unlimited certification) if and only if the following conditions are satisfied:

$$u(f(t); t) \geq u(f(s); t), \text{ for every } t \in T \text{ and } s \in M(t); \quad (2.1)$$

$$(u(f(t); t))_{t \in S} \in \text{INTIR}(S), \text{ for every } S \in \mathcal{C}. \quad (2.2)$$

Proof. Directly from the arguments above and the revelation principle for certifiable information in [Forges and Koessler \[2005\]](#). □

The first condition in the proposition guarantees that no agent type has an incentive to deviate to messages sent on the equilibrium path, while the second condition guarantees that the designer can punish observable deviations (i.e., when the agent certifies a set of types S which is not in $\{c^*(t) : t \in T\}$).

Notice that if there is a message certifying no information, i.e., $T \in \mathcal{C}$, then condition (2.2) simplifies to the standard individual rationality condition:

$$\exists \sigma \in \Delta(A) \text{ such that } u(f(t); t) \geq u(\sigma; t), \text{ for every } t \in T.$$

We say that an alternative $\underline{a} \in A$ is a *worst outcome* if $u(\underline{a}; t) \leq u(a; t)$ for every $a \in A$

and $t \in T$. When there exists a worst outcome, condition (2.2) is always satisfied whatever the allocation rule and the certifiability structure.

With limited certification, dynamic mechanisms could implement more allocation rules than direct mechanisms. To see this, consider the following example.⁵

Example 2.1. The set of types is $T = \{t^1, t^2, t^3\}$, the certifiability structure is $\mathcal{C} = \{\{t^1, t^2\}, \{t^2, t^3\}, T\}$ (which is not normal), the set of alternatives is $A = \{a_1, a_2, a_3, a_0\}$, and the agent's utility as a function of the agent's type $t \in T$ and the alternative $a \in A$ is given by

$$u(a; t) = \begin{array}{ccccc} & a_1 & a_2 & a_3 & a_0 \\ t^1 & \boxed{0} & 1 & 1 & -\beta \\ t^2 & 0 & \boxed{0} & 0 & -\beta \\ t^3 & 1 & 1 & \boxed{0} & -\beta \end{array}$$

where $\beta \geq 0$.

Consider the allocation rule $f(t^i) = a_i$, $i = 1, 2, 3$. Clearly, this allocation rule is implementable with unlimited certification: each type certifies the smallest event he is able to certify with the normalized certifiability structure $\mathcal{C}^* = \{\{t^1, t^2\}, \{t^2\}, \{t^2, t^3\}, T\}$, and the worst outcome a_0 is chosen off the equilibrium path (i.e., when the agent certifies the event T).

However it is not directly implementable with limited certification. Indeed, to induce action a_2 with probability one when the agent's type is t^2 , a direct mechanism should induce action a_2 with probability one after the agent has reported that his type is t^2 and has certified the event $\{t^1, t^2\}$ or $\{t^2, t^3\}$. If a_2 is induced after certifying $\{t^1, t^2\}$, then type t^1 has an incentive to mimic type t^2 ; if a_2 is induced after certifying $\{t^2, t^3\}$, then type t^3 has an incentive to mimic type t^2 .

As shown by [Bull and Watson \[2007\]](#), when normality is not satisfied (as in the above exam-

⁵Similar examples appear in [Glazer and Rubinstein \[2001, 2004\]](#), [Bull and Watson \[2007\]](#) and [Bull \[2008\]](#).

ple), the set of implementable allocation rules with limited certification can still be characterized in a canonical way, but with a three-stage mechanism: in the first stage, the agent reports his type to the designer; in the second stage, the designer sends a message to the agent; finally, in the last stage, the agent certifies some information to the designer. With such a dynamic mechanism, the allocation rule f in the previous example is implementable even with limited certification, as long as $\beta \geq 1$, with the following dynamic mechanism: the designer asks the agent who has reported type t^1 to certify $\{t^1, t^2\}$, he asks the agent who has reported type t^3 to certify $\{t^2, t^3\}$, and he asks the agent who has reported type t^2 to certify $\{t^1, t^2\}$ or $\{t^2, t^3\}$ with the same probability. If the agent fails to certify the event asked by the designer, then the worst outcome a_0 is implemented.

2.4 Implementation with Ambiguous Mechanisms

In the previous example, the allocation rule f is directly implementable with unlimited certification, it is not directly implementable with limited certification, but it is implementable with limited certification and a dynamic mechanism. However, an allocation rule which is implementable with unlimited certification is not necessarily implementable with a dynamic mechanism with limited certification. To see that, consider again Example 2.1. The allocation rule $f(t^i) = a_i$, $i = 1, 2, 3$, cannot be implemented with limited certification when $\beta < 1$. Indeed, if the designer asks type t^2 to certify $\{t^1, t^2\}$ with probability $p \geq 1/2$, then type t^1 has an incentive to mimic type t^2 , and if he asks type t^2 to certify $\{t^1, t^2\}$ with probability $p \leq 1/2$, then type t^3 has an incentive to mimic type t^2 .

Assume now that the agent is averse to ambiguity in the sense of maxmin expected utility [Gilboa and Schmeidler, 1989], and that the designer can use ambiguous communication strategies. More precisely, dynamic mechanisms are extended to allow nature (communication by the designer) to choose sets of probability distributions over messages instead of a single

probability distributions over messages. Following the terminology of [Riedel and Sass \[2013\]](#), the designer can use mechanisms in which Nature uses Ellsberg communication strategies (i.e., sets of probability distributions over messages) instead of standard strategies (i.e., probability distributions over messages). Then at each information set, the agent evaluates his local strategy with the minimal expected utility from this strategy, where the minimum is taken over all possible probability distributions compatible with the designer's (ambiguous) communication strategy.⁶ An allocation rule f is *implementable with an ambiguous mechanism* if there exists a dynamic mechanism, an ambiguous communication strategy for the designer and an optimal strategy for the agent inducing f .

To illustrate this extension, consider the previous example and assume that the designer ambiguously asks the agent who reported type t^2 in the first stage to certify either $\{t^1, t^2\}$ or $\{t^2, t^3\}$ according to a *set* of probability distributions $P \subseteq \Delta(\{\{t^1, t^2\}, \{t^2, t^3\}\})$. If $P = \{(p, 1 - p)\}$ we are back to a standard communication strategy for the designer as above. Consider the same mechanism as above, but assume that $P = \Delta(\{\{t^1, t^2\}, \{t^2, t^3\}\})$: the designer asks type t^2 to certify either $\{t^1, t^2\}$ or $\{t^2, t^3\}$ with an imprecise probability, and as before, if the agent fails to certify the event asked by the designer, then the worst outcome a_0 is implemented. With this mechanism, the maxmin expected utility of type t^1 who mimics type t^2 is

$$\min_{p \in [0,1]} (pu(a_2; t^1) + (1 - p)u(a_0; t^1)) = u(a_0; t^1) = -\beta \leq 0,$$

and his maxmin expected utility if he reports his true type to the designer is simply $u(a_1; t^1) = 0$.

Similarly, the maxmin expected utility of type t^3 who mimics type t^2 is

$$\min_{p \in [0,1]} (pu(a_0; t^3) + (1 - p)u(a_2; t^3)) = u(a_0; t^3) = -\beta \leq 0,$$

⁶As in [Bose and Renou \[2014\]](#) we assume prior-by-prior updating (full Bayesian updating), and that the agent is a consistent planner ([Siniscalchi, 2011](#)), i.e., at every information set he maximizes his minimal expected utility at this information set given the strategies he will actually follow.

and his maxmin expected utility if he reports his true type to the designer is $u(a_3; t^3) = 0$. It is therefore optimal for the agent to report his true type: the allocation rule f is implementable with an ambiguous mechanism⁷ and limited certification under the same conditions on β as the implementation condition without ambiguity and with unlimited certification. The next proposition shows that this is general.

Proposition 2.2. If the allocation rule $f : T \rightarrow A$ is implementable (with unlimited certification) then it is also implementable with an ambiguous mechanism with limited certification if the agent is averse to ambiguity in the sense of maxmin expected utility.

Proof. Consider an environment $\langle T, \mathcal{C}, A, u \rangle$ and an allocation rule f which is implementable (with unlimited certification). We construct below a dynamic and ambiguous mechanism together with an optimal strategy for the agent which implements f with limited certification. The mechanism is very similar to the three-stage mechanism of Bull and Watson [2007], with the difference that in the second stage the designer sends a message to the agent using an ambiguous strategy (a *set* of probability distributions over messages) instead of a mixed strategy (a *single* probability distribution over messages).

In the first stage, the agent reports his type with a cheap talk message $s \in T$ to the designer.

In the second stage, after receiving a message $s \in T$ from the agent, the designer sends a message $R(s) \in \mathcal{C}(s)$ to the agent, interpreted as a request to the agent for the certification of the event $R(s)$, according to the ambiguous communication strategy $\Delta(\mathcal{C}(s))$.

In the third stage, the agent certifies the event $R(s)$ requested by the designer in the second stage.

Finally, the alternative $f(s)$ is implemented if the agent has certified the event requested by the designer. Otherwise, if the event $S \in \mathcal{C}$ certified by the agent in the third stage does not

⁷This mechanism is constructed such that the agent faces ambiguity at the first stage (report) and a decision problem under certainty at the last stage (certification). Therefore, consistent plans consist in choosing a report that maximizes the minimal expected utility given the designer's ambiguous communication strategy and his anticipated certification decisions.

correspond to the event $R(s)$ requested by the designer in the second stage, then the alternative is chosen according to the same distribution $\sigma(S)$ used in the original mechanism.

Consider an agent type t . We show below that the agent has no incentive to report another type $s \neq t$ in the first stage, and in the third stage he has no incentive to certify an event different than the one requested by the designer when the agent is able to certify the event requested by the designer.

In the third stage, if the agent is able to certify the event requested by the designer and certifies it then he gets $u(f(t); t)$. If he certifies another event $S \in \mathcal{C}(t)$ he gets $u(\sigma(S); t)$, which is smaller than $u(f(t); t)$ by condition (2.2) in Proposition 2.1. When he can, it is therefore optimal for the agent to certify the event requested by the designer. If the agent cannot certify the event requested by the designer, then it is optimal to certify an event in $\arg \max_{S \in \mathcal{C}(t)} u(\sigma(S); t)$.

In the first stage, if the agent reports a type $s \in T$ where $\mathcal{C}(s) \subseteq \mathcal{C}(t)$, then $s \in M(t)$. So by condition (2.1) in Proposition 2.1 we have $u(f(t); t) \geq u(f(s); t)$, meaning that type t has no incentive to report s . If he reports a type s such that $\mathcal{C}(s) \not\subseteq \mathcal{C}(t)$, then it is possible (according to the ambiguous communication strategy of the designer) that the designer requests the agent to certify an event $R(s) \notin \mathcal{C}(t)$; in this case the expected utility of the agent is $\max_{S \in \mathcal{C}(t)} u(\sigma(S); t)$, which is smaller than $u(f(t); t)$ by condition (2.2) in Proposition 2.1. Hence, the minimal expected utility of the agent (where the minimum is taken over all communication strategies of the designer) is smaller than $u(f(t); t)$, and therefore the agent has no incentive to deviate from truthfully reporting his own type.⁸

□

In general, the converse of the previous proposition does not hold: an allocation rule might be implementable with an ambiguous mechanism (and limited certification) without being im-

⁸Notice that, in general, a deviation to a non-deterministic strategy for an ambiguity averse player may be beneficial while no pure strategy is; however, it is readily observed that the strategy of the agent in the mechanism considered in the proof of the proposition is still optimal if he can use non-deterministic strategies.

plementable (with unlimited certification). To see that, consider the following example.

Example 2.2. The set of types is $T = \{t^1, t^2\}$, the (normal) certifiability structure is $\mathcal{C} = \{\{t^1\}, \{t^2\}, T\}$ and the agent's utility as a function of the agent's type $t \in T$ and the alternative $a \in A = \{a_1, a_2, a_3, a_4\}$ is given by

$$\begin{array}{rcccc}
 & a_1 & a_2 & a_3 & a_4 \\
 u(a; t) = & t^1 & \boxed{0} & 2 & -1 & 2 \\
 & t^2 & 2 & \boxed{0} & 2 & -1
 \end{array}$$

Since the certifiability structure is normal, implementation with limited certification is equivalent to implementation with unlimited certification. From Proposition 2.1 we know that the allocation rule $f(t^i) = a_i$, $i = 1, 2$ is not implementable because

$$(u(f(t^1); t^1), u(f(t^2); t^2)) = (0, 0) \notin \text{INTIR}(T).$$

This allocation rule is however implementable with an ambiguous mechanism inducing the alternatives a_3 and a_4 with the whole set of probability distributions over $\{a_3, a_4\}$ when the agent certifies the event T . Indeed, with such a mechanism, any type t who deviates by not certifying the event $\{t\}$ gets

$$\min_{p \in [0,1]} (pu(a_3; t) + (1 - p)u(a_4; t)) = -1 < u(f(t); t).$$

In the previous example it is interesting to note that the introduction of ambiguity in the choice of alternatives *does* enlarge the set of implementable allocation rules. This contrasts with the case without certifiable information, where ambiguity in the choice of alternatives does not allow to implement more allocation rules than without ambiguity (see, for example, [Bose and Renou, 2014](#)); in particular, without information certification, introducing ambiguity is useless

(for implementing a social choice function) if there is only one agent.

The next proposition shows that what happens in the previous example cannot happen if there exists a worst outcome:

Proposition 2.3. Assume that there exists a worst outcome and that the agent is averse to ambiguity in the sense of maxmin expected utility. If the allocation rule $f : T \rightarrow A$ is implementable with an ambiguous mechanism and limited certification, then it is also implementable (with unlimited certification).

Proof. Consider an environment $\langle T, \mathcal{C}, A, u \rangle$ with a worst outcome \underline{a} , and an allocation rule f which is implementable with an ambiguous mechanism and limited certification. With such a mechanism and the associated optimal strategy of the agent inducing the allocation rule f , the outcome is by definition $f(t)$ with probability one for every type t of the agent. In order to prove the result, we only need to show that f satisfies the conditions of Proposition 2.1. First, modify this mechanism by assigning alternative \underline{a} to every terminal node reached with probability 0 given the agent's strategy. Clearly, with this new mechanism the strategy of the agent is still optimal. Now modify the mechanism further by replacing each certification node by a sequence of successive certification nodes, and modify the strategy of the agent in such a way that he certifies all sets in $\mathcal{C}(t)$ (both along and off the equilibrium path). This is equivalent to asking the agent of type t to certify $c^*(t)$ at each of his certification nodes. The allocation rule f is still implemented with this new mechanism and strategy, and the strategy of the agent is still optimal. Now, if the agent's type is t and he follows the strategy above he gets $f(t)$ with probability one, and if he deviates he gets either $f(s)$ with $s \in M(t)$, or the worst outcome \underline{a} . Since such a deviation should not be profitable, the conditions of Proposition 2.1 are satisfied, and therefore a direct mechanism with unlimited certification can be used to implement the same allocation rule. \square

An important corollary of the two previous propositions is that if there is a worst outcome,

then the set of allocation rules that are implementable with an ambiguous mechanism (and limited or unlimited certification) exactly coincides with the set of implementable allocation rules (with unlimited certification).

2.5 Discussion

Implications for communication games with certifiable information. Our results also have implications for sender-receiver communication games with certifiable information: if there is a fully revealing equilibrium under the certifiability configuration \mathcal{C}^* , then under any certifiability configuration \mathcal{C} whose normalized certifiability configuration is \mathcal{C}^* , there is also a fully revealing equilibrium if communication is bilateral and dynamic and the receiver can use ambiguous communication strategies. Indeed, since full information revelation induces the first best for the receiver, such an ambiguous communication strategy is obviously a best response for the receiver. For instance, in Example 2.1, if the receiver's utility is given by

$$\begin{array}{ccccc}
 & a_1 & a_2 & a_3 & a_0 \\
 t^1 & 1 & -2 & -2 & 0 \\
 t^2 & -2 & 1 & -2 & 0 \\
 t^3 & -2 & -2 & 1 & 0
 \end{array}$$

then the ambiguous mechanism used by the receiver to implement f is a best response to the agent's strategy.

Multiple agents. Proposition 2.2 can be extended to environments with multiple agents, by considering (weak) implementation, and without making specific assumption on agents' utility functions and prior beliefs. Each agent i 's set of types is T_i , $i \in \{1, \dots, n\}$, and they have a common prior $p \in \Delta(T)$, where $T = \times_i T_i$. The utility of agent i is $u_i(a, t)$, for $a \in A$ and $t \in T$.

The individually rational payoffs can be generalized as follows in settings with n agents. For every agent i , for every $S_i \subseteq T_i$, let

$$\text{INTIR}_i(S_i) = \{x \in \mathbb{R}^{S_i} : \exists \sigma(S_i) : T_{-i} \rightarrow \Delta(A), x_t \geq \sum_{t_{-i} \in T_{-i}} p(t_{-i} | t_i) u_i(\sigma(S_i, t_{-i}); t_i, t_{-i}) \forall t_i \in S_i\},$$

be the individually rational payoffs for agent i given S_i . That is, a payoff profile is individually rational given S_i if for every combination of types t_{-i} of agents other than i , there exists a strategy $\sigma(S_i, t_{-i}) \in \Delta(A)$ for the designer that gives a lower interim expected payoff to every type in S_i . Then, Proposition 2.1 is generalized as follows: An allocation rule $f : T \rightarrow A$ is implementable (with unlimited certification) if and only if the following conditions are satisfied for every $i \in N$:

$$\sum_{t_{-i} \in T_{-i}} p(t_{-i} | t_i) u_i(f(t); t) \geq \sum_{t_{-i} \in T_{-i}} p(t_{-i} | t_i) u_i(f(s_i, t_{-i}); t), \text{ for every } t_i \in T_i \text{ and } s_i \in M_i(t_i); \quad (2.3)$$

$$\left(\sum_{t_{-i} \in T_{-i}} p(t_{-i} | t_i) u_i(f(t_i, t_{-i}); t_i, t_{-i}) \right)_{t_i \in S_i} \in \text{INTIR}_i(S_i), \text{ for every } S_i \in \mathcal{C}_i. \quad (2.4)$$

In addition, we say that f is implementable with ambiguous mechanisms if there exists a dynamic mechanism, an ambiguous communication strategy of the designer and a profile of strategies for the agents such that, at each information set, each player maximizes his minimum expected utility given others' strategies and his own strategy at the other information sets. The proof of Proposition 2.2 uses the same mechanism as in the single-agent case: if every agent $j \neq i$ is truthful in the first stage and certifies the event requested by the designer in the third stage, then agent i faces the same decision problem as in the single-agent case except that his utility $u_i(\cdot; t)$ is replaced by his interim expected utility $\sum_{t_{-i} \in T_{-i}} p(t_{-i} | t_i) u_i(\cdot; t_i, t_{-i})$.

However, Proposition 2.3 does not extend when there are multiple agents; this has already been observed in the literature when information is not certifiable (see, in particular, [Bose and](#)

Renou, 2014); the reason is that ambiguous communication could modify agents' perceptions about other agents' types, and could therefore relax their incentive constraints. A simple example is provided below:

Example 2.3. There are two agents, and each agent $i \in \{1, 2\}$ has two possible types, $t_i \in \{t^1, t^2\}$, that are uniformly distributed. Information is not certifiable. The utilities of the agents as a function of the type profile and alternative a_1 or a_2 are given by:

a_1	t^1	t^2	a_2	t^1	t^2
	$0, 1$	$2, 0$		$-1, 0$	$0, 1$
	$0, 0$	$0, 1$		$0, 1$	$0, 0$

Clearly, the allocation rule

$$f(t^1, t^2) = \begin{cases} a_1 & \text{if } t^1 = t^2 \\ a_2 & \text{if } t^1 \neq t^2, \end{cases}$$

is not implementable with any unambiguous mechanism: when agent 1's type is t^1 and he reports type t^2 to the designer, his expected payoff is $(1/2)u_1(a_2; t^1, t^1) + (1/2)u_1(a_1; t^1, t^2) = -(1/2)1 + (1/2)2$ instead of $(1/2)u_1(a_1; t^1, t^1) + (1/2)u_1(a_2; t^1, t^2) = 0$ (adding a worse outcome would not modify this conclusion). However, if agent 2 first uses an ambiguous communication strategy before agent 1 reveals his type to the designer, then agent 1 has no incentive to deviate anymore if the communication strategy of agent 2 to agent 1 is sufficiently imprecise; for example, if agent 2 sends messages m_1 and m_2 to agent 1 according to the whole set of probability distributions $\Delta(m_1, m_2)$ whatever his type, then after any such message the set of beliefs of agent 1 about agent 2's type would be whole set, and the maxmin utility derived from a deviation of agent 1 when his type is t^1 would be $\min_{p \in [0, 1]} (-p + 2p) = -1 < 0$.⁹

⁹Equivalently, agent 2 can first reveal his type to the designer, who then sends messages m_1 and m_2 to

Other ambiguity averse preferences. Throughout the chapter, we assumed that agents are ambiguity averse in the sense of maxmin expected utility. Our main result is that ambiguous mechanisms and limited certification allow to implement all allocation rules that are implementable with unlimited certification. In short, ambiguity completely eliminates the effect of certification limitation. This result clearly relies on the extreme ambiguity aversion of MEU, and it does not necessarily hold for less extreme preferences. For example, consider the case of α -maxmin expected utility (α -MEU). An agent who has α -MEU preferences maximizes the weighted average of the minimal and maximal expected utilities with α ($1 - \alpha$, respectively) being the weight of the minimum (maximum, respectively). We can easily check that if an allocation rule is implementable with the three-stage ambiguous mechanisms proposed in the proof of Proposition 2.2 for a given α , then it is also implementable with the same mechanism for any α' greater than α . Clearly, the converse does not hold. For instance, under α -MEU, the ambiguous mechanism with limited certification used in Example 2.1 will implement the allocation rule $f(t^i) = a_i, i = 1, 2, 3$ only if $\alpha \geq \frac{1}{1+\beta}$. Therefore, when α is not large enough, ambiguity does not completely eliminate the effect of certification limitation, and implementable allocation rules will depend more finely on the environment and on the agent's sensitivity to ambiguity.

agent 1 according to the previous ambiguous communication strategy. Thus, the mechanism described here is consistent with our definition of ambiguous mechanisms.

Chapter 3

Seeking advice

3.1 Introduction

Decision makers have the difficult task of choosing among alternative policies and projects without direct access to the relevant information. Acquiring this information may be difficult and very costly or even impossible for the decision maker. One way to solve this issue is to simply rely on the advice of an informed party. Obviously, if they have identical preferences then this would completely eliminate the problem. Unfortunately, these preferences are usually misaligned due to a variety of reasons. For example, in a firm, an informed employee may be biased in favor of a certain project because it might benefit his career more than other projects. When information is *hard*, the decision maker can still learn it by consulting a biased agent if his bias is known. Well established results of unraveling (see for example [Milgrom \[1981\]](#) and [Grossman \[1981\]](#)) may apply in this setting allowing the elicitation of all information despite the bias. However, if the bias is unknown the previous argument could fail and it might not be sufficient to consult one agent.

In this chapter, the decision maker takes advantage of the presence of multiple informed agents with unknown biases in order to learn more information than she would when facing

only one agent. By consulting more agents, the decision maker increases her chances of learning information but agents' silence can be informative enough to make her stop seeking advice and choose an action without acquiring the missing information.

In organizations for example, information is usually dispersed in the hierarchy due to inherent structural reasons and to specialization (see for example [Marschak and Radner \[1972\]](#), [Walker \[1998\]](#) and [Dessein \[2002\]](#)). For instance, the most informed members about the suitability of a given project are not those who have to decide whether to invest in the project but rather those who would have to implement it subsequently. Moreover, informed agents might not have the same goals as decision makers who have to rely on their advice but have the possibility of consulting more than one of them.

We consider a model with two possible states and two alternative actions for the decision maker. Choosing the right action in this context amounts to guessing the true state. There are two types of agents: each agent is biased in favor of one of the two actions regardless of the state. Agents know the true state and, if consulted, can either reveal it or remain silent. Information is hard and agents cannot lie. The decision maker chooses at each period whether to consult an agent or stop and select an action. In the main version of the model, if a consulted agent chooses to reveal the state, the decision maker processes the presented information at a cost which represents the effort and time she has to allocate (to read a report or attend a meeting, for example). In Section 3.4.2, we consider two other versions where the decision maker incurs a constant cost per consultation or discounts future payoff. We show that we obtain essentially the same results as in the main version.

In equilibrium, the decision maker starts seeking advice only if her uncertainty about the state of the world is high enough. Her goal is to reduce her uncertainty before making a decision. She keeps consulting agents until she either learns all the information -by encountering an agent who is willing to reveal it- or her uncertainty is reduced to a level that does not justify consulting more agents. It is important to note that learning does not happen only when the agent reveals

information. If agents types (biases) are not equiprobable, encountering an agent who withholds information is also informative: the decision maker updates her belief by putting a higher weight on the state where the majority type withholds information. Therefore, the decision maker can stop seeking advice after consulting a certain number of agents who did not reveal the state. In this case, she can also make a choice that differs from the one she would have made based on her prior beliefs.

If the decision maker does not learn the state by the time she stops seeking advice, it is possible that she makes a bad decision. This can happen if she encounters only agents from the minority (in terms of bias distribution) and they withhold information. Agents of the minority are therefore able to influence the decision maker so that she chooses their favorite action when she should not, provided that she does not encounter any agent of the majority.

We observe that the decision maker requires a lower level of uncertainty about the state before choosing an action if the processing cost or uncertainty about agents preferences decreases. Consider a level of uncertainty about the state such that she is indifferent between choosing an action and seeking more advice. If the processing cost decreases she would strictly prefer consulting more agents as she expects a lower loss. The same applies if uncertainty about agents preferences decreases because she expects a higher gain as agents silence becomes more informative.

When the cost decreases, the decision maker consults more agents before stopping to choose an action. However, she generically consults less agents when she is less uncertain about their bias. Although she requires a lower level of uncertainty about the state, the increase of agents silence informativeness has a greater effect on the number of consultations.

We characterize the maximum number of consulted agents in equilibrium and we show that it increases indefinitely when the processing cost vanishes or uncertainty about agents preferences approaches its maximum. However, the average number of consultations always has a finite limit.

In section 3.5, we consider an alternative model with a richer state space that allows agents to reveal more or less precise information about the state. The state space is an interval and the message can be any interval containing the true state. We characterize the unique equilibrium outcome when the decision maker can consult only one agent and we show that it is the same as the outcome that would obtain if the agent can only reveal the state or remain silent. In other words, if sequential consultation is not allowed, the richness of the message space has no impact on the outcome. When sequential consultation is allowed, we find an equilibrium where the decision maker consults exactly one agent and uses the possibility of further consultation as a threat to extract more precise information. The consulted agent sends a message that is just precise enough to make the decision maker stop which allows him to influence her decision in his favor to a certain extent. We show the existence of such an equilibrium for a low enough consultation cost and equiprobable types.

Related Literature The issue of dispersed information within organizations has been studied in the literature and it is often better to delegate decision rights to informed agents rather than try to elicit information through communication (see [Aghion and Tirole \[1997\]](#) and [Dessein \[2002\]](#) for example). We depart from this literature by considering a setting where it is best to keep decision rights because the divergence between the preferences of the decision maker and the agents is extreme.

This chapter is also related to the literature on strategic communication. Since the work of [Crawford and Sobel \[1982\]](#) on cheap talk games and of [Milgrom \[1981\]](#) on persuasion games, an extensive literature has been developed on both types of sender-receiver games. In the present work, the sender is the agent, the receiver is the decision maker and the information is verifiable as in persuasion games. In the cheap talk literature, [Le Quement \[2016\]](#) models the interaction between consumers and product reviewers through sequential consultation of informed senders who use cheap talk messages and who either share the receiver's preferences

or are biased toward high actions. Our contribution is to study sequential consultation in the context of hard information and explore its effect on learning and decision making.

Uncertainty over the agent's preferences, formulated similarly to [Wolinsky \[2003\]](#), is an important feature of our model: the decision maker has two alternative actions (0 or 1) and the agent either wants action 1 (type H) or action 0 (type L). This uncertainty causes the failure of the unravelling argument which was presented by [Milgrom \[1981\]](#) and [Grossman \[1981\]](#) and states that full revelation occurs in equilibrium because the decision maker assumes the lowest possible value (of the state variable) that is consistent with reported information. The result holds precisely because the sender's goal is to persuade the decision maker that this value is as high as possible, which is only true for type H in our model while type L has the opposite goal. Other authors such as [Dziuda \[2011\]](#) introduce uncertainty over the level of strategic sophistication of the sender by having a proportion of non-strategic (or honest) senders who always report all their private information. In the cheap talk literature, some authors studied models where the bias is uncertain (see for example [Li \[2005\]](#) and [Dimitrakas and Sarafidis \[2005\]](#)).

The agent's preferences are independent of his private information. [Matthews and Postlewaite \[1985\]](#), [Milgrom and Roberts \[1986\]](#), [Shin \[1994\]](#), [Glazer and Rubinstein \[2004\]](#) and [Sher \[2011\]](#) among others also consider state-independent preferences in models of verifiable messages games. However, some papers in persuasion games literature consider state-dependent preferences, see for example [Seidmann and Winter \[1997\]](#), [Forges and Koessler \[2005\]](#), [Giovannoni and Seidmann \[2007\]](#) and [Mathis \[2008\]](#).

3.2 Model

A decision maker (DM) has to choose an action a in $A = \{0, 1\}$. Her payoff depends on her action and on the state of the world s which is in $S = \{0, 1\}$. The probability of state 1 is

denoted π . *DM* does not observe s but she can sequentially consult informed agents in the organization. Information is hard and the agent's messages are verifiable (or certified) in the following sense: if the state is s , then he can certify s or remain silent.

DM wants to choose the action that matches the state. Her utility function $u_D : A \times S \rightarrow \mathbb{R}$ is given by:

$$u_D(a, s) = \mathbb{1}_{\{a=s\}}.$$

Let $a^*(\pi) = \mathbb{1}_{\{\pi \geq \frac{1}{2}\}}$ denote *DM*'s optimal action when she believes that the state is 1 with probability π .

There is uncertainty over agents' preferences which are independent of the state. Similarly to [Wolinsky \[2003\]](#), each agent can be of type L or H with p denoting the probability of type L . Type L strictly prefers action 0, while type H strictly prefers action 1. The agent's payoff is 1 if his preferred action is chosen and $b < 1$ if the other action is chosen.

Before the game starts, *Nature* draws the state s according to the prior distribution. At the beginning of each period t in \mathbb{N} , *DM* decides whether to seek advice or choose an action. If she chooses to seek advice, *Nature* draws an informed agent (whose type is L with probability p), who either discloses s or remains silent. The game ends when *DM* chooses an action.

Revelation is costless for the agent. In the main model, *DM* incurs the cost of processing the information denoted c if an agent reveals the state but bears no cost if he is silent. In section 3.4.2, we consider variants of the model where the *DM* incurs a consultation cost or discounts future payoff and we show that we obtain similar results. In section 3.4.3, we examine the case where *DM* can choose to avoid the processing cost by learning that the agent wants to reveal the state without receiving the message.

3.3 Equilibrium analysis

In this section, we study the Perfect Bayesian Equilibria (PBE) of this game. The nature of the equilibrium depends on agent's time preferences. If b is positive, agents prefer immediate decisions: conditional on knowing the action to be chosen at the end, both types prefer DM to implement it immediately. If b is negative, each type prefers an immediate decision conditional on his favorite action being chosen but would rather cause DM to delay her decision otherwise.

We assume without loss of generality that the decision maker chooses to stop whenever indifferent between stopping and continuing to seek advice.

Remark 3.1. An agent's equilibrium strategy depends only on b and his type. In equilibrium, type L (respectively, H) reveals information in state 0 (respectively, 1) regardless of b because it guarantees that DM chooses his preferred action immediately. Given agents time preferences and b non zero¹, type L (respectively, H) is not indifferent between revealing the state and remaining silent in state 1 (respectively, 0). Consequently, agents' equilibrium strategies are pure.

Lemma 3.1. *There exists no PBE such that only one type withholds information in one state (type L in state 1 or type H in state 0).*

Proof. Consider a PBE such that type L withholds information in state 1 while type H reveals information in state 0. In that case, the decision maker would learn that the state is 1 when the agent remains silent and therefore she would choose action 1. This would give type H incentive to deviate and remain silent in state 0. □

Lemma 3.1 ensures that there are two types of PBE:

¹If b is zero, agents are indifferent between an immediate and a delayed decision conditional on their preferred action not being chosen. Consequently, we obtain multiple equilibria, but they are of the same nature as those of positive and negative b . For that reason, we discard this case.

1. Full revelation PBE: both types reveal information in both states.
2. Partial revelation PBE: type L (respectively, H) reveals information if and only if the state is 0 (respectively, 1).

We focus mainly on the setting where agents prefer to delay the bad action and biases are not equiprobable. In this case, we observe learning in equilibrium (partial revelation PBE) and a possibility for the minority to influence DM by withholding information. But first, we analyze the situations where agents prefer immediate decisions or biases are equiprobable.

3.3.1 Agents prefer immediate decisions

Assume that b is positive so that agents prefer immediate decisions whether DM implements their favorite action or not.

Proposition 3.1. *If agents strictly prefer immediate decisions then there exists a unique PBE. In equilibrium, both types reveal information in both states (full revelation). DM consults at most one agent:*

- if $c \geq \frac{1}{2}$ she chooses action $a^*(\pi)$ without seeking advice.
- if $c < \frac{1}{2}$:
 - if $c < \pi < 1 - c$, she consults exactly one agent, learns the state and chooses the right action.
 - otherwise, she chooses action $a^*(\pi)$ without seeking advice.
- Off equilibrium path, i.e. if the agent reveals no information, she continues seeking advice.

Proof. First, we check that this strategy profile is indeed a PBE. Given full revelation, the decision maker is willing to consult an agent if and only if her expected payoff $(1 - c)$ is greater

than her expected payoff if she chooses $a^*(\pi)$ which is $\max\{\pi, 1 - \pi\}$. Given the decision maker's strategy, full revelation is the agent's best response because of the strict preference for immediate decisions and the fact that the decision maker continues seeking advice until she learns the state.

Second, we prove that this is the only PBE. In a full revelation PBE, the decision maker has to continue seeking advice off equilibrium path because otherwise at least one type in one state would deviate to withholding information. In a partial revelation PBE, if $p = \frac{1}{2}$, DM's belief would not change after encountering a silent agent and she would continue seeking advice until she learns the state, which gives incentive to agents to deviate to full revelation given their preference for immediate decisions. If $p \neq \frac{1}{2}$, DM's belief would change after encountering a silent agent but this belief change can potentially benefit only one of the two types withholding information which gives the other type incentive to deviate and reveal the state. \square

Since both types prefer immediate decisions in both states, we obtain full revelation in equilibrium when the decision maker seeks advice. This is case if the processing cost is low enough ($c < \frac{1}{2}$) and uncertainty about the state is high enough ($c < \pi < 1 - c$).

3.3.2 Agents prefer delayed bad outcomes

If b is negative, agents want the decision to be delayed if it is not their favorite action. The equilibrium behavior changes so that each type reveals the state only if the corresponding action is good. Given these preferences, there can be no full revelation PBE.

Proposition 3.2. *If $p = \frac{1}{2}$ and agents strictly prefer delayed bad outcomes then there exists a partial revelation PBE and it is unique. In equilibrium, type L (respectively H) reveals the information only in state 0 (respectively 1). DM's equilibrium strategy is:*

- if $c \geq \frac{1}{2}$ she chooses action $a^*(\pi)$ without seeking advice.

- if $c < \frac{1}{2}$:
 - if $c < \pi < 1 - c$, she seeks advice until she learns the state and then chooses the right action.
 - otherwise, she chooses action $a^*(\pi)$ without seeking advice.

Proof. In a partial revelation PBE, DM 's belief remains unchanged when she encounters a silent agent because $p = \frac{1}{2}$. Therefore, for every belief π , she would either choose $a^*(\pi)$ without seeking advice or seek advice until she learns the state. Her payoff if she chooses $a^*(\pi)$ without seeking advice is $\max\{\pi, 1 - \pi\}$. Her payoff if she seeks advice until she learns the state is $1 - c$. In order to conclude, we only need to compare the two payoffs. \square

Note that although strategies are different, the outcome in this case is similar to that of the full revelation equilibrium of Proposition 3.1: DM learns the state and chooses the right action if the prior is in $(c, 1 - c)$ and she chooses the optimal action based on her prior without consultation otherwise. The only difference is the timing of her decision as she might consult more than one agent in this setting.

If the cost c is low enough and uncertainty about the state is high enough, DM would continue to seek advice indefinitely until she learns the state. The reason is that withholding information is not informative at all because $p = \frac{1}{2}$. If p is different from $\frac{1}{2}$, agent's silence becomes informative. Let $\pi_0 = \pi$ and π_t denote DM 's (updated) belief that the state is 1 after encountering silent agents for t periods:

$$\pi_t = \frac{p^t \pi}{p^t \pi + (1 - p)^t (1 - \pi)}.$$

Proposition 3.3. *If $p > \frac{1}{2}$ and agents strictly prefer delayed bad outcomes then there exists a partial revelation PBE and it is unique. In equilibrium, type L (respectively H) reveals the information only in state 0 (respectively 1). DM 's equilibrium strategy is:*

- if $c \geq p$ she chooses action $a^*(\pi)$ without seeking advice.
- if $c < p$, there exist two beliefs $\underline{\pi}$ and $\bar{\pi}$ such that $\underline{\pi} < \frac{1}{2} < \bar{\pi} = \frac{p(1-c)}{p(1-c)+(1-p)c}$ and at each period t :
 - if $\underline{\pi} < \pi_t < \bar{\pi}$, she seeks advice.
 - otherwise, she stops and chooses action $a^*(\pi_t)$.

Proof. See Appendix. □

Note that if $p < \frac{1}{2}$, we obtain an analogous result by symmetry (by relabeling the states, for example). In the remainder of the section, we focus on the case where $p > \frac{1}{2}$ without loss of generality. Under the assumptions of Proposition 3.3, the decision maker learns information on equilibrium path both from agents who reveal information and those who remain silent. Her belief evolves over time and it is possible that she stops seeking advice after encountering silent agents only.

More specifically, if c is smaller than p and DM 's prior belief $\pi_0 = \pi$ is in $(\underline{\pi}, \bar{\pi})$, she starts seeking advice. If the agent is silent she updates her belief to $\pi_1 = \frac{p\pi_0}{p\pi_0+(1-p)(1-\pi_0)}$ which is greater than π_0 (because $p > \frac{1}{2}$). If π_1 is below $\bar{\pi}$, she seeks advice again. This process is repeated until she either learns the state or arrives at a belief π_t above $\bar{\pi}$ after encountering silent agents for t periods. If DM learns the state, she chooses the right action and her payoff is $1 - c$. However, if she does not learn the state and stops at a belief over $\bar{\pi}$ she chooses action $a = 1$. Therefore, if $s = 1$ and DM encounters only agents of type L (the majority type for $p > \frac{1}{2}$), she still chooses the right action at the end. But if $s = 0$ and she encounters only agents of type H (the minority type), she ends up making a bad decision.

Learning on the equilibrium path makes it possible for agents of the minority type to mislead DM to make the wrong decision, by remaining silent. A major difference between this case and the previous one ($p = \frac{1}{2}$) is the fact that DM never continues to seek advice indefinitely.

For any given p and c , there exists a maximal number of consultations in equilibrium.

Let $\tilde{\pi}_n = \frac{(1-p)^{n-1}(1-c)}{(1-p)^{n-1}(1-c)+p^{n-1}c}$ for n in \mathbb{N} . This sequence is defined so that if DM seeks advice at belief $\tilde{\pi}_n$ and the agent is silent, the updated belief after consultation is $\tilde{\pi}_{n-1}$. Moreover, $\tilde{\pi}_0$ coincides with $\bar{\pi}$. Therefore, if π is in $(\underline{\pi}, \bar{\pi}) \cap [\tilde{\pi}_n, \tilde{\pi}_{n-1})$ (for non zero n) then DM seeks advice for at most n periods: if she does not learn the state by the n -th consultation, her posterior belief is above $\bar{\pi}$, therefore she stops seeking advice and chooses action $a = 1$. The expected payoff of the decision maker at a belief π in $(\underline{\pi}, \bar{\pi}) \cap [\tilde{\pi}_n, \tilde{\pi}_{n-1})$ is:

$$v_n(\pi, p, c) = \pi(p^n + (1 - p^n)(1 - c)) + (1 - \pi)(1 - (1 - p)^n)(1 - c).$$

By rearranging the terms, we get $v_n(\pi, p, c) = a_n(p, c)\pi + b_n(p, c)$ where

$$a_n(p, c) = p^n c + (1 - p)^n (1 - c)$$

$$b_n(p, c) = (1 - (1 - p)^n)(1 - c).$$

For every n , the payoff functions v_n and v_{n+1} coincide at $\tilde{\pi}_n$ (ensuring the continuity of DM 's expected payoff). Furthermore, a_n is strictly positive for all n and $a_{n+1}(p, c)$ is strictly smaller than $a_n(p, c)$ for all c and all p in $(\frac{1}{2}, 1)$. These observations allow us to define DM 's expected payoff of seeking advice as $v(\pi, p, c) = \max_{n \in \mathbb{N}} v_n(\pi, p, c)$ (see Figure 3.1) .

Proposition 3.4. *Under the assumptions of Proposition 3.3 and $c < p$, we have:*

- $\bar{\pi}$ is increasing in p and decreasing in c .
- $\underline{\pi}$ is decreasing in p and increasing in c .

Proof. We have $\bar{\pi} = \frac{p(1-c)}{p(1-c)+(1-p)c}$ and therefore $\bar{\pi}$ is increasing in p and decreasing in c . In order to determine how $\underline{\pi}$ varies with the parameters, we use the variations of DM 's expected payoff.

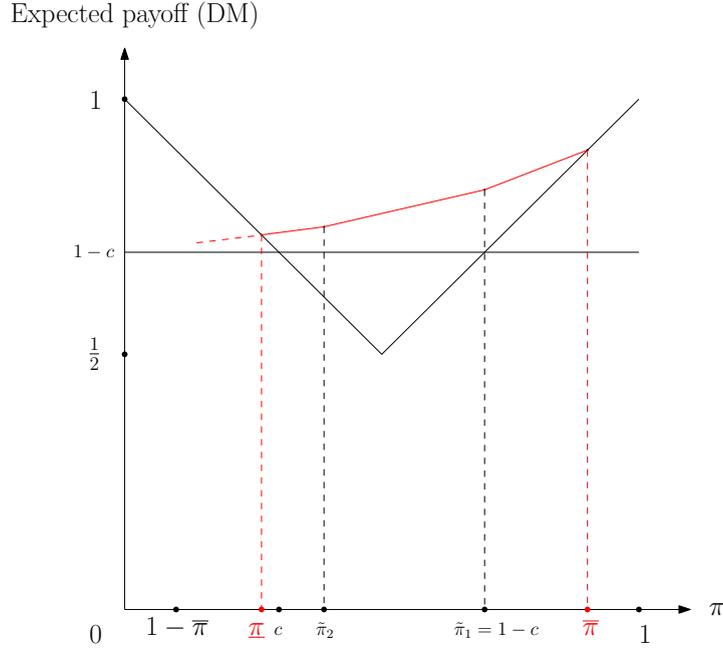


Figure 3.1: DM 's expected payoff (for some fixed p and c). The function v is represented in red.

Note that for all n , v_n is decreasing in c and increasing in p . Therefore, v is also decreasing in c and increasing in p . For any given p and c , v is continuous, strictly increasing and piecewise linear in π . In addition, $\underline{\pi}(p, c)$ is the unique belief π for which $v(\pi, p, c) = 1 - \pi$. Therefore, an increase in p (or a decrease in c) induces a decrease in $\underline{\pi}(p, c)$. \square

Proposition 3.4 asserts that the interval of beliefs such that DM seeks advice is enlarged (in the sense of inclusion) if the cost of processing information or the uncertainty about the agent's type decreases. When c decreases, DM has a larger incentive to seek advice. When p increases, the uncertainty about types decreases and agent's silence is seen as more informative which makes it more interesting to seek advice. In fact, if the true state is 0, an increase in p leads to an increase of the probability of learning the state and a decrease of the probability of wrongfully choosing action 1. If the true state is 1, an increase in p leads to an increase of the probability of not learning the state and choosing the action 1 without incurring the cost of

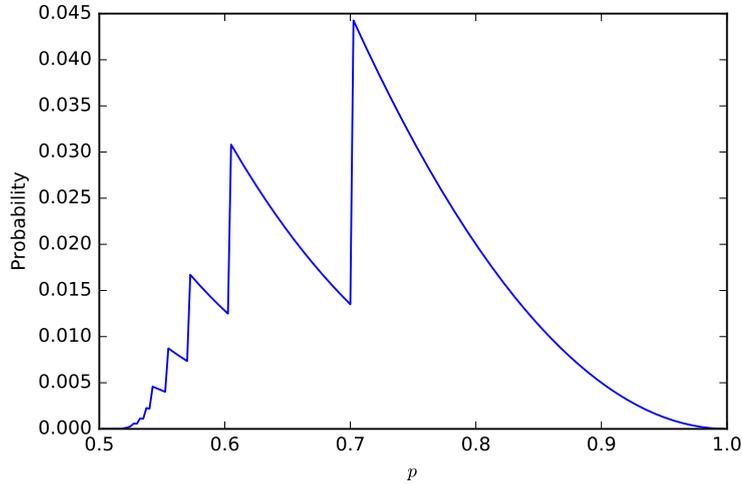


Figure 3.2: Probability of a bad decision for $\pi = 0.5$ and $c = 0.3$

information processing. In other words, the decrease of uncertainty is beneficial in both states.

Probability of bad decisions If DM 's prior belief is in $(\underline{\pi}, \bar{\pi}) \cap [\tilde{\pi}_n, \tilde{\pi}_{n-1})$, she consults at most n agents before choosing an action. Therefore, her final decision is:

- In state 1: action 1 with certainty.
- In state 0:
 - action 1 if she encounters n silent agents.
 - action 0 otherwise.

DM chooses the wrong action if and only if the state is 0 and she encounters n (silent) agents of type H . Note that only the minority (type H for $p > \frac{1}{2}$) benefits from DM 's mistake.

The probability of this mistake is $(1-\pi)(1-p)^n$. In order to understand how this probability changes when the parameters p and c change, we need to determine the way they affect n . Recall that $\tilde{\pi}_n = \frac{(1-p)^{n-1}(1-c)}{(1-p)^{n-1}(1-c)+p^{n-1}c}$. It is readily verifiable that $\tilde{\pi}_n$ is decreasing in c and p for n strictly larger than 1. We also know that $\tilde{\pi}_0 = \bar{\pi}$ is increasing in p and decreasing in c (Proposition 3.4).

Finally, $\tilde{\pi}_1 = 1 - c$ is independent of p and decreasing in c . As a consequence, starting with given parameters (p, c) and a belief π in $(\underline{\pi}, \bar{\pi}) \cap [\tilde{\pi}_n, \tilde{\pi}_{n-1})$ for some n , we can say that if p or c increases then the maximum number of consultations n at π weakly decreases.

Note that the probability of a bad decision depends on c only through n . However, it depends on p through the effect on the distribution of types and also through n . If the cost c decreases, this probability decreases because of the induced increase of the maximum number of consultations n . A small increase in p (so that π remains in $[\tilde{\pi}_n, \tilde{\pi}_{n-1})$) reduces this probability as the proportion of type H agents is reduced. However, the effect of a large increase in p is not clear due to countervailing effect on the maximum number of consultations n : when p increases, DM spends less time seeking advice (i.e. n decreases) which can increase her probability of making a bad decision. Figure 3.2 illustrates this idea for $\pi = \frac{1}{2}$ and $c = \frac{3}{10}$. The points where the mistake probability jumps up are the values of p where the maximum number of consultations (at this prior) decreases by 1. Between the points of discontinuity, the maximum number of consultations is constant and the probability decreases when p increases as explained above.

As stated in the proof of Proposition 3.4, DM 's expected payoff increases with p . This holds regardless of its impact on the probability of bad decisions. The reason is that DM always benefits from an increase of p in state 1 through two effects. First, it increases the probability of choosing action 1 without learning the state and incurring the processing cost. Second, if the change in p is large enough to decrease the maximum number of consultations, then this probability is even larger. In fact, it jumps up at the same discontinuity points as the probability of bad decisions.

We already established that an increase of the probability of a bad decision benefits the minority (type H) and hurts the majority (type L)². The proportion of minority agents is $(1 - p)$. As we can see in Figure 3.2, the minority would prefer this proportion to be large

²This is true if we ignore the effect of the timing of the decision and only consider the decision itself.

enough so that DM has a high probability of only encountering agents of their type, but not too large that she consults too many agents.

Proposition 3.5. *Under the assumptions of Proposition 3.3 and $c < p$, we have*

$$1 - \bar{\pi} < \underline{\pi} < c.$$

Proof. For given p and c , v is strictly increasing in π therefore $v(\underline{\pi}) < v(\bar{\pi})$. We also know that $v(\underline{\pi}) = 1 - \underline{\pi}$ and $v(\bar{\pi}) = \bar{\pi}$, which implies that $1 - \bar{\pi} < \underline{\pi}$.

In order to prove the second inequality, we show that v is above $1 - c$ for all beliefs π . For any given n , we have

$$v_n(\pi, p, c) - (1 - c) = \pi p^n c - (1 - \pi)(1 - p)^n(1 - c)$$

which is strictly positive for all π above $\tilde{\pi}_n$ (and therefore for all beliefs where it coincides with v). To conclude, apply this inequality to the particular case of $\underline{\pi}$:

$$v(\underline{\pi}) = 1 - \underline{\pi} > 1 - c.$$

□

Remark 3.2. DM 's expected payoff is strictly larger than $1 - c$, which is her expected payoff from seeking advice if $p = \frac{1}{2}$. The difference is due to the additional information carried by agents silence when $p \neq \frac{1}{2}$.

Proposition 3.5 gives upper and lower bounds to $\underline{\pi}$. Given that $\bar{\pi}$ converges to $1 - c$ when p decreases toward $\frac{1}{2}$, we deduce that $\underline{\pi}$ converges to c in that case. These bounds also allow us to approximate the maximum number of consultations.

Corollary 3.1. *Under the assumptions of Proposition 3.3 and $c < p$, the maximum number of consultations n^* is either n_c or $n_c + 1$, where*

$$n_c = \left\lceil 1 + 2 \frac{\ln \frac{1-c}{c}}{\ln \frac{p}{1-p}} \right\rceil.$$

Proof. See Appendix. □

n_c is the maximum number of consultations at belief $\pi = c$. In equilibrium, if DM 's prior belief is $1 - \bar{\pi}$ she chooses action 0 without consulting any agent. But if she were to seek advice at that prior, her updated belief after encountering one silent agent would be c . Therefore, the maximum number of consultations at any belief between $\underline{\pi}$ and c is either n_c or $n_c + 1$ (in particular, $\underline{\pi}$ by Proposition 3.5). Note that n_c increases if the processing cost c decreases or the uncertainty about agents type decreases (i.e. p increases). In particular, if c decreases toward 0 or p decreases toward $\frac{1}{2}$, n_c (and as a consequence n^*) goes to infinity. That is, DM is willing to seek advice longer as the cost vanishes or the uncertainty about types increases.

Combining this observation with the fact that $\underline{\pi}$ and $\bar{\pi}$ converge respectively to c and $1 - c$, we obtain the equilibrium of Proposition 3.2 at the limit when p decreases toward $\frac{1}{2}$.

Average number of consulted agents The average number of consulted agents for a prior belief π in $(\underline{\pi}, \bar{\pi}) \cap [\tilde{\pi}_n, \tilde{\pi}_{n-1})$ is:

$$n_{c,p}(\pi) = \sum_{k=1}^{n-1} k(\pi(1-p)p^{k-1} + (1-\pi)p(1-p)^{k-1}) + n(\pi p^{n-1} + (1-\pi)(1-p)^{n-1}).$$

It can be shown that

$$n_{c,p}(\pi) = \pi \frac{1-p^n}{1-p} + (1-\pi) \frac{1-(1-p)^n}{p}.$$

When the cost c vanishes, the maximum number of consulted agents (n) goes to infinity but the average number of consulted agents $n_{c,p}(\pi)$ has a finite limit value $n_p(\pi) = \frac{\pi}{1-p} + \frac{1-\pi}{p}$. Similarly, when p goes to $\frac{1}{2}$, n diverges³ but $n_{c,p}(\pi)$ limit value is 2 which is the average number of consulted agents in the equilibrium of Proposition 3.2.

3.4 Extensions

3.4.1 Commitment

In this section, we show that commitment cannot help the decision maker improve her expected payoff. In other words, she would not gain from ignoring the information she learns from agents silence or making dynamically inconsistent decisions.

If DM has commitment power, she can choose (conditional on her prior belief) a maximum number of consultations in advance and the action she would implement if she does not learn the state, which we subsequently call the *default action*. Assuming we know which action that is, we can determine the optimal maximum number of consultations and the resulting expected payoff. Then, we compare the obtained payoffs for the two possible default actions with the equilibrium payoff and conclude by observing that commitment does not help DM .

Although DM updates her belief along the equilibrium path and makes dynamically consistent decisions, her equilibrium strategy can be seen as the combination of a maximum number of consultations and a default action for each prior, where the default action is 1. The previous analysis shows that $v(\pi, p, c) = \max_{n \in \mathbb{N}} v_n(\pi, p, c)$, therefore the equilibrium strategy specifies the optimal maximum number of consultations for the default action 1.

We now need to determine the optimal maximum number of consultations for the default action 0. Let $w_n(\pi, p, c)$ denote the expected payoff of seeking advice for at most n periods and

³Assuming $c > \frac{1}{2}$ and π is in $(c, 1 - c)$ in order to guarantee DM seeks advice at prior π in the limit.

choosing action 0 if all n agents are silent. We can write $w_n(\pi, p, c) = \alpha_n(p, c)\pi + \beta_n(p, c)$ where

$$\begin{aligned}\alpha_n(p, c) &= -((1-p)^n c + p^n (1-c)) \\ \beta_n(p, c) &= (1-c) + (1-p)^n c.\end{aligned}$$

For every n , the payoff functions w_n and w_{n+1} coincide at $\hat{\pi}_n$ defined as:

$$\hat{\pi}_n = \frac{(1-p)^{n-1} c}{(1-p)^{n-1} c + p^{n-1} (1-c)}.$$

Note that the sequence $\hat{\pi}_n$ is strictly decreasing in n . In addition, α_n is strictly negative for all n and $|\alpha_{n+1}(p, c)| < |\alpha_n(p, c)|$ for all c and all p in $(\frac{1}{2}, 1)$. Therefore, for every belief π the $\max_{n \in \mathbb{N}} w_n(\pi, p, c)$ is either $w_0(\pi, p, c) = 1 - \pi$ or $w_\infty(\pi, p, c) = 1 - c$. In other words, if the default action is 0, it is optimal to either not seek advice and choose action 0 or seek advice indefinitely until she learns the state. We know that DM 's equilibrium payoff is greater than $1 - \pi$ and $1 - c$ which implies that DM has no incentive to commit to 0 as default action. In conclusion, commitment has no benefit for the decision maker regardless of the default action.

3.4.2 Cost of waiting for information

In this section, we examine two variants of the model with different cost structures. Throughout the chapter, we studied a model that is well suited to situations where the information is difficult to process. In this section, we call it the c -model. Here, we focus on contexts where the information is easy to process once obtained and the opportunity cost of additional consultations is dominant. We show that similar results hold in these settings.

Consultation cost: κ -model Consider the case where DM does not incur any cost to process the information, but she incurs a cost $\kappa > 0$ every time she consults an agent. We

obtain an equilibrium of the same form as in the c -model. Proposition 3.6 states analogous results to Propositions 3.1, 3.2 and 3.3. As before, $\bar{\pi}$ is determined explicitly and $\underline{\pi}$ is defined implicitly. We observe that $\bar{\pi}$ is increasing in p and decreasing in κ , which is similar to the result of Proposition 3.4. Also, $\underline{\pi}$ is increasing in κ due to the fact that DM 's expected payoff is decreasing in κ . However, the effect of a change in p on DM 's expected payoff and as a consequence on $\underline{\pi}$ is ambiguous. DM 's expected payoff at a belief π in $(\underline{\pi}, \bar{\pi}) \cap [\tilde{\pi}_n, \tilde{\pi}_{n-1})$ is⁴

$$v_n(\pi, p, \kappa) = \pi(1 - \kappa \frac{1 - p^n}{1 - p}) + (1 - \pi)(1 - (1 - p)^n)(1 - \frac{\kappa}{p}).$$

It can be divided into two components: $v_\infty(\pi, p, \kappa)$ the expected payoff of a naive DM who would seek advice until she learns the state (without updating her beliefs when the agent is silent) and $v_n^+(\pi, p, \kappa)$ her additional gain from learning.

$$v_\infty(\pi, p, \kappa) = \pi(1 - \frac{\kappa}{1 - p}) + (1 - \pi)(1 - \frac{\kappa}{p})$$

$$v_n^+(\pi, p, \kappa) = \pi \frac{\kappa p^n}{1 - p} - (1 - \pi)(1 - \frac{\kappa}{p})(1 - p)^n$$

If p increases, $v_n^+(\pi, p, \kappa)$ increases: the benefit from learning is greater if the bias of the agent is less uncertain. However, the effect of such an increase on $v_\infty(\pi, p, \kappa)$ depends on π : for π small enough, the state is more likely to be 0 so an increase in the proportion of type L agents is beneficial whereas the opposite is true for π large enough. The overall effect on DM 's expected payoff is not clear and $\underline{\pi}$ is not always in the region where v_∞ is increasing in p .

However, we observe numerically that $\underline{\pi}$ is decreasing in p , which is the same result as in the c -model (see Proposition 3.4). Figure 3.3 shows the typical shapes of $\underline{\pi}$ as a function of p .

The difference between the two plotted cases is that for κ above $\frac{1}{4}$, there are values of p

⁴See proof of Proposition 3.6.

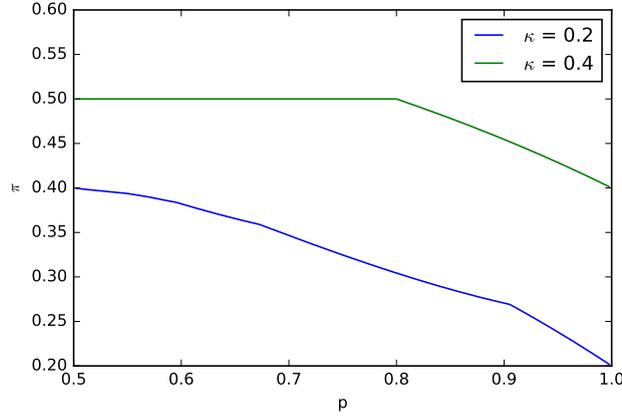


Figure 3.3: $\underline{\pi}$ as a function of p for $\kappa = 0.2$ and $\kappa = 0.4$.

such that there is no consultation in equilibrium ($\underline{\pi}$ is set equal to $\frac{1}{2}$ in that case) whereas this is impossible if κ is below $\frac{1}{4}$.

Moreover, the probability of bad decisions is expressed in the same way as in the c -model so that the analysis and conclusions about the probability of bad decisions hold in this setting as well. The difference is that DM 's expected payoff from consultation may not be monotonic in p for all priors and consultation cost κ . Also, the sequence of beliefs $\tilde{\pi}_n$ is now given by:

$$\tilde{\pi}_n = \frac{(1-p)^n(p-\kappa)}{(1-p)^n(p-\kappa) + p^n\kappa}.$$

It follows that for $n \geq 1$, $\tilde{\pi}_n$ is decreasing in p and κ so that we reach a similar conclusion about the way the parameters affect the maximum number of consultations at a given prior π in $(\underline{\pi}, \bar{\pi}) \cap [\tilde{\pi}_n, \tilde{\pi}_{n-1})$ for some n : if p or κ increases then the maximum number of consultations weakly decreases. This observation allows us to conclude that if κ increases, the probability of bad decisions decreases: DM is willing to consult more agents. The effect of a change in p is exactly the same as in the c -model.

Proposition 3.6. *1. If agents strictly prefer immediate decisions then there exists a unique PBE. In equilibrium, both types reveal information in both states (full revelation). DM*

consults at most one agent in equilibrium:

- if $\kappa \geq \frac{1}{2}$ she chooses action $a^*(\pi)$ without seeking advice.
- if $\kappa < \frac{1}{2}$ she consults exactly one agent if π is in $(\kappa, 1 - \kappa)$, learns the state and chooses the right action. Otherwise, she chooses action $a^*(\pi)$ without consultation.
- Off equilibrium path, i.e. if the agent reveals no information, she continues seeking advice.

2. If agents strictly prefer delayed bad outcomes then there exists a partial revelation PBE and it is unique. In equilibrium, type L (respectively H) reveals the information only in state 0 (respectively 1).

(a) If $p = \frac{1}{2}$, DM's equilibrium strategy is:

- if $\kappa \geq \frac{1}{4}$ she chooses action $a^*(\pi)$ without seeking advice.
- if $\kappa < \frac{1}{4}$ she seeks advice until she learns the state if π is in $(2\kappa, 1 - 2\kappa)$. Otherwise, she chooses action $a^*(\pi)$ without seeking advice.

(b) If $p > \frac{1}{2}$, DM's equilibrium strategy is:

- if $\kappa \geq \frac{p}{2}$ she chooses action $a^*(\pi)$ without seeking advice.
- if $\kappa < \frac{p}{2}$, there exist two beliefs $\underline{\pi}$ and $\bar{\pi}$ such that $\underline{\pi} < \frac{1}{2} < \bar{\pi} = 1 - \frac{\kappa}{p}$ and at each period t :
 - if $\underline{\pi} < \pi_t < \bar{\pi}$, she seeks advice.
 - otherwise, she stops and chooses action $a^*(\pi_t)$.

Proof. See Appendix. □

Discounting future payoff: δ -model Here DM incurs no processing cost as well, but she has a discount factor δ in $(0, 1)$ which expresses her preference for the present. We can

interpret an increase in δ as a reduction of the cost of waiting. Indeed, the analysis and results are identical to the κ -model with this interpretation. Proposition 3.7 states analogous results to Proposition 3.6. We also observe numerically that $\underline{\pi}$ is decreasing in p .

Proposition 3.7. *1. If agents strictly prefer immediate decisions then there exists a unique PBE. In equilibrium, both types reveal information in both states (full revelation). DM consults at most one agent in equilibrium:*

- if $\delta \leq \frac{1}{2}$ she chooses action $a^*(\pi)$ without seeking advice.
- if $\delta > \frac{1}{2}$ she consults exactly one agent if π is in $(1 - \delta, \delta)$, learns the state and chooses the right action. If $\pi \leq 1 - \delta$ or $\pi \geq \delta$, she chooses action $a^*(\pi)$ without seeking advice.
- Off equilibrium path, i.e. if the agent reveals no information, she continues seeking advice.

2. If agents strictly prefer delayed bad outcomes then there exists a partial revelation PBE and it is unique. In equilibrium, type L (respectively H) reveals the information only in state 0 (respectively 1).

(a) If $p = \frac{1}{2}$, DM's equilibrium strategy is:

- if $\delta \leq \frac{2}{3}$ she chooses action $a^*(\pi)$ without seeking advice.
- if $\delta > \frac{2}{3}$ she seeks advice until she learns the state if π is in $(\frac{2(1-\delta)}{2-\delta}, \frac{\delta}{2-\delta})$. Otherwise, she chooses action $a^*(\pi)$ without seeking advice.

(b) If $p > \frac{1}{2}$, DM's equilibrium strategy is:

- if $\delta \leq \frac{1}{1+p}$ she chooses action $a^*(\pi)$ without seeking advice.
- if $\delta > \frac{1}{1+p}$, there exist two beliefs $\underline{\pi}$ and $\bar{\pi}$ such that $\underline{\pi} < \frac{1}{2} < \bar{\pi} = \frac{\delta p}{1-\delta(1-p)}$ and at each period t :

- if $\underline{\pi} < \pi_t < \bar{\pi}$, she seeks advice.
- otherwise, she stops and chooses action $a^*(\pi_t)$.

Proof. See Appendix. □

3.4.3 Processing cost avoidance

In the main model, DM has to process the information if it is presented. In this section, we relax this assumption: if an agent is willing to reveal the state, DM chooses whether to accept the report, consult another agent or implement an action. We study the issues that arise in this case and provide a solution for p close to $\frac{1}{2}$.

First, we observe that the equilibrium of Proposition 3.3 is not an equilibrium in this context. The reason is that if DM encounters an agent that wants to reveal the state at time t , DM has an incentive to deviate from accepting the report to the following strategy: reject the report and follow the strategy of Proposition 3.3 at time $t + 1$. The expected payoff of this strategy is exactly the equilibrium expected payoff which is higher than $1 - c$ (see Remark 3.2).

Second, there cannot be an equilibrium where type L (respectively H) reveals the information only in state 0 (respectively 1). In fact, given this strategy DM 's belief π increases if the agent is silent and decreases if he wants to reveal the state. DM would learn for free and she would have an incentive to learn indefinitely without incurring the processing cost and as a consequence she would never choose an action.

This shows the limits of the model in this context. Indeed, if the number of consultations is large it is reasonable to take a consultation cost into account. Let $\kappa > 0$ denote this cost with κ small in comparison to c .

Note that Proposition 3.2 holds here. Namely, if $p = \frac{1}{2}$ DM seeks advice at beliefs π in $(c, 1 - c)$ until she learns the state. If we take κ into account, then this interval is replaced with $(c + 2\kappa, 1 - c - 2\kappa)$ because the average number of consultations is 2. The equilibrium expected

payoff at these beliefs is $1 - c - 2\kappa$.

Consider $p = \frac{1}{2} + \varepsilon$. For small enough $\varepsilon > 0$, there exists a similar equilibrium to that of Proposition 3.3: the expected payoff from consultation is $1 - c - 2\kappa + O(\varepsilon)$ which is smaller than $1 - c$ for small ε . Thus, it is optimal for *DM* to process information when it is presented if p is close to $\frac{1}{2}$.

3.5 Alternative model: information precision

In the binary state space model, the analysis of sequential consultation is tractable but the information structure is not flexible. In this section, we assume the state space is a compact interval and the verifiable messages structure is such that the agent can certify any compact interval containing the true state. This structure is rich enough to allow him to reveal the state, remain silent or to make a true statement of any intermediary precision about the state.

We characterize an equilibrium in the basic game, i.e. the game without sequential consultation, and show that all PBE are outcome equivalent. In the game with sequential consultation, we find an equilibrium where *DM* consults one agent (under certain conditions) and where the possibility of sequential consultation serves as a threat that allows her to extract more precise information. In fact, agents equilibrium messages are just precise enough to make her stop and choose an action rather than continue to seek advice from other agents.

DM has to choose an action a in $A = [0, 1]$. Her payoff depends on her action and on the state of the world s which is uniformly distributed in $S = [0, 1]$:

$$u_D(a, s) = 1 - l(a, s) = 1 - (a - s)^2,$$

where l is the loss function. As before, agents are either of type *L* or *H* with utility functions depending only on *DM*'s action. The payoff of type *L* (respectively, *H*) is strictly decreasing (respectively, increasing) in a . Information is verifiable in the following sense: if the state is s ,

then the set of available messages is

$$\mathcal{M}(s) = \{[x, y] \text{ s.t. } 0 \leq x \leq s \leq y \leq 1\}.$$

In the basic game, *Nature* draws s according to the uniform prior on S and then draws an agent whose type τ is L with probability p and H with probability $1 - p$. The agent observes s and sends a message m in $\mathcal{M}(s)$ to the decision maker who then chooses an action a in A . Let $\mathbf{m}_\tau(s)$ denote the message that an agent of type τ sends when the state is s , and let $\mathbf{a}(m)$ be DM 's action after receiving message m . The quadratic loss function of DM implies that on the equilibrium path $\mathbf{a}(m)$ is equal to the expected value of s given the message m .

In the game with sequential consultation, *Nature* draws s according to the uniform prior on S . Then, DM can start consulting agents. At the beginning of every period t in \mathbb{N} , she decides whether to consult an agent or to stop and choose an action. If she chooses to consult an agent, *Nature* draws an informed agent from the same distribution (type L or H with probability p and $1 - p$ respectively), who sends a message in $\mathcal{M}(s)$ to DM . If she stops and chooses an action the game ends. Notice that similarly to the basic game, DM 's action is also the expected value of s given all the messages she received.

3.5.1 Game without sequential consultation

In this section, we study the basic game where a single agent is drawn and the decision maker has to choose an action after receiving a message from that agent. The goal of this section is to construct a Perfect Bayesian Equilibrium (PBE) of this game and to show that its outcome is the unique PBE outcome.

Let \hat{a} be an action in $[0, 1]$ and consider DM 's strategy $\hat{\mathbf{a}}$ and agents strategies $\hat{\mathbf{m}}_L$ and $\hat{\mathbf{m}}_H$

defined as follows^{5 6}:

$$\hat{\mathbf{a}}([x, y]) = \begin{cases} y & \text{if } y < \hat{a} \\ \hat{a} & \text{if } x \leq \hat{a} \leq y \\ x & \text{if } \hat{a} < x \end{cases}$$

$$\hat{\mathbf{m}}_{\mathbf{L}}(s) = \begin{cases} [0, s] & \text{if } s \leq \hat{a} \\ [0, 1] & \text{if } \hat{a} < s \end{cases}$$

$$\hat{\mathbf{m}}_{\mathbf{H}}(s) = \begin{cases} [0, 1] & \text{if } s < \hat{a} \\ [s, 1] & \text{if } \hat{a} \leq s \end{cases}$$

In words, the strategy $\hat{\mathbf{a}}$ selects the action \hat{a} whenever it is justifiable, i.e. whenever it is in the certified interval, and it selects the closest action to \hat{a} in the certified interval otherwise. It is therefore easy to check that $\hat{\mathbf{m}}_{\mathbf{L}}$ and $\hat{\mathbf{m}}_{\mathbf{H}}$ are best responses of types L and H to $\hat{\mathbf{a}}$ respectively. Therefore $(\hat{\mathbf{a}}, \hat{\mathbf{m}}_{\mathbf{L}}, \hat{\mathbf{m}}_{\mathbf{H}})$ is a PBE if and only if $\hat{\mathbf{a}}$ is sequentially rational.

Proposition 3.8. *The strategy profile $(\hat{\mathbf{a}}, \hat{\mathbf{m}}_{\mathbf{L}}, \hat{\mathbf{m}}_{\mathbf{H}})$ is a PBE if and only if*

$$\hat{a} = \frac{\sqrt{p}}{\sqrt{p} + \sqrt{1-p}}.$$

Proof. If $(\hat{\mathbf{a}}, \hat{\mathbf{m}}_{\mathbf{L}}, \hat{\mathbf{m}}_{\mathbf{H}})$ is a PBE then on the equilibrium path, action \hat{a} is sequentially rational when $[0, 1]$ is certified, i.e. when the agent remains silent. Given the uniform prior over S and the probability p that the agent's type is L , this condition is equivalent to (for $p \in (0, 1)$):

⁵The strategy $\hat{\mathbf{a}}$ is reminiscent of the strategy used by [Giovannoni and Seidmann \[2007\]](#) to construct the pooling equilibrium of Proposition 3.3. in a setting with multiple types who have single-peaked utilities and prefer the decision maker to be uninformed rather than fully informed.

⁶The message $[0, s]$ (respectively $[s, 0]$) used in $\hat{\mathbf{m}}_{\mathbf{L}}$ (respectively $\hat{\mathbf{m}}_{\mathbf{H}}$) could be replaced by any message of the form $[x, s]$ (respectively $[s, y]$) with $0 \leq x \leq s$ (respectively $s \leq y \leq 1$) without affecting the argument or the result.

$$\hat{a} = \frac{(1-p) \int_0^{\hat{a}} s ds + p \int_{\hat{a}}^1 s ds}{(1-p)\hat{a} + p(1-\hat{a})}$$

which can be rearranged under the following form: $p(1-\hat{a})^2 = (1-p)\hat{a}^2$ or equivalently

$$\frac{\hat{a}}{1-\hat{a}} = \sqrt{\frac{p}{1-p}}$$

and therefore

$$\hat{a} = \frac{\sqrt{p}}{\sqrt{p} + \sqrt{1-p}}.$$

This expression can be extended to the cases where p is equal to 0 or 1. If $p = 0$, i.e. the agent is of type H with certainty, then $\hat{a} = 0$ is the only value for which $\hat{\mathbf{a}}$ is sequentially rational: it is the skepticism strategy implied by the unravelling argument which holds in this case. The same is true if $p = 1$, in which case $\hat{a} = 1$.

In order to prove the converse assertion, we have to show that $\hat{\mathbf{a}}$ is sequentially rational. For all messages that are not sent in equilibrium, there exists a belief, consistent with the certified information, that justifies the action selected by $\hat{\mathbf{a}}$ (i.e., this strategy selects an action in the certified interval). On equilibrium path, given strategies $\hat{\mathbf{m}}_{\mathbf{L}}$ and $\hat{\mathbf{m}}_{\mathbf{H}}$, it is sequentially rational to choose action x (respectively y) if $[x, 1]$ (respectively $[0, y]$) is certified. Finally, we have already shown that action \hat{a} is sequentially rational when $[0, 1]$ is certified. \square

In the remainder, let $\hat{a} = \frac{\sqrt{p}}{\sqrt{p} + \sqrt{1-p}}$. The outcome $\hat{\omega}$ of this equilibrium as a function of the state s and agent's type τ is given by the following mapping:

$$\hat{\omega}(s, \tau) = \begin{cases} \hat{a} & \text{if } (s \leq \hat{a}, \tau = H) \text{ or } (s \geq \hat{a}, \tau = L) \\ s & \text{otherwise.} \end{cases}$$

Proposition 3.9. *$\hat{\omega}$ is the unique PBE outcome of the game.*

Proof. Consider a PBE strategy profile $(\mathbf{a}, \mathbf{m}_L, \mathbf{m}_H)$ and let ω denote the outcome of this PBE. Let $a_S = \mathbf{a}(S)$ be the action that is chosen if the agent certifies the whole set $S = [0, 1]$, i.e. if he remains silent⁷. The message $m = S$ is always available to the agent and therefore we have $\omega(s, L) \leq a_S$ and $\omega(s, H) \geq a_S$ for every s in S ⁸. The message $m = \{s\}$ is available only in state s . It follows that $\omega(s, L) \leq s$ and $\omega(s, H) \geq s$ for all s .

Assume that for some state $\bar{s} \geq a_S$, $\omega(\bar{s}, L) = \tilde{a} < a_S$. Let $[x, y]$ be an interval that type L certifies with positive probability in state \bar{s} . This implies that $[x, y]$ contains \bar{s} . It follows from $\omega(\bar{s}, L) < a_S$ that $x < a_S$ and that there exists a state \underline{s} such that $x \leq \underline{s} < \tilde{a}$ and at least one type certifies $[x, y]$ in \underline{s} with positive probability in equilibrium. This is necessary to ensure that $\mathbf{a}([x, y]) = \tilde{a}$ is sequentially rational. However, it contradicts at least one of the conditions $\omega(\underline{s}, H) \geq a_S > \tilde{a}$ and $\omega(\underline{s}, L) \leq \underline{s} < \tilde{a}$. Therefore, we have

$$\omega(s, L) = a_S, \text{ for all } s \geq a_S.$$

Similarly, we can show that

$$\omega(s, H) = a_S, \text{ for all } s \leq a_S.$$

Assume that for some state $\bar{s} > a_S$, $\omega(\bar{s}, H) = \tilde{a} > \bar{s}$. Analogously to the previous argument, this would require the existence of some state $\bar{\bar{s}} > \tilde{a}$ and of a message that is certified with positive probability by type H in state \bar{s} and by at least one type in state $\bar{\bar{s}}$ so that choosing \tilde{a} after receiving this message is sequentially rational. This leads us again to a contradiction given that $\omega(\bar{\bar{s}}, H) \geq \bar{\bar{s}} > \tilde{a}$ and $\omega(\bar{\bar{s}}, L) \leq a_S < \tilde{a}$. Finally, we get

$$\omega(s, H) = s, \text{ for all } s > a_S.$$

⁷We restrict attention to pure strategies of the *DM* because it is never optimal to randomize given the quadratic loss function.

⁸The outcome ω is deterministic due to the strictly monotonic preferences of both agent types and to the fact that \mathbf{a} is pure.

Similarly, we can show that

$$\omega(s, L) = s, \text{ for all } s < a_S.$$

It remains to prove that $a_S = \hat{a}$. Let \mathcal{M}_S denote the set of messages sent in equilibrium for which DM chooses the action a_S . Note that if instead of receiving the messages contained in \mathcal{M}_S , DM only learned certifiably that the message is in \mathcal{M}_S , the sequentially rational action would still be a_S . Moreover, learning that the message is in \mathcal{M}_S is equivalent to learning that the following statement is true: $(s \leq a_S, \tau = H)$ or $(s \geq a_S, \tau = L)$. It follows that a_S satisfies the same condition as \hat{a} and is therefore equal to \hat{a} . In conclusion, $\omega = \hat{\omega}$.

□

Proposition 3.9 shows that in the game without sequential consultation, the equilibrium outcome is the same as the one we would obtain if the agents could either reveal the state or remain silent: the richness of the set of available messages does not affect communication. In the next section, we show that if sequential consultation is allowed, information of intermediate precision can be revealed in equilibrium even when DM consults only one agent.

3.5.2 Game with sequential consultation

In this section, DM can consult multiple agents sequentially. We focus on the case $p = \frac{1}{2}$ where in equilibrium, DM would not gain additional information from receiving the same message more than once (similarly to the binary state model with equiprobable types) which makes the analysis more tractable. In fact, learning happens only if new information is presented which precludes the kind of learning we observe in Proposition 3.3.

We look for a PBE where DM consults exactly one agent and uses the possibility of further consultations as a threat to extract more precise information. In this type of equilibrium, she consults one agent at cost κ and chooses an action based on the message he certifies. The agent's strategies in such equilibria depends only on the state s and agent's type τ , similarly

to equilibrium strategies in the basic game. We call such an equilibrium a *single-consultation PBE*.

Note that there exist multiple equilibria in this setting. Our goal is simply to construct an equilibrium where *DM* receives the most precise information while consulting only one agent and agents send the least informative message while giving *DM* incentive to stop consultation and choose an action. In other words, when an agent is consulted he captures *DM*'s attention and influences her decision in his favor but the resulting error is small enough to make her stop.

First, we show that the equilibrium of section 3.5.1 $(\hat{\mathbf{a}}, \hat{\mathbf{m}}_{\mathbf{L}}, \hat{\mathbf{m}}_{\mathbf{H}})$ is generically not an equilibrium in this setting. Then, we define an equilibrium satisfying the conditions above for a small enough consultation cost κ .

Claim 3.1. $(\hat{\mathbf{a}}, \hat{\mathbf{m}}_{\mathbf{L}}, \hat{\mathbf{m}}_{\mathbf{H}})$ defines a single-consultation PBE only if $\kappa = \frac{1}{24}$.

Proof. It follows from the analysis of section 3.5.1 that $\hat{\mathbf{m}}_{\mathbf{L}}$ and $\hat{\mathbf{m}}_{\mathbf{H}}$ are best responses of types *L* and *H* respectively to $\hat{\mathbf{a}}$. Therefore, $(\hat{\mathbf{a}}, \hat{\mathbf{m}}_{\mathbf{L}}, \hat{\mathbf{m}}_{\mathbf{H}})$ defines a single-consultation PBE if and only if (i) it is optimal for *DM* to seek advice and (ii) it is sequentially rational to stop after consulting exactly one agent and choose the action of strategy $\hat{\mathbf{a}}$.

(i) Seeking advice is optimal if and only if *DM*'s expected payoff is greater than the expected payoff from choosing an action without consultation. The optimal action in that case is the expected value of s which is $\frac{1}{2}$ given the uniform prior, and the expected loss is therefore:

$$\int_0^1 (s - \frac{1}{2})^2 ds = \frac{1}{12}.$$

If *DM* consults exactly one agent, she incurs the cost κ and given that she expects to follow the strategy $\hat{\mathbf{a}}$, her expected loss (in addition to the consultation cost) is

$$\int_0^1 \frac{1}{2} (s - \frac{1}{2})^2 ds = \frac{1}{24}.$$

In conclusion seeking advice is optimal if and only if $\kappa + \frac{1}{24} \leq \frac{1}{12}$, i.e. $\kappa \leq \frac{1}{24}$.

(ii) On the equilibrium path, it is sequentially rational to stop seeking advice and follow the strategy $\hat{\mathbf{a}}$ if the certified interval is of the form $[0, y]$ (with $y \leq \frac{1}{2}$) or $[x, 1]$ (with $x \geq \frac{1}{2}$) because these messages fully reveal the state in equilibrium. If the agent remains silent, DM 's updated belief is identical to her uniform prior⁹. DM 's expected loss if she stops seeking advice and follows strategy $\hat{\mathbf{a}}$ is therefore $\frac{1}{12}$.

In order to compute the expected loss of further consultations, note that if DM continues after receiving the message $[0,1]$ in the first period, she expects to continue at each period as long as she receives the same message (because her belief does not change) and to stop if she receives a different message, in which case she would learn the state. At each period, there is a probability $\frac{1}{2}$ to receive the message $[0,1]$ and a probability $\frac{1}{2}$ to receive a fully revealing message. Her expected loss is therefore

$$\kappa \sum_{n=1}^{+\infty} n \left(\frac{1}{2}\right)^n = 2\kappa.$$

Consequently, it is optimal to stop consulting and follow strategy \mathbf{a} on the equilibrium path if and only if $2\kappa \geq \frac{1}{12}$, i.e. $\kappa \geq \frac{1}{24}$. The conclusion follows. □

If κ is too large, DM simply chooses the prior optimal action without consultation, whereas if it is too small she would continue seeking advice if she receives a non-informative message in the first period. In the former case there is an equilibrium in which DM chooses action $\frac{1}{2}$ without asking for information (expecting that agents follow strategies $\hat{\mathbf{m}}_{\mathbf{L}}$ and $\hat{\mathbf{m}}_{\mathbf{H}}$). In order to find an equilibrium in the latter case, we first determine how informative a message has to be in order to make it optimal for DM to stop seeking advice regardless of her belief.

⁹The fact that $p = \frac{1}{2}$ simplifies the analysis by eliminating free learning similarly to what we observe in the binary state space model.

Let η denote the largest message size for which it is always optimal to stop and choose an action: it is the largest d such that if $[x, y]$ with $y - x \leq d$ is certified, it is optimal for DM to stop consultation for every possible belief.

Lemma 3.2. $\eta = 2\sqrt{\kappa}$.

Proof. Suppose $[x, y]$ such that $y = x + d$ is certified. Let Z be the random variable that describes DM 's belief after receiving this message. If DM stops, she chooses the action $\mathbb{E}(Z)$ and her expected loss is $\text{Var}(Z)$. Let $W = (Z - \frac{x+y}{2})^2$. By definition, we have $\text{Var}(Z) \leq \mathbb{E}(W)$. Given that Z is a.s. in $[x, y]$, we have $W \leq \frac{d^2}{4}$ and therefore $\text{Var}(Z) \leq \frac{d^2}{4}$. The expected loss of continuing to seek advice after receiving the message $[x, y]$ is greater or equal to κ as DM has to incur the consultation cost. Therefore, if $\frac{d^2}{4} \leq \kappa$ it is optimal to stop seeking advice for every possible belief.

For Z such that $\mathbb{P}(Z = x) = \mathbb{P}(Z = y) = \frac{1}{2}$ we have $\mathbb{E}(Z) = \frac{x+y}{2}$ and $\text{Var}(Z) = \frac{d^2}{4}$. The expected loss of continuing to seek advice after receiving the message $[x, y]$ is equal to κ if DM expects to learn the state with certainty. For these beliefs, it is not optimal to stop seeking advice if $\frac{d^2}{4} > \kappa$. The conclusion follows. \square

Remark 3.3. Lemma 3.2 implies that there can be no fully revealing equilibrium. A consulted agent can always make DM stop if he sends a message of smaller size than η . For such a message, DM 's action is necessarily in the interval. Consequently, in any state at least one of the two types benefits from sending such a message.

For a small enough κ , consider the strategies \mathbf{m}_L^η and \mathbf{m}_H^η defined as follows:

$$\mathbf{m}_L^\eta(s) = \begin{cases} [0, s] & \text{if } 0 \leq s < \eta \\ [0, 2\eta] & \text{if } \eta \leq s \leq 2\eta \\ [s - \eta, s] & \text{if } 2\eta < s < 1 - \eta \\ [1 - 2\eta, 1] & \text{if } 1 - \eta \leq s \leq 1 \end{cases}$$

$$\mathbf{m}_{\mathbf{H}}^{\eta}(s) = \begin{cases} [0, 2\eta] & \text{if } 0 \leq s \leq \eta \\ [s, s + \eta] & \text{if } \eta < s < 1 - 2\eta \\ [1 - 2\eta, 1] & \text{if } 1 - 2\eta \leq s \leq 1 - \eta \\ [s, 1] & \text{if } 1 - \eta < s \leq 1 \end{cases}$$

If type L (respectively, H) follows strategy $\mathbf{m}_{\mathbf{L}}^{\eta}$ (respectively, $\mathbf{m}_{\mathbf{H}}^{\eta}$) he certifies the interval $[s - \eta, s]$ (respectively, $[s, s + \eta]$) in state s if s is far enough from the boundaries. By definition of η , given such a message it is optimal for DM to stop seeking advice. Near the boundaries, $\mathbf{m}_{\mathbf{L}}^{\eta}$ and $\mathbf{m}_{\mathbf{H}}^{\eta}$ mimic $\hat{\mathbf{m}}_{\mathbf{L}}$ and $\hat{\mathbf{m}}_{\mathbf{H}}$ respectively.

In order to construct a single-consultation PBE where agents strategies are $\mathbf{m}_{\mathbf{L}}^{\eta}$ and $\mathbf{m}_{\mathbf{H}}^{\eta}$, we need to specify both on and off equilibrium path decisions of DM . On the equilibrium path, assuming that it is sequentially rational to stop seeking advice for information after exactly one period, DM 's optimal action would be:

$$\mathbf{a}^{\eta}(m) = \begin{cases} y & \text{if } m = [0, y] \text{ and } y < \eta \\ \eta & \text{if } m = [0, 2\eta] \\ x + \frac{\eta}{2} & \text{if } m = [x, x + \eta] \text{ and } \eta < x < 1 - 2\eta \\ 1 - \eta & \text{if } m = [1 - 2\eta, 1] \\ x & \text{if } m = [x, 1] \text{ and } x > 1 - \eta \end{cases}$$

For fully revealing messages, DM chooses the right action. For all other messages, she chooses the middle point of the certified interval which is optimal due to the symmetries of the model.

Off equilibrium messages can be separated in two categories. The first contains every off equilibrium message that certifies a subset of some equilibrium message. For such a message

$m = [x, y]$, DM can stop seeking advice and choose the action:

$$\mathbf{a}^\eta([x, y]) = \begin{cases} y & \text{if } y \leq \eta \\ \eta & \text{if } \eta \in [x, y] \subset [0, 2\eta] \\ x & \text{if } [x, y] \subset [\eta, 2\eta] \\ \frac{x+y}{2} & \text{if } y < x + \eta \text{ and } \eta < x < 1 - 2\eta \\ y & \text{if } [x, y] \subset [1 - 2\eta, 1 - \eta] \\ 1 - \eta & \text{if } 1 - \eta \in [x, y] \subset [1 - 2\eta, 1] \\ x & \text{if } x \geq 1 - \eta \end{cases}$$

These decisions are rational (actions are in the certified interval, so that there exists a belief that justifies them) and do not provide any incentive to agents to deviate from \mathbf{m}_L^η and \mathbf{m}_H^η . Indeed, they are similar to off equilibrium decisions of $\hat{\mathbf{a}}$ near the boundaries. For messages far from the boundaries, DM would select the middle point of the interval which means that agents have no incentive to certify an interval smaller than η .

The second category contains every message that is not a subset of an equilibrium message. These are exactly the messages $m = [x, y]$ such that $y - x > \eta$, $x < 1 - 2\eta$ and $y > 2\eta$. As shown in the proof of Lemma 3.2, the belief that the state is x with probability $\frac{1}{2}$ and y with probability $\frac{1}{2}$ justifies continuing to seek advice if DM expects to learn the true state in the next period, which is the case under this belief and given the strategies \mathbf{m}_L^η and \mathbf{m}_H^η . Extending the definition of DM 's strategy, let $\mathbf{a}^\eta(m)$ be *to continue seeking advice*, for such messages m .

Proposition 3.10. $(\mathbf{a}^\eta, \mathbf{m}_L^\eta, \mathbf{m}_H^\eta)$ is a single-consultation PBE if and only if $\kappa \leq \frac{1}{36}$.

Proof. First, we show that there are no profitable deviations from \mathbf{m}_L^η and \mathbf{m}_H^η for agents. As it is argued above, there is no incentive to send off equilibrium messages of the first category. We need to show it is also the case for the second category of off equilibrium messages. If an

agent certifies $[x, y]$ such that $y - x > \eta$, $x < 1 - 2\eta$ and $y > 2\eta$ he expects DM to keep seeking advice. He expects her to receive an equilibrium message in the second period. Let $[x', y']$ be its intersection with $[x, y]$ and assume DM chooses $\mathbf{a}^\eta([x', y'])$. This action is rational and the deviation is not profitable for the same reasons as before.

This guarantees that $(\mathbf{a}^\eta, \mathbf{m}_L^\eta, \mathbf{m}_H^\eta)$ is a single-consultation PBE if and only if the strategies are well defined, seeking advice and stopping after exactly one period is optimal on the equilibrium path.

The strategies are well defined if and only if $\eta \leq 1 - 2\eta$, which is equivalent to

$$\kappa \leq \frac{1}{36}.$$

DM 's expected loss in such an equilibrium is the consultation cost κ together with a convex combination of 0 (if the message is fully revealing near the boundaries), $\frac{\eta^2}{4}$ (for equilibrium messages of size η) and $\frac{\eta^2}{3}$ (if the message is $[0, 2\eta]$ or $[1 - 2\eta, 1]$). Therefore her expected loss is smaller than $\frac{7}{3}\kappa$ which is in turn smaller than $\frac{1}{12}$ (her expected loss if she does not consult) for all $\kappa \leq \frac{1}{36}$. As a consequence, it is optimal to start seeking advice.

Stopping after receiving an equilibrium message certifying an interval of size η is optimal (see Lemma 3.2). Also, it is optimal to stop and follow the strategy \mathbf{a}^η if the certified interval is of the form $[0, y]$ (with $y \leq \eta$) or $[x, 1]$ (with $x \geq 1 - \eta$). After receiving the message $[0, 2\eta]$, DM 's posterior belief is that the state is uniformly distributed in $[0, 2\eta]$. Her expected loss if she stops and chooses action η is $\frac{\eta^2}{3}$ while her expected loss from continuing to seek advice is 2κ (same computation as in the proof of Claim 3.1). It follows that it is optimal to stop after consulting one agent if $[0, 2\eta]$ is certified. The same applies for message $[1 - 2\eta, 0]$.

□

For κ small enough, we have an equilibrium where the certified interval is of size $2\sqrt{\kappa}$ and the loss is 2κ , for states far enough from the boundaries: DM incurs the consultation

cost κ and then chooses an action that is $\sqrt{\kappa}$ away from the true state so that she loses an additional κ . In contrast, if sequential consultation is not allowed and DM can consult only one agent, her expected loss is $\kappa + \frac{1}{24}$. Although DM does not actually consult more agents in the equilibrium we constructed, she benefits from the possibility of sequential consultation to extract information and get a higher payoff. In this equilibrium, she is less likely to choose the action that matches the state (it can happen only in states near the boundaries, with an ex-ante probability of $\eta = 2\sqrt{\kappa}$) but the expected loss is reduced in comparison with the basic game due to error size reduction.

3.6 Appendix

Proof of Proposition 3.3. Assume ($p > \frac{1}{2}$). First, we note that in equilibrium, given p and c , DM 's decision at any point in time depends only on her belief at that time and that her decisions are dynamically consistent.

We show that there exists no equilibrium such that DM seeks advice for every belief $\pi \in (0, 1)$. Such an equilibrium would require DM to seek advice until she learns the state for all prior beliefs in $(0, 1)$, which implies that c has to be equal to zero because DM 's expected payoff at prior π using this strategy is $1 - c$ while her payoff if she stops is $\max\{\pi, 1 - \pi\}$ which is higher for π high enough.

Consider the strategy that consists of seeking advice exactly once, i.e. consulting one agent and stopping right afterwards whether he reveals the state or not. The expected payoff of this strategy depends on whether the prior belief is below or above $1 - p$ because the posterior belief if the agent is silent would be respectively below or above $\frac{1}{2}$, namely:

$$\begin{aligned} \pi(1-p)(1-c) + (1-\pi)(p(1-c) + (1-p)) & \text{ if } \pi \leq 1-p \\ \pi(p + (1-p)(1-c)) + (1-\pi)p(1-c) & \text{ if } \pi > 1-p. \end{aligned}$$

Using these payoffs, we first observe that there exists an equilibrium such that DM never seeks advice if and only if $c \geq p$. Then we determine the set of beliefs Π_1 such that DM seeks advice if she expects to stop after one period regardless of the result. Π_1 is empty if $c \geq p$. Otherwise (i.e. if $c < p$), Π_1 is an interval $(\underline{\pi}_1, \bar{\pi}_1)$ containing $\frac{1}{2}$ such that

$$\bar{\pi}_1 = \frac{p(1-c)}{p(1-c) + (1-p)c}$$

and

$$\underline{\pi}_1 = \begin{cases} \frac{1-p(1-c)}{1+p+(1-2p)(1-c)} & \text{if } \frac{(1-p)^2}{p^2+(1-p)^2} \leq c < p \\ \frac{pc}{pc+(1-p)(1-c)} & \text{if } 0 < c < \frac{(1-p)^2}{p^2+(1-p)^2} \end{cases}$$

Consider a belief π_0 in $(0, 1)$ and the sequence of beliefs $(\pi_k)_{k \in \mathbb{Z}}$ such that:

$$\forall k, \pi_{k+1} = \frac{p\pi_k}{p\pi_k + (1-p)(1-\pi_k)}.$$

This is the unique sequence of beliefs that contains π_0 such that if DM 's belief is π_k and she encounters a silent agent, she updates her belief to π_{k+1} . Using the fact that DM does not seek advice indefinitely in equilibrium, the optimal decision is to stop and choose action $a^*(\pi_k) = 1$ for all k such that $\pi_k \geq \bar{\pi}_1$. Furthermore, if π_k is in Π_1 , the optimal decision is to continue seeking advice. Finally, if $\pi_k \leq \underline{\pi}_1$ and DM stops at π_k then she also stops at all $\pi_{k'}$ such that $k' < k$. Therefore, if Π_1 is not empty, there exist \underline{k} and \bar{k} with $\pi_{\underline{k}} \leq \underline{\pi}_1$ and $\pi_{\bar{k}} < \bar{\pi}_1 \leq \pi_{\bar{k}+1}$ such that DM seeks advice at belief π_k if and only if $\underline{k} \leq k \leq \bar{k}$. However, if Π_1 is empty, DM does not seek advice no matter her beliefs.

In the remainder of the proof, assume $c < p$ ensuring that Π_1 is not empty. Consider the sequence of beliefs $(\pi_k)_{k \in \mathbb{Z}}$ with $\pi_0 = \bar{\pi}_1$. For a given $n \geq 1$, consider a belief π in $[\pi_{-n}, \pi_{-n+1})$. The analysis above implies that DM 's equilibrium strategy at belief π is either to stop and choose action $a^*(\pi)$ or to seek advice for at most n periods¹⁰.

Let $v_n(\pi, p, c)$ denote the expected payoff of the decision maker at belief π from seeking advice for n periods at most:

$$v_n(\pi, p, c) = \pi(p^n + (1 - p^n)(1 - c)) + (1 - \pi)(1 - (1 - p)^n)(1 - c).$$

Define $\phi_{p,c} : (0, \frac{1}{2}) \rightarrow \mathbb{R}$ such that $\phi_{p,c}(\pi) = \frac{v_n(\pi, p, c)}{1 - \pi}$ if π is in $[\pi_{-n}, \pi_{-n+1}) \cap (0, \frac{1}{2})$. $\phi_{p,c}$ is continuous, strictly increasing and takes values above and below 1. As a consequence there exists a unique $\underline{\pi}$ such that $\phi_{p,c}(\underline{\pi}) = 1$. Moreover, $\underline{\pi} \leq \bar{\pi}_1$ with equality if and only if $\bar{\pi}_1 \geq \pi_{-1} = 1 - c$.

In conclusion, DM seeks advice if and only if her belief is in the interval $(\underline{\pi}, \bar{\pi})$ such that $\bar{\pi} = \bar{\pi}_1$ and $\phi_{p,c}(\underline{\pi}) = 1$.

□

Proof of Corollary 3.1. Let n_c denote the maximum number of consultations at the belief $\pi = c$. Proposition 3.5 ensures that

$$1 - \bar{\pi} < \underline{\pi} < c.$$

If DM seeks advice at prior belief $1 - \bar{\pi}$ (which does not happen in equilibrium), her posterior belief if the agent is silent is

$$\frac{p(1 - \bar{\pi})}{p(1 - \bar{\pi}) + (1 - p)\bar{\pi}} = c.$$

This guarantees that the maximum number of consultations at the belief $\underline{\pi}$ is either n_c or

¹⁰To seek advice for at most n periods is to start seeking advice and to stop after encountering (i) an agent who reveals the state or (ii) silent agents for n periods.

$n_c + 1$. n_c satisfies $\tilde{\pi}_{n_c} \leq c < \tilde{\pi}_{n_c-1}$ which is equivalent to

$$\left(\frac{p}{1-p}\right)^{n_c-2} < \left(\frac{1-c}{c}\right)^2 \leq \left(\frac{p}{1-p}\right)^{n_c-1},$$

thus,

$$n_c - 1 < 1 + 2 \frac{\ln \frac{1-c}{c}}{\ln \frac{p}{1-p}} \leq n_c,$$

which implies that $n_c = \left\lceil 1 + 2 \frac{\ln \frac{1-c}{c}}{\ln \frac{p}{1-p}} \right\rceil$.

□

- Proof of Proposition 3.6.**
1. The argument of the proof of Proposition 3.1 applies in this case. The only difference is that DM 's payoff from consultation is $1 - \kappa$.
 2. (a) The argument of the proof of Proposition 3.2 applies in this case. The only difference is that DM 's expected payoff from consultation is:

$$1 - \kappa \sum_{k=1}^{+\infty} n \left(\frac{1}{2}\delta\right)^n = 1 - 2\kappa.$$

- (b) The proof of this result follows the same steps as the proof of Proposition 3.3. Assume ($p > \frac{1}{2}$). In equilibrium, given p and κ , DM 's decision at any point in time depends only on her belief at that time and that her decisions are dynamically consistent. There exists no equilibrium such that DM seeks advice for every belief $\pi \in (0, 1)$. Such an equilibrium would require DM to seek advice until she learns the state for all prior beliefs in $(0, 1)$, which requires κ to be 0 because DM 's expected payoff at prior π using this strategy is lower than $1 - \kappa$ while her payoff if she stops is $\max\{\pi, 1 - \pi\}$ which is higher for π high enough.

Using the payoffs expressed in Figure 3.4, we first check that there exists an equilib-

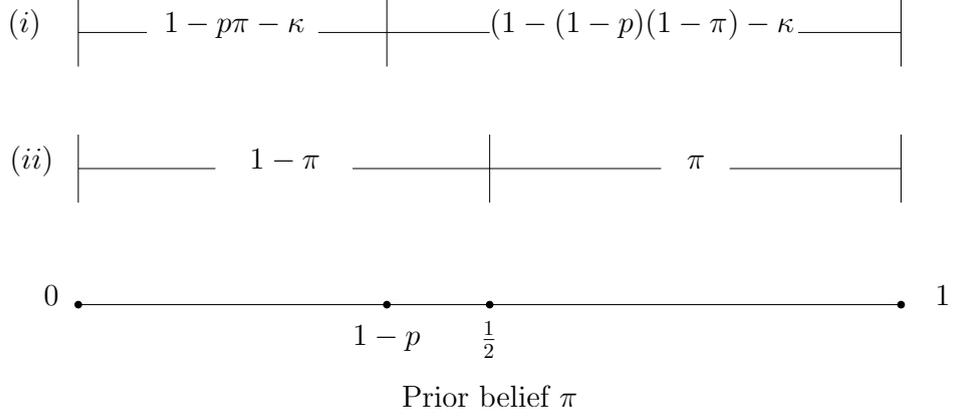


Figure 3.4: Expected payoffs: (a) DM consults exactly one agent, (b) DM chooses $a^*(\pi)$ without consultations

rium such that DM never seeks advice if and only if $\kappa \geq \frac{p}{2}$. Then we determine the set of beliefs Π_1 such that DM seeks advice if she expects to stop after one period regardless of the result. Π_1 is empty if $\kappa \geq \frac{p}{2}$. Otherwise, Π_1 is an interval $(\underline{\pi}_1, \bar{\pi}_1)$ containing $\frac{1}{2}$ such that

$$\bar{\pi}_1 = 1 - \frac{\kappa}{p}$$

and

$$\underline{\pi}_1 = \begin{cases} \frac{\kappa}{1-p} & \text{if } \kappa < (1-p)^2 \\ \frac{\kappa+1-p}{2-\kappa} & \text{if } (1-p)^2 \leq \delta < \frac{p}{2} \end{cases}$$

Consider a belief π_0 in $(0, 1)$ and the sequence of beliefs $(\pi_k)_{k \in \mathbb{Z}}$ such that:

$$\forall k, \pi_{k+1} = \frac{p\pi_k}{p\pi_k + (1-p)(1-\pi_k)}.$$

This is the unique sequence of beliefs that contains π_0 such that if DM 's belief is π_k and she encounters a silent agent, she updates her belief to π_{k+1} . Using the same argument as in the proof of Proposition 3.3, we conclude that if Π_1 is not empty, there exist \underline{k} and \bar{k} with $\pi_{\underline{k}} \leq \underline{\pi}_1$ and $\pi_{\bar{k}} < \bar{\pi}_1 \leq \pi_{\bar{k}+1}$ such that DM seeks advice

at belief π_k if and only if $\underline{k} \leq k \leq \bar{k}$. However, if Π_1 is empty, DM does not seek advice no matter her beliefs.

In the remainder of the proof, assume $\kappa < \frac{p}{2}$ ensuring that Π_1 is not empty. Consider the sequence of beliefs $(\pi_k)_{k \in \mathbb{Z}}$ with $\pi_0 = \bar{\pi}_1$. For a given $n \geq 1$, consider a belief π in $[\pi_{-n}, \pi_{-n+1})$. The analysis above implies that DM 's equilibrium strategy at belief π is either to stop and choose action $a^*(\pi)$ or to seek advice for at most n periods.

Let $v_n(\pi)$ denote the expected payoff of the decision maker at belief π from seeking advice for at most n periods:

$$\begin{aligned} v_n(\pi, p, \kappa) &= \pi + (1 - \pi)(1 - (1 - p)^n) \\ &\quad - \kappa \left(\sum_{k=1}^{n-1} k(\pi(1 - p)p^{k-1} + (1 - \pi)p(1 - p)^{k-1}) + n(\pi p^{n-1} + (1 - \pi)(1 - p)^{n-1}) \right) \\ &= \pi \left(1 - \kappa \frac{1 - p^n}{1 - p} \right) + (1 - \pi)(1 - (1 - p)^n) \left(1 - \frac{\kappa}{p} \right) \end{aligned}$$

Define $\phi_{p,\kappa} : (0, \frac{1}{2}) \rightarrow \mathbb{R}$ such that $\phi_{p,\kappa}(\pi) = \frac{v_n(\pi, p, \kappa)}{1 - \pi}$ if π is in $[\pi_{-n}, \pi_{-n+1}) \cap (0, \frac{1}{2})$. It is possible to show that $\phi_{p,\delta}$ is continuous. Moreover, it is strictly increasing for $\kappa \leq 1 - p$ and has a V -shape otherwise. It takes values above and below 1 and its limit at $\pi = 0$ is $1 - \frac{\kappa}{p} < 1$. As a consequence there exists a unique $\underline{\pi}$ such that $\phi_{p,\kappa}(\underline{\pi}) = 1$. Moreover, $\underline{\pi} \leq \underline{\pi}_1$ with equality if and only if $\underline{\pi}_1 \geq \pi_{-1}$.

In conclusion, DM seeks advice if and only if her belief is in the interval $(\underline{\pi}, \bar{\pi})$ such that $\bar{\pi} = \bar{\pi}_1$ and $\phi_{p,\kappa}(\underline{\pi}) = 1$.

□

Proof of Proposition 3.7. 1. The argument of the proof of Proposition 3.1 applies in this case. The only difference is that DM 's payoff from consultation is δ .

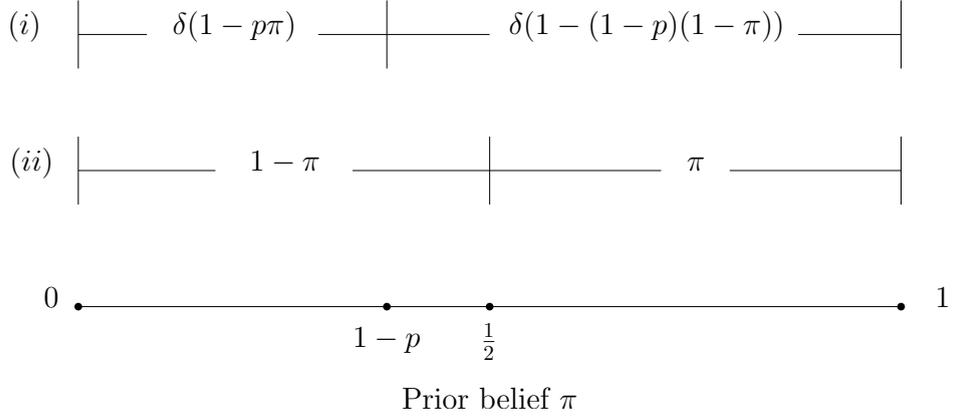


Figure 3.5: Expected payoffs: (a) *DM* consults exactly one agent, (b) *DM* chooses $a^*(\pi)$ without consultations

2. (a) The argument of the proof of Proposition 3.2 applies in this case. The only difference is that *DM*'s expected payoff from consultation is:

$$\sum_{k=1}^{+\infty} \left(\frac{1}{2}\delta\right)^k = \frac{\delta}{2-\delta}.$$

- (b) The proof of this result follows the same steps as the proof of Proposition 3.3. Assume ($p > \frac{1}{2}$). In equilibrium, given p and δ , *DM*'s decision at any point in time depends only on her belief at that time and that her decisions are dynamically consistent. There exists no equilibrium such that *DM* seeks advice for every belief $\pi \in (0, 1)$. Such an equilibrium would require *DM* to seek advice until she learns the state for all prior beliefs in $(0, 1)$, which requires δ to be 1 because *DM*'s expected payoff at prior π using this strategy is

$$\delta\pi(1-p) \sum_{k=0}^{+\infty} (\delta p)^k + \delta(1-\pi)p \sum_{k=0}^{+\infty} (\delta(1-p))^k = \delta\left(\pi \frac{1-p}{1-\delta p} + (1-\pi) \frac{p}{1-\delta(1-p)}\right)$$

while her payoff if she stops is $\max\{\pi, 1-\pi\}$ which is higher for π high enough.

Using the payoffs expressed in Figure 3.5, we first check that there exists an equilibrium such that DM never seeks advice if and only if $\delta \leq \frac{1}{1+p}$. Then we determine the set of beliefs Π_1 such that DM seeks advice if she expects to stop after one period regardless of the result. Π_1 is empty if $\delta \leq \frac{1}{1+p}$. Otherwise, Π_1 is an interval $(\underline{\pi}_1, \bar{\pi}_1)$ containing $\frac{1}{2}$ such that

$$\bar{\pi}_1 = \frac{\delta p}{1 - \delta(1 - p)}$$

and

$$\underline{\pi}_1 = \begin{cases} \frac{1-\delta p}{1+\delta(1-p)} & \text{if } \frac{1}{1+p} < \delta \leq \frac{p}{1-p(1-p)} \\ \frac{1-\delta}{1-\delta p} & \text{if } \frac{p}{1-p(1-p)} < \delta < 1 \end{cases}$$

Consider a belief π_0 in $(0, 1)$ and the sequence of beliefs $(\pi_k)_{k \in \mathbb{Z}}$ such that:

$$\forall k, \pi_{k+1} = \frac{p\pi_k}{p\pi_k + (1-p)(1-\pi_k)}.$$

This is the unique sequence of beliefs that contains π_0 such that if DM 's belief is π_k and she encounters a silent agent, she updates her belief to π_{k+1} . Using the same argument as in the proof of Proposition 3.3, we conclude that if Π_1 is not empty, there exist \underline{k} and \bar{k} with $\pi_{\underline{k}} \leq \underline{\pi}_1$ and $\pi_{\bar{k}} < \bar{\pi}_1 \leq \pi_{\bar{k}+1}$ such that DM seeks advice at belief π_k if and only if $\underline{k} \leq k \leq \bar{k}$. However, if Π_1 is empty, DM does not seek advice no matter her beliefs.

In the remainder of the proof, assume $\delta > \frac{1}{1+p}$ ensuring that Π_1 is not empty. Consider the sequence of beliefs $(\pi_k)_{k \in \mathbb{Z}}$ with $\pi_0 = \bar{\pi}_1$. For a given $n \geq 1$, consider a belief π in $[\pi_{-n}, \pi_{-n+1})$. The analysis above implies that DM 's equilibrium strategy at belief π is either to stop and choose action $a^*(\pi)$ or to seek advice for at most n periods.

Let $v_n(\pi)$ denote the expected payoff of the decision maker at belief π from seeking

advice for at most n periods:

$$\begin{aligned}
v_n(\pi, p, \delta) &= \delta\pi((1-p) \sum_{k=0}^{n-1} (\delta p)^k + \delta^{n-1} p^n) + \delta(1-\pi)p \sum_{k=0}^{n-1} (\delta(1-p))^k \\
&= \delta\pi((1-p) \frac{1 - (\delta p)^n}{1 - \delta p} + \delta^{n-1} p^n) + \delta(1-\pi)p \frac{1 - (\delta(1-p))^n}{1 - \delta(1-p)}
\end{aligned}$$

Define $\phi_{p,\delta} : (0, \frac{1}{2}) \rightarrow \mathbb{R}$ such that $\phi_{p,\delta}(\pi) = \frac{v_n(\pi,p,\delta)}{1-\pi}$ if π is in $[\pi_{-n}, \pi_{-n+1}) \cap (0, \frac{1}{2})$.

It is possible to show that $\phi_{p,\delta}$ is continuous, strictly increasing and takes values above and below 1. As a consequence there exists a unique $\underline{\pi}$ such that $\phi_{p,\delta}(\underline{\pi}) = 1$.

Moreover, $\underline{\pi} \leq \underline{\pi}_1$ with equality if and only if $\underline{\pi}_1 \geq \pi_{-1}$.

In conclusion, DM seeks advice if and only if her belief is in the interval $(\underline{\pi}, \bar{\pi})$ such that $\bar{\pi} = \bar{\pi}_1$ and $\phi_{p,\delta}(\underline{\pi}) = 1$.

□

Bibliography

Philippe Aghion and Jean Tirole. Formal and real authority organizations. *Journal of Political Economy*, 105(1):1–29, February 1997.

Elchanan Ben-Porath and Barton L Lipman. Implementation with partial provability. *Journal of Economic Theory*, 147(5):1689–1724, 2012.

Subir Bose and Ludovic Renou. Mechanism design with ambiguous communication devices. *Econometrica*, 82:1853–1872, 2014.

Jesse Bull. Mechanism Design with Moderate Evidence Cost. *The BE Journal of Theoretical Economics (Contributions)*, 8(5), 2008.

Jesse Bull and Joel Watson. Hard evidence and mechanism design. *Games and Economic Behavior*, 58:75–93, 2007.

Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6): 1431–1451, 1982.

Raymond Deneckere and Sergei Severinov. Mechanism design with partial state verifiability. *Games and Economic Behavior*, 64(2):487–513, 2008.

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

- Alfredo Di Tillio, Nenad Kos, and Matthias Messner. The design of ambiguous mechanisms. *mimeo*, 2012.
- Vassilios Dimitrakas and Yianis Sarafidis. Advice from an expert with unknown motives. 2005. SSRN.
- Wioletta Dziuda. Strategic argumentation. *Journal of Economic Theory*, 146:1362–1397, 2011.
- Françoise Forges and Frédéric Koessler. Communication equilibria with partially verifiable types. *Journal of Mathematical Economics*, 41(7):793–811, 2005.
- Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.
- Francesco Giovannoni and Daniel J. Seidmann. Secrecy, two-sided bias and the value of evidence. *Games and Economic Behavior*, 59(2):296–315, 2007.
- Jacob Glazer and Ariel Rubinstein. Debates and decisions: On a rationale of argumentation rules. *Games and Economic Behavior*, 36(2):158–173, 2001.
- Jacob Glazer and Ariel Rubinstein. On optimal rules of persuasion. *Econometrica*, 72(6):1715–1736, 2004.
- Jacob Glazer and Ariel Rubinstein. A study in the pragmatics of persuasion: a game theoretical approach. *Theoretical Economics*, 1:395–410, 2006.
- Jerry R. Green and Jean-Jacques Laffont. Partially verifiable information and mechanism design. *Review of Economic Studies*, 53(3):447–456, 1986.
- Sanford J. Grossman. The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics*, 24(3):461–483, 1981.

- Sergiu Hart, Ilan Kremer, and Motty Perry. Evidence games: Truth and commitment. *American Economic Review*, Forthcoming, 2016.
- Navin Kartik and Olivier Tercieux. Implementation with evidence. *Theoretical Economics*, 7(2):323–355, 2012.
- Frederic Koessler and Eduardo Perez-Richet. Evidence based mechanisms. *mimeo*, 2014.
- Matthias Lang and Achim Wambach. The fog of fraud – mitigating fraud by strategic ambiguity. *Games and Economic Behavior*, 81:255–275, 2013.
- Mark Le Quement. The (human) sampler’s curses. *American Economic Journal Micro (forthcoming)*, 2016.
- Ming Li. To disclose or not to disclose? cheap talk with uncertain biases. 2005. *mimeo*.
- Barton L. Lipman and Duane J. Seppi. Robust inference in communication games with partial provability. *Journal of Economic Theory*, 66:370–405, 1995.
- Giuseppe Lopomo, Luca Rigotti, and Chris Shannon. Uncertainty in mechanism design. *mimeo*, 2013.
- Jacob. Marschak and Roy Radner. *Economic theory of teams*. New Haven, Yale University Press, 1972.
- Jerome Mathis. Full revelation of information in sender–receiver games of persuasion. *Journal of Economic Theory*, 143(1):571–584, 2008.
- Steven Matthews and Andrew Postlewaite. Quality testing and disclosure. *The RAND Journal Economics*, 16(3):328–340, 1985.
- Paul Milgrom and John Roberts. Relying on the information of interested parties. *The RAND Journal Economics*, 17(1):18–32, 1986.

- Paul R. Milgrom. Good news and bad news: Representation theorems and applications. *Bell Journal of Economics*, 12(2):380–391, Autumn 1981.
- Roger B. Myerson. Optimal coordination mechanisms in generalized principal–agent problems. *Journal of Mathematical Economics*, 10(1):67–81, June 1982.
- Roger B. Myerson. *Game Theory, Analysis of Conflict*. Harvard University Press, 1991.
- Frank Riedel and Linda Sass. Ellsberg games. *Theory and Decision*, pages 1–41, 2013.
- Daniel J. Seidmann and Eyal Winter. Strategic information transmission with verifiable messages. *Econometrica*, 65(1):163–169, 1997.
- Itai Sher. Credibility and determinism in a game of persuasion. *Games and Economic Behavior*, 71:409–419, 2011.
- Itai Sher. Persuasion and dynamic communication. *Theoretical Economics*, 9:99–136, 2014.
- Itai Sher and Rakesh Vohra. Price discrimination through communication. *Theoretical Economics*, 10:597–648, 2015.
- Hyun Song Shin. News management and the value of firms. *The RAND Journal Economics*, 25(1):58–71, 1994.
- Nirvikar Singh and Donald Wittman. Implementation with partial verification. *Review of Economic Design*, 6(1):63–84, 2001.
- Marciano Siniscalchi. Dynamic choice under ambiguity. *Theoretical Economics*, 6(3):379–421, 2011.
- Roland Strausz. Mechanism design with partially verifiable information. *mimeo*, 2016.

Martin Walker. Management accounting and the economics of internal organization: a review essay. *Management Accounting Research*, 9:21–30, March 1998.

Asher Wolinsky. Information transmission when the sender's preferences are uncertain. *Games and Economic Behavior*, 42(6):319–326, 2003.

Information certifiable : révélation et persuasion

Mehdi Ayouni

Résumé

Dans de nombreuses situations, les décideurs sont amenés à choisir une action ou une politique sans être parfaitement informés. De plus, il est parfois difficile ou coûteux d'acquérir directement les informations manquantes. Dans ce cas, ils peuvent solliciter l'aide des institutions ou individus informés. Ces derniers peuvent essayer d'influencer la décision en leur faveur en cachant ou en ne présentant qu'une partie de l'information. Dans cette thèse, j'examine des situations où certaines informations sont certifiables ou vérifiables. En d'autres termes, un agent informé peut prouver certaines déclarations en présentant des preuves ou le décideur peut vérifier l'exactitude de ces déclarations. Puisque la vérification peut être coûteuse ou prenante, le décideur ne peut souvent vérifier qu'une partie de l'information reçue. Ces contraintes déterminent la quantité d'information qui peut être vérifiée avant la prise de décision. Les deux premiers chapitres portent sur des modèles adaptés aux situations où le décideur doit évaluer une déclaration ou répondre à une demande faite par un individu ou une institution. Dans le troisième chapitre, je considère un cadre légèrement différent où le décideur peut consulter des agents informés de manière séquentielle avant de choisir une action.

MOTS CLÉS : Information certifiable, Communication stratégique, Persuasion, Mécanismes de communication, Théorie des jeux.

Abstract

In many situations, decision makers do not observe all relevant information which undermines their ability to choose the best action or policy. Moreover, it can be difficult or costly to directly acquire the missing information. In such cases, the decision maker may acquire information from privately informed parties with potentially different objectives. The issue is that they may try to influence the outcome in their favor either by withholding or selectively reporting information. This thesis is focused on settings where at least some information is certifiable or verifiable. In other words, the informed party can prove certain statements by presenting hard evidence or the decision maker can verify the accuracy of received claims and documents. Since verification can be costly or time consuming, the decision maker might be able to only partially check the claim. These constraints determine the amount of information that can be verified before the decision is made. The first two chapters focus on models that best describe settings where the decision maker has to evaluate a claim or respond to a request made by an individual or institution. In the third chapter, I consider a slightly different framework where the decision maker may seek advice by sequentially consulting informed agents.

KEYWORDS: Certifiable information, Strategic communication, Persuasion, Communication mechanisms, Game theory.