Université Paris - Panthéon - Assas



École doctorale d'Économie et Gestion

Thèse de doctorat en Sciences économiques

Politique, jeux et communication stratégique

soutenue le 5 Décembre 2025

Emmanuel Lagrée

Sous la direction de Lucie Ménager

Membres du jury :
Sylvain Béal, Professeur à l'Université Marie et Louis Pasteur – rapporteur
Nicolas Klein, Professeur agrégé à l'Université de Montréal – rapporteur
Chantal Marlats, Professeure à l'Université de Tours – présidente du jury
Lucie Ménager, Professeure à l'Université Paris-Panthéon-Assas – directrice de thèse
Christoph Schottmüller, Professeur à l'Université de Cologne – examinateur

Avertissement

L'université n'entend donner aucune approbation ni improbation aux opinions émises dans cette thèse ; ces opinions doivent être considérées comme propres à leur auteur.

À Juliette

Remerciements

L'université française a ceci de merveilleux qu'elle nous laisse faire à peu près ce qu'on veut. À l'issue de mon master, j'avais donc eu la chance d'étudier les mathématiques appliquées (surtout), la philosophie et les sciences cognitives. Le désir de faire une thèse bien ancré, il a été nécessaire de partir à la recherche d'un sujet qui pût mêler ces trois disciplines disparates. La recherche a alors ceci de merveilleux qu'elle résout aisément cette quadrature du cercle : cela a même un nom, la théorie des jeux.

S'inscire en thèse en théorie des jeux supposa alors de choisir une discipline de rattachement : l'économie, les mathématiques ou l'informatique. L'économie était tout indiquée car elle seule se situe réellement à l'intersection de mes divers intérêts. C'est ici que Lucie entra en scène : elle ne la quitta plus. C'est au cours de la thèse que des personnages qui n'ont définitivement rien de secondaire sont apparus pour me prêter une main secourable. L'occasion est donnée de les remercier.

Lucie, donc, et toutes tes qualités. Tu auras été ma guide tout au long de ces trois années et demi, mémoire de master compris. Ce n'a pas été une mince affaire. En effet, il fallait m'instruire en théorie des jeux mais également en économie car je n'y connaissais presque rien : cela faisait désordre. Je te remercie pour m'avoir aiguillé par petites touches, sans jamais rien m'imposer. Cela m'a permis de construire mon sujet de thèse, puis mes articles, au gré de mes lectures. Cette liberté est inestimable. Tu as su également me remettre dans le droit chemin quand mes idées étaient mauvaises (fréquemment) et tu m'as initié aux secrets de la modélisation. J'ai appris à tes côtés que le plus difficile n'est pas systématiquement de résoudre le modèle mais plutôt de bien le poser. Merci d'avoir adressé tes conseils avec patience et enthousiasme : deux qualités qu'on apprécie toujours chez les autres quand rien ne va chez soi. Aussi, et cela compte autant que le reste, nous avons bien ri.

Ensuite, merci Christoph de m'avoir accueilli à l'université de Cologne pour le semestre

d'hiver de ma troisième année. C'était une période fabuleuse du point de vue de la recherche comme de l'expérience. Grâce à toi, Cologne est désormais ma ville allemande préférée. Merci pour toute l'aide que tu m'as apportée sur mon deuxième article ; et pour m'avoir gracieusement prêté ton modèle. J'espère que les modifications que j'ai apportées sont pour le meilleur, j'ose croire qu'elles ne sont pas pour le pire. C'est un plaisir de t'avoir dans mon jury. *Dankeschön!*

Merci Sylvain de m'avoir invité à un *workshop* sur la théorie des jeux coopératifs à Besançon. Cette rencontre au début de ma thèse a été riche de conséquences. Grâce à toi, j'ai eu l'occasion de découvrir une belle et inconnue littérature. J'ai pu discuter avec toi de mes orientations possibles de recherche et, finalement, essayer un modèle qui mêle coopération et bayésianisme sous ton regard bienveillant. Il était donc tout naturel de te demander de faire partie de mon jury et je te remercie d'avoir accepté.

J'adresse un grand merci à mes deux autres membres du jury. Nicolas tout d'abord dont j'ai suivi le cours doctoral d'expérimentation stratégique il y a maintenant plus de deux ans. Cette même expérimentation stratégique forme le coeur de mon troisième (projet d') article : j'aurai pris mon temps mais j'y serai venu. Chantal ensuite, qui a joué un rôle bien plus déterminant que le seul titre d'examinatrice ne le laisserait supposer. Tu as suivi toute ma thèse aux côtés de Lucie et tu as fait preuve de la même patience et du même enthousiasme. Merci pour tout !

Bien qu'elle soit largement un exercice de l'esprit, une thèse ne saurait se faire sans quelques dispositions matérielles et financières. À ce titre, je tiens à remercier Antoine Billot pour m'avoir accepté au Lemma – pour y trouver un bureau, des séminaires, etc. Par extension, un grand merci à l'école doctorale EGIC. C'est du fond du coeur que je te remercie, Laurence, pour ta disponibilité et ta capacité à aplanir toutes les difficultés. En particulier, puisque la passion de la papeterie est la chose la mieux partagée du monde (des doctorants), merci d'avoir su répondre à nos envies les plus précises ou les plus fantasques.

J'espère avoir rempli efficacement ordres de mission et autres papiers de sorte que tu n'as eu qu'à les avaliser. Nous avons aussi parlé de natation et de chats, merci. Merci aux appariteurs et au personnel de nettoyage dont la contribution est transparente : rendre nos vies si simples. Quant au nerf de la guerre, un grand merci au Labex MME-DII et au *Deutscher Akademischer Austauschdienst*.

Derniers adultes que je remercie : Victor et José pour avoir formé mon comité de suivi individuel. Ce CSI n'a eu lieu qu'une seule fois (record en cours, je pense), mais quelle fois ! J'ai réalisé à quel point il est salutaire de présenter ses travaux à des chercheurs dont ce n'est pas la spécialité.

Viennent désormais les camarades. Redha, Ève et Lucas ont été indispensables, du premier au dernier jour. Ensemble, nous avons : débattu des mérites respectifs de la théorie et des données, bu du thé, appris à enseigner des travaux dirigés, bu (mais pas du thé), arpenté Paris, pesté contre le système des revues, évalué les mérites d'une organisation mafieuse de la société, vécu. Tom, Kieran et Vincent forment une belle relève. À l'inverse, merci à Chloé, Théo et François qui m'ont initié au métier de doctorant avant de voguer vers d'autres cieux.

Merci à mes amis *aus Köln*, Matteo et Rastislav, avec qui j'ai débattu plus ou moins des mêmes sujets, souvent en anglais, malheureusement plus rarement en allemand.

Ensuite, avec quelques autres doctorants, nous avons essayé de relancer le séminaire junior de théorie des jeux de Paris. Nous le faisons tourner depuis deux ans. Je n'ai aucun doute que cela tiendra dans le temps. Merci à vous : David, Julian et tous les autres.

Merci à Pierre pour avoir vécu les mêmes affres que moi, en même temps, depuis les tout premiers cours magistraux de philosophie en L1. Nous pouvons nous féliciter de notre parallélisme. Soutenant quatre jours avant toi, je serai à jamais le premier, mes excuses.

Pour tout le reste, merci Juliette.

Résumé : Politique, jeux et communication stratégique. Cette thèse étudie la transmission stratégique d'information dans des jeux de science politique. Elle consiste en trois parties indépendantes. La première partie est consacrée à l'étude de la coopération mal informée qui apparaît dans les projets publics à grande échelle. Je développe un modèle de jeux coalitionnels simples où les joueurs connaissent pour sûr uniquement les valeurs des coalitions auxquelles ils appartiennent. Chaque coalition peut dévoiler sa valeur à un planificateur social si ses membres le décident à l'unanimité. Le planificateur social met ensuite à jour ses croyances à propos du véritable jeu coalitionnel et alloue les paiements en fonction. Je caractérise sous quelles conditions il est optimal de certifier toutes et seulement les coalitions gagnantes. Dans la seconde partie, j'étudie un modèle de communication facile (*cheap talk*) répétée entre des experts et un parlementaire. Les experts ont chacun un biais idéologique et sont inégalement compétents. Ils peuvent être licenciés librement s'ils sont insatisfaisants. Le parlementaire souhaite identifier les experts compétents et qu'ils le conseillent sincèrement. Je montre que la compétition idéologique, c'est-à-dire avoir des experts avec des biais opposés, est bénéfique au Parlement seulement quand (i) la compétence des experts est une information privée et (ii) même les experts compétents peuvent se tromper. La troisième partie est préliminaire. Je définis un modèle d'expérimentation stratégique où des politiciens font face à la possibilité d'une catastrophe naturelle. Ils disposent de trois actions : construction de protection privée, conduite de recherche publique ou attentisme (sans coût). Je résous actuellement le problème de décision en stratégies markoviennes.

Descripteurs : Jeux répétés, Jeux coalitionnels, Communication stratégique, Expérimentation stratégique, Bayésianisme.

Abstract: Politics, games and strategic communication. This thesis studies the strategic transmission of information in political science games. It consists in three independent parts. The first part is devoted to the study of misinformed cooperation as found in public large-scale projects. I develop a model of simple coalition games where players know for sure only the worth of the coalitions they belong to. Each coalition may disclose its worth to a social planner should its members unanimously agree. The social planner then updates her belief about the coalition game at play and allocates payments accordingly. I characterise under which conditions it is optimal to certify all and only winning coalitions. In the second part, I study a model of repeated cheap talk between experts and a Member of Parliament. The experts have each an ideological bias and are unequally competent. They can be fired freely should they be unsatisfactory. The Member of Parliament wants to identify competent experts and to have them advise sincerely. I prove that ideological competition, namely having experts with opposed biases, is beneficiary to the Parliament only when (i) the competence of experts is private information and (ii) even competent experts may be wrong. The third part is at an early stage. I develop a model of strategic experimentation where politicians face the possibility of a natural disaster. They must choose between three actions: they may build private protection, conduct public research or wait (costless). I currently solve the decision problem in Markovian strategies.

Keywords: Repeated games, Coalitional games, Strategic communication, Strategic experimentation, Bayesianism.

Main acronyms and abbreviations

- MP: Member of Parliament
- p.: page
- PDE: partial differential equation
- Prop.: Proposition
- resp.: respectively
- SSM: strictly strongly monotonic
- s.t.: such that
- TU-game: coalitional game with transferable utility

Contents

1 Introduction **11**

1.1 Game theory and Public Action 11

1.1.1 The undertaking of large-scale projects and the problem of misinformed cooperation 12

1.1.2 Members of Parliament, experts and the problem of sincerity 16

1.1.3 Natural disasters, the precaution principle and the problem of resource waste 19

1.2 Theoretical background 22

1.2.1 A brief overview of economics of information 23

1.2.2 Games in normal form 25

1.2.3 Informational asymmetries and Bayesian games 27

1.2.4 Repeated games 28

1.2.5 Coalitional Games with transferable utility 33

2 Certification of Coalitions in Games of Imperfect Information **37**

2.1 Introduction 37

2.2 A Two-stage Game of Certification 43

2.3 Equilibrium Analysis 46

2.4 Extensions 57

2.5 Proofs 60

3 Ideological Competition and the Taming of Experts **69**

3.1 Introduction . . . 69

3.2 The model . . . 75

3.3 First results . . . 77

3.4 Is competition beneficial among experts? . . . 79

3.5 Equilibrium with constant payments for D . . . 84

3.6 Extensions and concluding remarks . . . 85

3.7 Proofs . . . 88

3.8 An alternative model: Multiple states of the world . . . 103

4 Strategic Experimentation and Protection **109**

4.1 Introduction . . . 109

4.2 Model . . . 113

4.3 The Bellman equation . . . 115

4.4 Contradictions and the absence of smooth pasting. . . . 118

5 Conclusion **121**

List of Tables **127**

Appendice **127**

A Résumé en français **129**

A.1 Introduction . . . 129

A.2 La certification de coalitions dans les jeux à information imparfaite . . . 131

A.3 Compétition idéologique et domestication des experts . . . 136

A.4 Expérimentation stratégique et protection . . . 142

Bibliography **146**

1 Introduction

1.1 Game theory and Public Action

Game theory is the study of interactive decisions between agents. It rests on a simple idea. A decision is interactive whenever the payment of an agent depends on the decisions made by others (and not only on hers, or chance). Such a situation can be modelled as a *game* between *players* that are (instrumentally) rational. Instrumental rationality means that players optimise their choices given both their knowledge and their constraints.

Game theory lies at the intersection of applied mathematics and social sciences. It has emerged in the 1940s with the works of Von Neumann and Morgenstern (1944) and Nash (1950b) and has blossomed since then. It is now both a field of research *per se* and a toolbox for other disciplines (e.g. political science or biology). Game theory plays a prominent role in economics: the framework of interacting agents driven by their interests (and methodological individualism) fits the core axioms of the discipline.

The key concept to study and solve a game is the *Nash equilibrium* (Nash, 1950b).[1] A Nash equilibrium is a course of strategies, one per player, such that no player prefers to change their decision unilaterally. It is proven to exist under mild conditions and it provides a benchmark for decision makers and agents as to which courses of actions are possible, stable and expected to

[1] It is *the* solution concept. In his historical account, Myerson (1999) argues that "the formulation ... has had a fundamental and pervasive impact in economics and the social sciences which is comparable to that of the discovery of the DNA double helix in the biological sciences." (p. 1067) Buchanan (2001) argues alike. For different opinions, however, see Misyak and Chater (2014) who argue that players conduct "virtual bargaining" instead of playing equilibria and Rosenthal (1981) who advocates for decision-theoretic tools when players lack information.

emerge from strategic players.

Since the finite one-shot games under full information studied by Nash, game theory has investigated several paths of research–e.g. partial information, dynamic setups, notions of co-operation, strategic communication. New concepts beyond the standard Nash equilibrium have been developed accordingly.

This thesis is driven by two fields of interest. On the one hand, it investigates issues of political science: broadly, it studies the decisions of a misinformed public decision maker and their interactions with better informed individuals. On the other hand, it offers game-theoretic answers to these issues: it relies on the economics of asymmetric information and the strategic transmission of information.

Over the years of my PhD thesis, I developed a keen interest in public action. I therefore formulate the narratives in framework that stage public actors and institutions (the State, municipalities, Members of Parliament...) and identify their preferred outcomes that consist in equilibria.[2] Nevertheless, the strength of a model lies in its ability to account for several situations–be they "private" or "public"–so long as the players genuinely have the preferences and incentives that the model ascribes. The scope of the models I develop is hence wider than the narrative they relate to. The remaining of the section describes the questions that I tackle and the political science issues I am interested in. It justifies the undertaking of game-theoretic solutions. A description of the results follows.

1.1.1 The undertaking of large-scale projects and the problem of misinformed cooperation

Designing and achieving large-scale public projects is notoriously difficult. They seem to suffer a manifold of issues: delays, additional costs, outright cancellations. Such issues can have, broadly,

[2]The theory of *mechanism design* provides tools to answer such questions. It studies what rules should be implemented to achieve a good result according to some predefined metric. It is the reverse problem of game theory–which takes rules as a given and study the behaviour of players. I focus on game theory in its strict acceptance.

two orthogonal explanations. On the one hand, a project can be difficult to carry on due to material or "concret" characteristics: complex engineering, exogenous events (Covid 19...), unstable technology, etc. On the other hand, there can be institutional characteristics: disagreements between several contractors, multiplicity of sources of funding, diluted responsibility, etc.

To the former explanation belongs for instance the construction of the European Pressurized Reactor of Flamanville (Manche, France). The construction started in 2007 and was supposed to end in 2012 (for a previsional cost of 3.4 billion euros). The commission date was finally 2024 for an estimated cost of 23.7 billion euros (Mercereau, 2025). Likewise "High Speed 2", a major high-speed railway project in England, aimed at connecting London to northern England *via* Birmingham. The estimated costs of the project exploded (from 36 billion in 2009 to 106 billion euros today, see Davies, 2024). It motivated the British government to first cancel the branch from Birmingham to Leeds and Sheffield, and then to cancel the branch from Birmingham to Liverpool and Manchester.[3]

To the latter explanation belong countless projects that involve several heterogeneous bodies, some that have political priority while others know better the preferences of the citizens. For instance, water supply in France is under the jurisdiction of the communes. They can form groups should they opt to. The State (represented by the "*préfet*") is a facilitator and may take the upper hand should a commune fail to its duty. There are five administrative levels that may play a role and five different sources of funding accordingly. These are: communes, *établissements publics de coopération intercommunale* (groups of communes), *départements*, *régions*, the State. The levels range from the best to the least informed of the needs of the local inhabitants and from the least to the most powerful. (In the sequel, I consider only communes, groups of communes and the State. I rule out *départements* and *régions*. The European Union may also finance investments but is left aside.)

In this maze of administrative powers, who does what? And who wants what? The explicit delimitation of responsibilities is determined by the law. Besides, the preferences and the pos-

[3]This project is also impeded by institutional disagreements between the British government and both the city of Liverpool and Scotland but this is left aside.

sibilities remain often hidden to some parties. Communes stand the closest to the citizens and know best how to supply water to them. In contrast, the State is remote and its sole purpose is therefore to fulfill the needs, all other things being equal, of the communes. It is benevolent and supports the communes fairly.

Very often communes form groups to provide water: it is more efficient and cheaper. In particular, it is indicated when communes are not "big" or rich enough to supply water on their own but are able, as a group, to do so. A first difficulty arises when two communes disagree on the course of actions to take. Indeed, they are equally powerful legal- or administrative-wise, one cannot overcome the other; the (neutral) State has no reason to favour one to the detriment of the other.

Also, when a commune cooperates with two peers, it needs not know perfectly what the abilities of the latter communes are. It may know, say, that the three of them are able (as a group) to build a water treatment plant and to supply water to their citizens and that associated with only one of the two peers, it cannot. It may, however, ignore whether the two latter by themselves are able to build the plant. In other words, each commune has limited knowledge about the possibilities of cooperation of its peers. In principle, the State lacks all the more information.

This raises the following questions: What groups should communes form? How do they share the costs of the project within the group? How can the State efficiently and fairly support the communes? Put in a more abstract manner, how can entities that cooperate share the dividends of cooperation when there is uncertainty about their respective abilities? How can the social planner best adjudicate between needs and resources of the entities? How can information be revealed throughout the process of cooperating–in particular to the social planner? What is a fair share between misinformed cooperating parties?

Chapter 2 is entitled "Certification of coalitions in games of imperfect information." It develops a model of coalitional games under asymmetric information and provides answers to the aforementioned questions. Cooperative game theory studies games where enforcing contracts and agreements can be signed. Coalitional games form a subpart thereof. A coalitional game repre-

sents what each coalition (subgroup of the grand population) can earn, should all the members of the coalition cooperate and act together.

Given a coalitional game, there are two main questions: Which partition of the grand population forms? How do players share the fruits of cooperation within a coalition? Several solution concepts have been developed to tackle this problem.

Usually, the coalitional game that is played is common knowledge: everyone knows what the game at play is, knows that everyone ones... infinitely so.[4] The purpose of Chapter 2 is to remove this assumption and to assume that players know the worth of the coalitions they belong to and not necessarily the worth of others.

I define a strategic Bayesian game where players possess this information and the underlying coalitional game is *simple* (coalitions are either winning or losing). They have the opportunity to *certify* the worth of the coalitions they belong to if they unanimously decide to (within each coalition). Afterwards, a social planner updates her belief about the coalitional game according to Bayes' rule, forms the grand coalition and allocates the dividends of cooperation according to some predetermined solution concept. The purpose of a player is solely to shape the beliefs of the social planner to his advantage.

I show that it is not necessarily a perfect Bayesian equilibrium to certify all and only the winning coalitions. Sometimes, players prefer to "hide" the worth of a good coalition because it makes them better off (that is, with respect to the other coalitions). I then consider specific priors under which all and only winning coalitions are certified and explain why the relevant coalitions to look at are the so called *minimal winning*.

This chapter contributes to the literature in two ways. First it defines a simple model of informational asymmetry in coalitional games. Second it shows that players may have incentive to "hide" good news and gives sufficient conditions where this behaviour does not arise. The model, however, stumbles on the curse of dimensionality. If there are n players, there are 2^n

[4]Lewis (1969) provides the first formal definition of common knowledge. Aumann (1976) rephrases the infinite regression model of common knowledge into static partitions. Ménager (2023) surveys this literature. For a reasonable weakening of the infinite epistemic conditions, see Paternotte (2011); for a probabilistic account thereof, Paternotte (2017).

coalitions and so 2^{2^n} simple games are possible from the perspective of the social planner.[5]

1.1.2 Members of Parliament, experts and the problem of sincerity

Our societies rely on technical and scientific knowledge of increasing complexity. In the course of policy-making, Members of Parliament (MPs) must therefore gain and acquire precise knowledge to vote suited bills. They face a difficulty. On the one hand, their decisions should be as well-informed (or as strongly evidence-based) as possible. On the other hand, they cannot develop personal expertise on every topic.

One way to avoid this issue is to have each MP specialise in a different field. The drawback is that only some MPs turn out to be "relevant" on specific topics: this weakens the deliberating nature of a Parliament. Another way is to rely on external expertise provided by groups of interest, think tanks or academics. Trade unionists and representatives of employers have insights that are valuable to the MPs. There is then an epistemic delegation between MPs and experts: the former vote after listening to technical recommendations and opinions of the latter.[6]

This need for information by the MPs explains the increasing importance given to experts. For example, around $3,000$ hearings are conducted by the specialized commissions of the French *Assemblée nationale* every term (part of which, it is true, have a different purpose, like making the Government accountable). Likewise, the European Parliament commissions about one thousand technical reports per year to external experts and praises itself for its link with academia (Maciejewski, 2018).

This major phenomenon within policy-making calls for norms and good practices and requires

[5]Technically, it makes slightly less than 2^{2^n} because the empty set has worth zero and I impose the grand coalition to have worth one. Yet it is immaterial to the computational issue. To fix ideas, if there are 8 players, that makes around 2^{256} possibilities. This is similar to the estimated number of atoms in the observable universe (10^{80}, the Eddington number). If there are 9 players, there are $2^{512} = 2^{256} \times 2^{256}$ possibilities, namely 10^{80} different universes that each contain 10^{80} atoms. Obviously too many.

[6]The European Parliament adopts a mixed approach to this issue. There is division of cognitive labour between MEPs who all develop specialized skills. They nonetheless extensively rely on external expertise. Elomäki and Haapala (2024) describes this mixed approach and explains how the notion of expertise itself is *constructed* endogenously.

strict rules to fulfill them. Indeed, value neutrality of the experts is not to be taken for granted. A lobbyist of an employer federation possesses knowledge about the economic and labour system that would be useful for the MPs. When they promote lower business taxes, it may be the case that they believe it will ultimately galvanize the labour market to the profit of the whole society. Yet it may be the case that they push their own agenda regardless of the preferences of the MPs. Symmetrically, interpreting a trade unionist that advocates for a higher minimum wage is not straightforward.

All in all, it seems there must be transparency, balance between different views and opinions, sincerity from the experts. Precise rules are designed to enforce the norms.

As an example, take the case of the *Assemblée nationale*. Most of expertise relationships take place in specialized committees rather than in plenary sessions. Summoned experts are compelled to show up as they face prosecutions if they do not. They are invited either by the president of the committee (who is in general a member of the majority) or by its vice-president (who is a member of the opposition): this way, representativeness of the whole political spectrum is ensured. There is *ideological competition* between experts. (Implicit here is the fact that there are only two parties. In multi-partite systems like in the European Parliament, satisfaction of the preferences of smaller parties is also ensured.)

One can identify the political preferences of the experts easily: one knows what the lobbyist or the trade unionist prefer for themselves. Likewise, the preferences of a Marxist economist or a fellow of the School of Chicago are well known.

Experts must tell the truth. They face prosecutions if it is proven they lied. However intentional deception is hard to establish because one cannot be sure, in general, whether bad advice is due to insincerity or incompetence. Even the best experts turn out to be wrong here and there.

To hedge against both incompetence and insincerity, experts are not hired on a fixed term contract. They can be dismissed at any time should they prove unsatisfactory. If so, a dismissed expert is replaced by a peer that has aligned political preferences. In spite of the precariousness of their positions, experts usually maintain themselves in place and repeatedly provide advice to

the Members of Parliament.

One can henceforth question whether MPs efficiently discipline experts to tell the truth and identify quickly those that are either incompetent or insincere. There are two channels through which the Members of Parliament can discipline the experts. First experts are disciplined by the fear of dismissal and replacement by an ideological peer. Second experts are disciplined by their present counterparts: a trade unionist cannot push to much their agenda because if they do, the lobbyist in front easily contradicts them.

Chapter 3 is entitled "Ideological competition and the taming of experts." It assesses the efficacy of the rules to have the experts be sincere. In particular, it tackles the following questions: Is ideological competition beneficial to the Members of Parliament? Or is the "peer channel" sufficient to have experts be sincere? To that end, I develop a model of repeated cheap talk with transparent motives but private competence. To simplify the problem, I shall assume that the politically-driven Members of Parliament can be averaged as a representative decision maker who has no ideological bias.

The model is a repeated game of *cheap talk* between a long-lived decision maker and experts. It relates to the literature of repeated games and strategic communication. A repeated game is the repetition of a (fixed) stage game throughout time. Time is discrete. Players seek to maximise their discounted flow of payment. Cheap talk messages have three properties: they are unverifiable *ex interim*, costless and non binding. The model of this chapter builds on and extends Schottmüller (2019) but examines different questions and purposes.

In the model, at every period a decision maker must take an action. Her payment depends on an ignored state of the world that changes at every period. She asks for advice to experts who receive signals about the state. An expert has preferences that are independent from the state of the world and may be either competent or not. His preferences are common knowledge and independent of the state of the world (so called transparent motives). His competence is hidden information. At any time, the decision maker may fire an expert and a new one with the same preferences steps in. Experts have dynamic incentives to tell the truth and are therefore disciplined by the threat of being replaced by peers.

There is ideological competition when the decision maker interviews two experts with opposed preferences simultaneously. One can expect that direct ideological competition is a channel of discipline. The question I tackle is the following: Does ideological competition improve the expected payment of the decision maker? Or is the threat to fire and replace experts sufficient to ensure high payments to the decision maker?

Players use *quasi-markovian* strategies: a strategy is a function of the *reputation* of the expert[7] and of the outcome of the current period. I show that ideological competition is *not* needed provided either there is no hidden information–there are only competent experts–, or competent experts are perfect–their signals are direct observations of the state of the world.

Indeed, even without competition, perfect sincerity for the experts is sustainable in (subgame perfect or markovian) equilibrium. In the former case, this holds true because firing an expert is done at zero cost as the ones that step in are known to have the same competence: firing is credible. In the latter case, this holds because whenever an expert sends a wrong message in the "sincerity profile," his reputation falls immediately to zero as no mistake happens. Such experts with zero reputation are optimally dismissed.

When there are both hidden information and imperfect competence, however, for every equilibrium without ideological competition, there exists an equilibrium with competition where the expected payment of the decision maker is *strictly* higher. It goes through two channels: one direct, one indirect.

Chapter 3 then considers extensions of the model and an alternative that I developped in an earlier stage of the project.

1.1.3 Natural disasters, the precaution principle and the problem of resource waste

How best to protect society against natural disasters is a notoriously long-lived problem. It has received renewed attention due to the phenomena of global warming and climatic instability.

[7]The subjective probability that he is competent. It is the state variable.

Natural disasters differ qualitatively with regular private or public accidents for at least three reasons: they are massive, contain an element of irreversibility, and are very hard to predict. The lack of prediction comes from their scarceness: they happen rarely enough not to have reliable statistics. That makes them hardly insurable.

The only reasonable course of action against an incoming natural disaster is consequently to protect oneself.

Recent floods in Europe have put the topic on top of preoccupations. On July 14th and 15th 2021, a major flood caused the death of 269 people in Western Europe (mainly in North-Rhine-Westphalia, the Netherlands and Belgium). On October 29th and 30th 2024, 225 people died from a flood in the region of Valencia (Spain). Such floods caused major material damage (e.g. estimated to around 10 billion euros in Spain) that are insurable. Death casualties, though, are irreversible.

Both disasters caused a public uproar and unveiled many failures. For instance, the flood in Spain highlighted the failures of the emergency system and the dilution of responsibilities between the federal government and local authorities (here, the region of Valencia). Beyond such specific criticisms addressed to the government(s), questions boil down to two sets of interrogations: (i) How come were the public authorities unable to predict the occurrence of the flood? and (ii) How come were there no facilities ready to prevent or mitigate the damage of the flood?

The first question addresses the issue of (lack of) information. The acquisition of information about the natural disaster–if and when it shall occur–is deficient. The second question addresses the issue of (lack of) protection against the disaster should it occur. In other words, why cannot the authorities know if and when there will be a flood? And should there be a flood today, why is there no protection against it?

Yet the occurrence of the disaster is uncertain and building protection can prove wasteful *ex post*. The principle of precaution suggests to build protection nonetheless, not in spite of missing information but because so. (It is to be opposed to a cost-benefit analysis.) Budget or time constraints and the multiplicity of orthogonal issues suggests not to over-invest in highly

hypothetical risks.[8]

There are therefore three possible actions in face of a conceivable but highly hypothetical natural disaster: (i) build protection in spite of a possible waste *ex post*; (ii) try to lift uncertainty through research (but the detection of low signals is hard); (iii) adopt a wait-and-see attitude (to presumably work on other issues).

This framework can be generalized to several public authorities. Let us go back to the 2021 flood in Western Europe. Three parties (Germany, Belgium and the Netherlands) are concerned. To a large extent protection is private: when Belgium constructs a dam, its effect benefits solely to Belgium. On the other hand, research is public: once Belgium knows that the Rhine valley is subject to massive floods in the near future that justify building dams, Germany and the Netherlands are likewise aware.

The asymmetry between private protection[9] and public research, and the possibility for one authority to free ride on research expenditures of others make the situation a strategic game.

Chapter 4 is entitled "Strategic experimentation and protection." It formalizes this dilemma. The model is a game of *strategic experimentation* with three exponential bandits, namely three arms. Bandit games usually stage players who must choose between two actions (or arms) in continuous time. One arm is safe and provides a known flow of utility. The other arm is risky and can be either good or bad. The quality of the risky arm is unknown at first and can only be learned through trials.

I define a model where a disliked event (a natural disaster, say) may or may not happen in the future. If yes, its arrival is exponentially distributed and is independent of the players' strategies. The occurrence of the event ends the game and final payments are distributed. A player can build protection (first arm) against the event (following a precautionary principle) but it can be

[8]This dilemma is well pictured in 2018 by Édouard Philippe, then Prime minister of France. In response to the resignation of his minister for environment and his supposedly lack of activity against global warming, he highlights than there are other issues difficult to anticipate and equally important. Such issues one could reproach him for within ten years: he mentions the possibility of an unknown world-wide virus and Russia invading Poland... See < https://www.dailymotion.com/video/x82gmpz >.

[9]Obviously protection cannot be completely private. In particular, the construction of a dam benefits everyone downstream, no one upstream. It suggests that protection is private for Belgium but public for Germany. This enriching of the story (and the model accordingly) is not considered here. (Contrariwise, cities around a volcano presumably build purely private protection.)

a waste *ex post* if the event never happens. A player can also conduct research to learn the *type* of the event (second arm) to better tailor his future protection strategy. Finally, a player can wait (third arm). The first two actions are costly, waiting is done at no cost.

I focus on markovian strategies where the state variable is the tuple (subjective probability of the type, current stock of protection). In most of the existing literature, playing the safe arm freezes the updating of the subjective probability: it does not bring additional information. In contrast, in my model players learn also when they do not experiment. Indeed, suppose that the event has not happened so far. It may be because "times has not come yet" but it can also be because it will never happen. Therefore, the mere fact that the event does not occur decreases the subjective probability (the Bayesian updating function is standard, conducting research speeds up the process).

Besides, there is no notion of myopic threshold below which playing safe is optimal and above which playing risky is optimal in the short term. Indeed, neither conducting research nor building protection is worth it in the short term because these actions are oriented towards the future and their present gains are outweighed by their costs.

This project is at an early stage and I am busy with solving the decision problem: I wish to find the optimal course of action of a politician as a function of their belief that the natural disaster *can* occur in the future. Afterwards, I plan to extend the model to several players in order to measure the free rider effect and, if any, the encouragement effect.

1.2 Theoretical background

This section first sketches a brief overview of economics of information. It then defines the main concepts of game theory that will be useful throughout the thesis. The concepts fall into four parts: (i) games in strategic form, (ii) informational asymmetries and Bayesian games, (iii) repeated games, and (iv) coalitional games.

1.2.1 A brief overview of economics of information

The three chapters all belong to the extensive literature on *economics of information*. They share a common theoretical interest in the strategic transmission of information and, for the first two, asymmetries of information. Nevertheless they define different models. I briefly sketch the literature on economics of information and explain how the three articles relate to it. In particular, I outline the differences in how information is strategically transmitted or communicated.

A salient article to date the birth of economics of information is Hayek (1945). Hayek explains that prices reflect the scarcity of goods. The price mechanism permits decentralized information to emerge (Stigler (1961) pursues this strand of research: he investigates the ascertainment of market price). However, many real-life situations involve uncertainties that cannot be mediated by a price mechanism (and the market).

Provided there are asymmetries of information, namely different agents have different information, two famous problems may arise: adverse selection and moral hazard. Adverse selection happens before a transaction is made, moral hazard happens after.

Adverse selection is pictured in Akerlof (1970). Akerlof shows that when buyers are not able to distinguish between low- and high-quality products, they tend to pay no more than their expected valuation which is less that the high valuation. This leads sellers of high quality product to leave the market. All in all, this selection kills the market. For the market to exist, there need be certificates that allow high-quality products to be identified.

Two solutions were first proposed to avoid adverse selection: screening (Stiglitz, 1975) and signalling (Spence, 1973). Screening consists in offering a menu of contracts such that two different types (say, good (or skilled) and bad job applicants) will choose two different contracts. (There, the asymmetry of information is solved. One can consider partial screening.)

Signalling assumes that the informed party may take actions that credibly reveals their type. Spence (1973) takes the example of job application and higher education. Both a good and a bad applicant pretend that they are good. Then getting a degree helps a good applicant to reveal their type because the cost of higher education (in terms of effort) is lower for good applicants.

This highlights the following point: an action that tries to convey information should credibly

do so. This thesis studies by way of three models the strategic transmission of information.

A *message* is an action that transmits information. It can be a "hard" action as the choice of education in the signalling literature. It can be the simple utterance of a word (which fears being not credible) from a sender to a receiver. The rules of the model determine three characteristics of a message:

- A message can be costly or costless to transmit or receive;
- A message can be verifiable or not (by the receiver or a third party like a court);[10]
- A message can be binding or not (that is, it does not limit strategic choice to either party).

Verifiability means that the sender of a message do not lie as it cannot be profitable. Indeed if they did, they would be detected and lies are not credible. In some models and circumstances, lies can even be counterproductive as they are harshly punished (by prosecution, say). In many instances, certifiability by the sender and verifiability by the receiver are equivalent sincerity-wise: in both cases, a sender cannot lie. (If messages are certifiable, then a message that is not certified is *not* credible and it is as if the receiver could verify its credibility (Grossman, 1981).)

For instance, the signalling game in Spence (1973) stages costly, verifiable and non binding messages.

Chapter 2 stages a game of group certification that requires the unanimity of its members (for the certification to occur). Players can either tell the truth or stay silent. Messages are costless but are verifiable and binding. The binding characteristic comes from the fact that, whenever a member (of the group) decides not to certify, it prevents all the others to certify: their strategies are limited by the former.

The least demanding strategic communication is cheap talk (pioneered by Crawford and Sobel, 1982): messages are costless, unverifiable and non-binding (in a static setting). The sender is free, in principle, to lie as much as they want. In equilibrium, it may be optimal, however, to reveal part or all of the information they possess. It is the realm of Chapter 3.

Finally, strategic experimentation presents a different information framework: (i) there is not

[10]When the game is repeated, one should distinguish between verifiability *ex interim* and *ex post*. In Chapter 3, messages are not verifiable *ex interim* (that is, before the receiver takes an action) but are verifiable *ex post* (that is, between today and tomorrow).

asymmetric information–players face the same uncertainty; (ii) there is no plain communication but rather information is transmitted *via* hard actions. Experimenting is costly and verifiable (players observe perfectly one another's actions).[11] It is not binding: experimenting or not today does not preclude strategies in the future. It is the realm of Chapter 4.

To sum up, cheap talk puts less restrictions on communication or transmission of information than both certification and experimentation. The latter have their own orthogonal conditions. The remaining of the section defines formally the necessary tools.

1.2.2 Games in normal form

A *game* in *strategic* or *normal form* is a triple $(N, (A_i)_{i\in N}, (u_i)_{i\in N})$ where N is the finite and non empty set of players, A_i is the non empty set of actions of player $i \in N$ and $u_i : \prod_{j\in N} A_j \to \mathbb{R}$ is the payoff function. If $a = (a_j)_{j\in N}$ is the *profile* of actions, then $u_i(a)$ is the payment of player i. It can be monetary. Following the seminal book by Von Neumann and Morgenstern (1944), u_i can also be interpreted as the utility function of i. When all the A_i are finite, say that the game is finite. Players choose their actions simultaneously and independently.

Example 1.1. The most famous game is the Prisoner's Dilemma. Define this game as a team project where two players can either exert effort (E) or shirk (S). If both players exert efforts, they each earn a high payoff. If they both shirk, their payments are mediocre. The best for one player is shirk while the other exerts effort. The payoffs are given by Table 1.1.

		Player 1	
		E	S
Player 2	E	$(3,3)$	$(0,4)$
	S	$(4,0)$	$(1,1)$

Table 1.1: Matrix of payoffs of the Prisoner's Dilemma.

Write $-i$ instead of $N \setminus \{i\}$, $A_{-i} = \prod_{j\neq i} A_j$ and $a_{-i} = (a_1, ..., a_{i-1}, a_{i+i}, ..., a_n)$.

[11] Compare this with the notion of perfect monitoring in repeated games.

a_i is a *best response* to a_{-i} for player i if $u_i(a_i, a_{-i}) \geq u_i(b_i, a_{-i})$ for all $b_i \in A_i$: against a_{-i}, i cannot do strictly better than to play a_i. This allows us to define the *best response correspondence* of player i as the multivalued function $BR_i : A_{-i} \to A_i$ defined for all a_{-i} by $BR_i(a_{-i}) = \arg\max\{u_i(a_i, a_{-i}) : a_i \in A_i\}$.

A *Nash equilibrium* is a profile of actions such that no player wants to deviate unilaterally, that is, a fixed point of $BR = (BR_i)_{i \in N}$:

Definition 1.1. A profile $a^* \in A$ is a Nash equilibrium if, and only if, for all $i \in N$, for all $a_i \in A_i, u_i(a^*) \geq u_i(a_i, a^*_{-i})$. Equivalently, for all $i \in N, a^*_i \in BR_i(a^*_{-i})$.

In Example 1.1, both players prefer to shirk, no matter what the other does. It means that shirking is always a best response. The only Nash equilibrium is therefore the profile (S, S), which leads to the suboptimal payoff $(1, 1)$. Would they both exert effort, they would earn $(3, 3)$, which Pareto-dominates $(1, 1)$. The Prisoner's Dilemma highlights the tension between Pareto-optimality and individual optimisation.

In the thesis, I only consider games with finite sets of actions. In general, a finite game does not admit a Nash equilibrium (an example of which is the rock-paper-scissor game). A player may want to randomize their actions. This leads to the definition of mixed strategies. A *mixed strategy* σ_i of player i is a probability distribution over A_i. Write $\Delta(A_i)$ the set of all mixed strategies of player i. Actions a_i are called *pure* and can be identified (as Dirac masses) to the extreme points of the simplex $\Delta(A_i)$. There are two possible interpretations of a mixed strategy: either a player genuinely uses a random device to choose their (pure) action, or the mixed strategy models the lack of information of the other players and of the observer as to which (pure) action is chosen (Harsanyi, 1973).

Mixed strategies are chosen independently by the players and one can therefore define the mixed extension of u_i as

$$u_i(\sigma) = \sum_{a \in A} \sigma_1(a_1)...\sigma_n(a_n) \times u_i(a)$$

This yields a new game $(N, (\Delta(A_i))_{i \in N}, (u_i)_{i \in N})$. Each u_i is multilinear in σ, hence continuous. A finite game extended with mixed strategies admits at least one Nash equilibrium (Nash,

1950b). Let σ^* be a Nash equilibrium in mixed strategies, $i \in N$ and $a_i \in A_i$. If $\sigma_i^*(a_i) > 0$, then $a_i \in BR_i(\sigma_{-i}^*)$. In other words, pure actions that are given a positive probability yield the same maximal payoff.

1.2.3 Informational asymmetries and Bayesian games

In Subsection 1.2.2, all the players have the relevant information. This can be decomposed in two parts. The information is *complete* when all the players know the utility functions, the available strategies and the payoffs of all the players. This means that the preferences and abilities of the players are well identified. Information is incomplete otherwise.

Information is *perfect*, in sequential games,[12] when all the players have perfect knowledge (memory) of all the past events and strategies taken by everyone, including the initial move by "Nature" when relevant. Hence chess is of perfect information whereas poker is not (players know the moves of other players but ignore the initial draw of cards). The study of such games relies on two sets of postulates: (i) the definition of rational behaviours, (ii) the definition of rational expectancies about the behaviour of other players (Harsanyi, 1966).

This thesis investigates games with incomplete information, called *Bayesian games*. One models a game of incomplete information as a game with *imperfect* but complete information (Harsanyi, 1968a, Harsanyi, 1968b, Harsanyi, 2004). Indeed, before players choose their actions Nature draws the "type" of each player (her preferences and utility function, her available strategies...) according to some prior probability. Each player learns only her type.

Formally, a Bayesian game is a tuple $(N, (A_i)_{i\in N}, (\Theta_i)_{i\in N}, (u_i)_{i\in N}, \mu)$ where N and A_i are defined as above, Θ_i is the set of "types" of player i and μ is a common prior distribution on $\Theta = \prod_{i\in N} \Theta_i$. The payoff depends on the action profile and the types, $u_i : A \times \Theta \to \mathbb{R}$. At step 0, Nature draws $(\theta_i)_{i\in N}$ according to μ and i learns θ_i. In player i's perspective, the probability that $-i$ has type θ_{-i} is $\mu(\theta_{-i}|\theta_i)$. A (mixed) strategy for player i is a mapping $\sigma_i : \Theta_i \to \Delta(A_i)$.

One can define the *Bayesian Nash equilibrium*:

[12] Namely games where there is an order of plays like chess. I do not develop on this as I will later define and consider particular cases of sequential games.

Definition 1.2. A profile σ^* is a Bayesian Nash equilibrium if and only if for all $i \in N$, for all $\sigma_i : \Theta_i \to \Delta(A_i)$,

$$\mathbb{E}\big[u_i(\sigma^*, .)|\theta_i\big] \geq \mathbb{E}\big[u_i(\sigma_i, \sigma^*_{-i}, .)|\theta_i\big] \tag{1.1}$$

where the expectation is taken over Θ_{-i} according to the distribution $\mu(.|\theta_i)$. In words, a Bayesian Nash equilibrium is a Nash equilibrium of the game taken in expectation where players update their beliefs about the types of others according to Bayes' rule.

The result of Nash (1950b) extends immediately. Finiteness of actions and types ensures the existence of a Bayesian Nash equilibrium.

1.2.4 Repeated games

This short presentation builds largely on Chapters 1 to 4 of Mailath and Samuelson (2006).

Subsection 1.2.2 defines simultaneous games that take place only once. Games where an order of plays is involved are called *sequential* and are often described in *extensive form* (decision trees). I will specifically focus on two such specific sequential games: (i) repeated games in discrete time, and (ii) repeated games in continuous time as used in the literature on strategic experimentation. I restrict myself to *perfect monitoring*, that is, at the end of every period all the players observe the profile of actions that was played.[13] In general, players face a trade-off between their short term interest (to betray, say) and their dynamic incentives (to cooperate, say). Such cooperating behaviours may emerge from dynaminc incentives in equilibrium.

Repeated games in discrete time. A *repeated game* (in perfect information) is the repetition of a *stage game* $(N, (A_i)_{i\in N}, (u_i)_{i\in N})$ for a specified number of periods. It may be finitely or infinitely repeated. If the game is finitely repeated, the natural tool to solve the game is *backward induction*: look at what players optimally do in the last period T; given what they will do in T, identify what they optimally do in $T-1$; unfold to period 1. In period T, players do not have any dynamic incentives. In the T-times repeated Prisoner's Dilemma (see Table 1.1), there

[13] If players randomize, only the outcome of randomization is observed. The mixture is only guessed in equilibrium.

cannot be team cooperation because the players shirk in T, then given they shirk in T, the best is to shirk in period $T-1$ and so on, so forth.

The situation is conceptually different when the game is infinitely repeated. Backward induction does not hold and players may have dynamic incentives at any period. Several definitions are needed. Choices in the stage game are actions; strategies refer to behaviours in the repeated game.

The set of period $t \geq 0$ histories is $H_t := A^t$ with typical element h^t. A history of length t is therefore a list of t action profiles. A history of length $t+1$ is therefore an element of $H_t \times A =: H_{t+1}$. The set of histories of finite length is $H = \cup_{t \geq 0} H_t$. A pure strategy for player i is a mapping $\sigma_i : H \to A_i$, it prescribes a pure action to every possible history of finite length. A mixed strategy is a probability mixture of pure strategies. A *behaviour strategy* is a mapping $\sigma_i : H \to \Delta(A_i)$. The latter are easier to work with.

Kuhn's theorem (Kuhn, 1953) allows us to work with behaviour instead of mixed strategies:

Theorem 1.1. *Every behaviour strategy admits a realization-equivalent mixed strategy: the two induce the same distribution over outcomes. The converse holds true.*

A strategy profile σ induces a realisation profile $(a^t(\sigma))_{t \geq 0}$: in period t, action profile $a^t = (a_i^t)_{i \in N} \in A$ is played. The *discounted* payoff of player i is given *ex post* by

$$(1-\delta) \sum_{t \geq 0} \delta^t \times u_i(a^t(\sigma))$$

$\delta \in [0, 1[$ is the discount factor of the players (assumed common).[14] Let $\mathbb{P}_\sigma$ the distribution over pure action profiles $a \in A$ induced by the strategy profile $\sigma = (\sigma_i)_{i \in N}$. A profile σ^* is a Nash equilibrium if and only if no one has a strictly beneficial deviation *ex ante*: for all $i \in N$ and

[14]The discount factor is here exponential. It gives a highly tractable recursive structure to the game. There are two main interpretations of δ. On the one hand it can model the preference for the present of the players (I prefer one apple today to one apple tomorrow) due to impatience or risk-aversion. On the other hand it can model a finite game with random termination date: if the stage game ends at the end of the current period with probability $1-\delta$, then the non-discounted expected payoff of player i is given by the above formula. In the following, I take the first interpretation. Doyle (2013) surveys the literature on time preference and discounting.

$\sigma_i : H \to \Delta(A_i)$,

$$(1-\delta)\sum_{t\geq 0}\Big(\delta^t \sum_{a^t\in A}\mathbb{P}_{\sigma^*}(a^t)\times u_i(a^t)\Big) \geq (1-\delta)\sum_{t\geq 0}\Big(\delta^t \sum_{a^t\in A}\mathbb{P}_{(\sigma_i,\sigma^*_{-i})}(a^t)\times u_i(a^t)\Big) \tag{1.2}$$

.

Provided players are patient enough, cooperation is sustainable in equilibrium in the repeated Prisoner's Dilemma. Indeed, always exerting effort so long as the others cooperate and shirk until the end after the first deviation is a Nash equilibrium (if $\delta \geq 1/2$ with the figures of Table 1.1). Shirking is said to be *off-path* because only one history is realized where everyone always cooperates.

This example highlights a major fact of repeated games: the set of equilibrium outcomes is in general larger than in the stage game. Define the set of *feasible* payoff vectors as $\mathcal{V} = \{(u_i(a)_{i\in N} : a \in \Delta(A)\}$ and the mixed-action minimax payoff of player i as

$$\underline{v}_i = \min_{\alpha_{-i}\in\Delta(A)_{-i}} \max_{a_i\in A_i} u_i(a_i, \alpha_{-i}) \tag{1.3}$$

The quantity $\underline{v}_i$ is the maximal payoff that i can ensure no matter what $-i$ does in the stage game. Say that $u(a)$ is *individually rational for player* i when $u_i(a) \geq \underline{v}_i$.

Theorem 1.2. *(Folk theorem) Any element of the convex hull of $\mathcal{V}$ that is individually rational for every player can be attained by an equilibrium profile provided that δ is large enough. The converse is true.*

The set of equilibrium outcomes of the repeated game is therefore large. However, the concept of Nash equilibrium is too permissive as some equilibrium profiles violate basic notions of rational sequentiality. Indeed it may prescribe irrational behaviour at histories that are off the equilibrium path. Indeed, if we consider again the Prisoner's Dilemma of Example 1.1, the cure may seem worse than the disease: one punishes (by shirking) a single occurence of S by the other (off-path), precluding cooperation ever after. This is not forgiving enough and is deemed irrational.

The continuation game from history $h^t \in H$ is the game that begins in period t, after history

h^t takes place (until time t, recall that all histories are finite). A continuation strategy induced by h^t is given by $\sigma_i|_{h^t}(h^\tau) := \sigma_i(h^t h^\tau)$ for all $h^\tau \in H$ where $h^t h^\tau$ is the concatenation of the two histories: history h^t happens first and is followed by h^τ. Repeated games have a recursive structure, any continuation game of history h^t is strategically identical to the original game. Call such games *subgames*.

The concept of Nash equilibrium does not impose any notion of optimality on subgames that are not attained by the equilibrium strategy.[15] *Subgame perfection* imposes that the behaviour prescribed by the profile be an equilibrium after any history, be it on- or off-path:

Definition 1.3. A strategy profile σ^* is a subgame perfect equilibrium if for all histories $h^t \in H, \sigma^*|_{h^t}$ is a Nash equilibrium of the repeated game.

In principle, proving subgame perfection requires to check that one has a Nash equilibrium after any histories. This means countable infinitely many verifications. Yet it suffices to show that no one-shot deviation is profitable. A *one-shot deviation* of player i from strategy σ_i is a strategy $\hat{\sigma}_i$ that agrees with σ_i everywhere but on one single history $h^t \in H$. Fix the adversarial strategy σ_{-i}. This deviation is *profitable* if the payoff of i is larger conditional on history h^t being reached. (For finitely repeated games with perfect information, backward induction always yields a subgame perfect equilibrium in pure strategies.)

As such a history h^t may not be reached in Nash equilibrium, Nash equilibria may admit profitable deviations. The following proposition says that one can look only at profitable one-shot deviations to study subgame perfection:

Proposition 1.1. *A profile σ is subgame perfect if and only if no player has a profitable one-shot deviation.*

The proofs rests on the result that, due to exponential discounting, a profitable deviation must be profitable in finitely many periods. Reducing to a profitable deviation on one period is done by backward induction.

[15] For instance, in the aforementioned example, players never shirk because they would get the suboptimal payoff thereafter. Yet nothing is said about the optimality of their strategies *should* someone shirk.

All the definitions naturally extend to repeated games where the stage game is a Bayesian game. For instance, if a new state of the world ω is drawn at every period and observed after the actions are taken then a history of length t is a tuple $(a^1, \omega^1, ..., a^t, \omega^t)$. Players update their beliefs between periods. Their strategies must be measurable with respect to the information they have.

A natural simplification considers only *Markov strategies*, that is, memoryless present-dependent-only strategies, in repeated games with partial information. A strategy σ is Markovian if for any two histories $h^t, \hat{h}^t \in H$ of same length and same terminal state (in period t), $\sigma(h^t) = \sigma(\hat{h}^t)$. Strategies are therefore markovian if they ignore all the characteristics of the histories except the length and the current state. They are in addition *stationary* if they depend only on the current state (and not on the length). A Markovian equilibrium (resp. stationary Markovian equilibrium) is a subgame perfect equilibrium in markovian strategies (resp. stationary markovian strategies). The current state is often referred to as the *state variable* (see below).

Repeated games in continuous time. In this paragraph, I define a specific repeated game in continuous time with possible states of the world. It is the infinitesimal repetition of a stage game where a state of the world ω is drawn from $\Omega = \{\omega_0, \omega_1\}$ according to the prior probability p^0. Identify $\Delta(\Omega)$ with $[0, 1]$. Player i has action set A_i and $u_i : A \times \Omega \to \mathbb{R}$.

Bayesian updating goes as follows. Let $\mathcal{F}_t$ be the σ-algebra spanned by the action profile $(a_s : s < t)$ up to time t: actions may give additional information. Then $p^t = \mathbb{P}(\omega = \omega_1 | \mathcal{F}_t)$.

The state variable is $p \in \Delta(\Omega) = [0, 1]$ and a stationary markovian strategy for player i is a function $\sigma_i : \Delta(\Omega) \to \Delta(A_i)$ that is measurable at any time t, namely with respect to the filtration $(\mathcal{F}_t)_{t \geq 0}$.

The *discounted expected payoff* of player i under strategy profile σ is

$$V_i(\sigma) = \mathbb{E}_{\omega,\sigma}\Big[\int_0^{+\infty} re^{-rt} u_i(., \omega) dt\Big] \tag{1.4}$$

where the expectation is taken over ω and σ (as σ itself depends on (p^t) which itself is determined by the actions taken by the players). By the law of iterated expectations, one can write the payoff

as

$$V_i(\sigma) = \mathbb{E}_{(p^t)_{t\geq 0},\sigma}\Big[\int_0^{+\infty} re^{-rt}\Big(p^t u_i(.,\omega_1) + (1-p^t)u_i(.,\omega_0)\Big)dt\Big] \tag{1.5}$$

$r > 0$ is the exponential discount rate. Notions of stationary Markov equilibrium and subgame perfection are defined in the same way as in the discrete case.

1.2.5 Coalitional Games with transferable utility

This short presentation builds on Chapters 1 to 3 and 8 of Peleg and Sudhölter (2007).

Cooperative games. The distinction between *cooperative games* and strategic games lies in the possibility for the players to sign enforcing contracts in the former, while this does not hold in the latter. The first cooperative game was the so called bargaining game (Nash, 1950a): two players may cooperate and reach an allocation that Pareto dominates a so called "outside option" (what they get if negotiation fails). The question is: What allocation do the players reach? The main method is *axiomatic*: one solution satisfies a given small set of reasonable axioms (efficiency, symmetry, etc.). Bargaining can be extended to n players. It is implicit that, if the n players fail to reach an agreement, then they all end up with their respective outside options. It is an all or nothing process as argued e.g. by Hart and Kurz (1983) and Hart and Mas-Colell (1996).

Coalitional game theory takes a slightly different point of view: cooperation is a partial process where subgroups can bargain while ignoring the rest of the players. These subgroups of players are called *coalitions*. The coalitional game is said to have *transferable utility* (hereafter "TU-games") if the players are free to make any side payments (of real money or direct utility) they wish within coalitions that form.[16]

Definitions of TU-games. Formally, a *TU-game* is a couple (N, v) where N is a finite non empty set of players and $v : 2^N \to \mathbb{R}$ is a mapping from each *coalition* $S \subseteq N$ to a real number $v(S)$ with the additional requirement that $v(\emptyset) = 0$. The quantity $v(S)$ denotes the maximal

[16] As Montero (2023) points out, this model bypasses the issue of moral hazard *within* coalitions.

utility (or money) that the members of S can collectively gather when they cooperate. It is called the *worth* of S. Call $\mathcal{G}^N$ the set of all TU-games on N and $\mathcal{G}$ the union of the $\mathcal{G}^N$. N is called the *grand coalition*; if N forms then everyone cooperates.

A TU-game does not govern how players cooperate to form coalitions or to share the dividends thereof. It is a description of all the possible payoffs of cooperation.

Let us write i instead of $\{i\}$ for singletons and $(i_1...i_k)$ as a short hand for $\{i_1, ..., i_k\}$ when no confusion is possible. The *marginal contribution* of player i to coalition $S \subseteq N \setminus i$ in game (N, v) is the quantity $v(S \cup i) - v(S)$. Several classes of TU-games are of interest:

Definition 1.4. Take (N, v) and write v instead.

- v is non-negative if $\forall S, v(S) \geq 0$;
- v is simple if it takes values in $\{0, 1\}$;
- v is monotonic if $\forall S \subseteq T, v(S) \leq v(T)$: the more, the better;
- v is superadditive if $\forall S, T \subseteq N$ such that $S \cap T = \emptyset, v(S \cup T) \geq v(S) + v(T)$;
- v is convex if $\forall S, T \subseteq N, v(S \cup T) + v(S \cap T) \geq v(S) + v(T)$.

Convexity implies superadditivity. Superadditivity implies monotonicity. The converse statements do not hold.

Definitions of solutions. Take a TU-game (N, v). There are two main questions. First, what coalitions will form (granted that each player can be part of at most one)? Within each coalition, how do the members share its worth? The major stand taken by the literature is the axiomatic method which prescribes a division of the grand coalition (that forms). This division is determined by the chosen set of axioms.

An *allocation* is a vector $x \in \mathbb{R}^N$ with the short-hand that $x(S) = \sum_{i \in S} x_i$ for all $S \subseteq N$. x_i represents the share of the worth allocated to $i \in N$. For a monotonic game, x is *feasible* if $x(N) \leq v(N)$: one cannot spend more than there is.

A *solution concept* on a set of players N is a multivalued mapping from the set of TU-games $\mathcal{G}^N$ to $\mathbb{R}^N$ such that any allocation given by the solution concept is feasible. If the solution

concept is single-valued, it is an *allocation rule*. A solution concept (in general) is defined on the set of TU-games $\mathcal{G}$ such that its restrictions to all the sets $\mathcal{G}^N$ are solution concepts (for N).

Axiomatic, the Shapley value, the core. Let φ be a solution concept. The two most famous ones are the *Shapley value* and the *core*. They rest on simple axioms:

- *Efficiency.* φ is efficient if for all (N, v) and $x \in \varphi(N, v), x(N) = v(N)$. Each allocation is therefore Pareto-optimal;
- *Dummy player.* Player i is dummy for game (N, v) if all her marginal contributions are null: $v(S \cup i) = v(S)$ for all $S \subseteq N \setminus i$. φ satisfies the dummy player axiom if it gives a null payment to a dummy player;
- *Symmetry.* Players i and j are symmetric in (N, v) if they have the same marginal contributions to all $S \subseteq N \setminus \{i, j\}$. φ satisfies symmetry if it gives the same payment to symmetric players;
- *Linearity.* φ satisfies linearity if it is linear (on the real vector space $\mathcal{G}$);[17]
- *Coalitional rationality.* φ is coalitionally rational if for all $(N, v), x \in \varphi(N, v)$ and $S \subseteq N, x(S) \geq v(S)$: no coalition can benefit from breaking away from the grand coalition.

The Shapley value (Shapley (1953)) is the allocation rule Sh defined for all (N, v) by

$$Sh(N, v)_i = \sum_{S \subseteq N \setminus i} \frac{s!(n-s-1)!}{n!} \big(v(S \cup i) - v(S)\big) \tag{1.6}$$

where s is the cardinal of coalition S.

Proposition 1.2. *The Shapley value is the only solution concept that satisfies Efficiency, Dummy player, Symmetry and Linearity.*

The Shapley value is defined on the marginal contributions only. It allocates to each player her expected marginal contribution to the coalitional game (if random orders of players are taken uniformly).[18]

[17] If φ is multivalued, linearity is defined set-wise.

[18] Moretti and Patrone (2008) discuss the "ubiquity" of the Shapley value.

The *core* $\mathcal{C}(N,v)$ of a coalitional game (N,v) is the set of feasible and coalitionally rational allocations (Gillies, 1959):

$$\mathcal{C}(N,v) = \left\{x \in \mathbb{R}^N : \forall S \subseteq N, x(S) \geq v(S), x(N) = v(N)\right\} \tag{1.7}$$

The core is well-defined but may be empty. The main drawback of the Shapley value lies in its lack of stability. Indeed, when the core of a game is not empty, the Shapley value may not belong to the core. Then, a coalition could deviate. This undermines the implicit assumption that the grand coalition forms. Such an example is given by the glove game (Shapley and Shubik, 1969).

Example 1.2. There are three players, one owns a left glove, the two others own a right glove. The purpose is to form a pair of gloves. Define (N,v) by $N = \{1,2,3\}, v(12) = v(13) = v(123) = 1, v(S) = 0$ else. Then $Sh(N,v) = (2/3, 1/6, 1/6)$ and $\mathcal{C}(N,v) = \{(1,0,0)\}$. Indeed for every allocation where $x_2 > 0$, players 1 and 3 can deviate together and earn more (for instance an additional $x_2/2$ for each). The roles of 2 and 3 are symmetric.

Conversely, the core seems too demanding. Consider the extended glove games with $2n+1$ players, n of which have a left glove, $n+1$ of which have a right glove. Each pair is worth 1. Then the core is reduced to the singleton that gives 1 to lucky left owners and 0 to unlucky right owners. However if n is big it seems that all the players should almost receive the same share.

Convexity of the coalitional game is a sufficient condition to ensure the non-emptiness of the core and then the Shapley value is the barycentre of the core.

Finally, the egalitarian solution is defined for all (N,v) by $Eq_i(N,v) = v(N)/|N|$. It satisfies Efficiency, Linearity and Symmetry but not the Dummy axiom.

2 Certification of Coalitions in Games of Imperfect Information

2.1 Introduction

Some projects must be carried out collectively because the resources that are necessary to the development of the project exceed the capacities of single stakeholders. As an example, consider several cities that need a water treatment plant. Cities are nearby enough so that a single plant suffices to fulfill their needs, thus building more than one plant is a waste of resources. The resources necessary to build the plant (e.g. land, human resources, raw material, money) are spread across the cities; therefore cities must cooperate.

Coalitional game theory is suitable to model this issue. Players form coalitions: the *worth* of a coalition of players is 1 (say, the coalition is "winning") if the coalition is able, as a whole and through cooperation, to carry out the project; the worth is 0 if the coalition is unable to do so (say, the coalition is "losing"). In the framework of the example, a coalition of cities is winning if, and only if, it is able to build a water treatment plant without external assistance (of the remaining cities). The coalitional game is exogenous in the sense that the worth of a coalition only transcribes objective abilities of the members and the way they fit in together.[1]

To fix the ideas, take two cities represented by their mayors Adam and Bob. Suppose that the coalitional game that encodes the situation is $\overline{v}$ defined by $\overline{v}(\text{Adam}) = 0, \overline{v}(\text{Bob}) =$

[1]The worth of a coalition does not encode the different bargaining powers of the players which are left aside in this paper.

$\overline{v}$(Adam, Bob) = 1. In words, Adam's city is not self-sufficient while Bob's city is.

If Adam and Bob were to decide on their own which coalition undertakes the construction of the treatment plant and to whom it will benefit, then it could be inefficient in terms of social welfare. Indeed, Bob could build a plant and restrict its use to his city only whereas the remaining city (Adam's) does not have the capacity to build another plant and runs out of water. In general, one wants to avoid two more cases: (i) more than one plant is built, (ii) no plant is built because the cities do not reach an agreement, even in the long run.

That is why the intervention of a social planner, here incarnated by the State or whichever political or administrative power that stands above the cities, is necessary to ensure that exactly one treatment plant is built and that every city has access to it. This goes through the formation of the grand coalition: all the cities are compelled to participate to the project. The social planner next distributes the costs (or benefits) of the treatment plant, that is, the dividends of cooperation, to the cities according to a cooperative solution. In the example, should a social planner use the Shapley value, we have: $Sh_{\text{Bob}}(\overline{v}) = 1, Sh_{\text{Adam}}(\overline{v}) = 0$.

Usually, cooperative game theory assumes that the coalitional game at play, call it $\overline{v}$, is common knowledge, both to the social planner and the players. In our example, it means that each city knows exactly which coalitions of cities–be it a member of the coalition or not–are winning, which are losing; and so for the social planner. This assumption seems unjustified here: Adam knows better than his peers or the State the quality of the land of his city. Likewise, mayors Adam and Bob know how well the different assets they have fit in together and how fruitful their cooperation can be; something Clarke, the mayor of a third city, or the social planner can ignore.

One can therefore wonder what happens when the social planner has a belief upon $\overline{v}$: it is a prior μ on a set of coalitional games $\mathcal{G}$. The social planner could dictate the formation of the grand coalition (N) and impose an allocation to the extent of her beliefs: $\varphi(\mathbb{E}_\mu[v]) \in \mathbb{R}^N$ where φ is a predetermined function. Consider again the two-cities example and suppose that the social planner believes three coalitional games are equally likely: $v_1 = \overline{v}$ (defined as before), v_2 with v_2(Bob) $= 0, v_2$(Adam) $= v_2$(Adam, Bob) $= 1$ and v_3 with v_3(Adam) $= v_3$(Bob) $=$

v_3(Adam, Bob) = 1. In v_3, both cities are self-sufficient whereas in v_1 (resp. v_2) only Bob is (resp. Adam). Then given the belief of the social planner, the payments are 1/2 for each player. Consequently, the ignorance of the social planner makes Bob worse off (and conversely it makes Adam better off).

On the one hand, if Bob has the possibility or the opportunity to prove the stand-alone worth of his city to the social planner, then he certifies $\overline{v}$(Bob) = v_1(Bob) = 1 to the social planner because he prefers her to believe that $\overline{v}$ must be either v_1 or v_3 (his share of the Shapley value moves from 1/2 to 3/4). By the same token, *if* the stand-alone worth of Adam were 1, he would certify it to the social planner. Yet Adam does not want to certify his stand-alone worth because he prefers that the social planner does not know that $\overline{v} = v_1$. Because she knows the incentive of Adam and she observes that he does not certify his stand-alone worth, the social planner understands that Adam is not self-sufficient. All in all she gains full information (namely that $\overline{v} = v_1$) and can allocate $Sh(v_1)$ to the players.

This improves Bob's payment. Naturally, such proof must be reliable to be efficient. For instance, cheap talk cannot be credible (here) because it would be in the interest of Bob to say his stand-alone worth is 1, regardless of whether it is true or not. Hence, Bob must prove his worth in a more credible manner: it must be true, reliable and easily verifiable. This is what I call the *certification* of a coalition. One can assume that Bob can gather enough reliable evidence to truthfully convince the social planner that his city is capable.

Consequently, the question I address in this paper is: Under what conditions do players have an incentive to certify their private information to the social planner?

To answer this question, I define a two-stage game of certification. n players of a coalitional game $\overline{v}$ and the social planner share a prior belief μ over a set $\mathcal{G}$ of simple and superadditive coalitional games. In addition, each player knows the worth of the coalitions he belongs to, *i.e.* his type. In the first step of the game, the worth of each coalition can be certified if its members unanimously agree to. In the second step, the social planner updates her belief given the certifications (it yields a new belief $\hat{\mu}$), imposes the formation of the grand coalition and

allocates $\varphi(\mathbb{E}_{\hat{\mu}}[v])$ to the players (where φ is a predetermined function). Both the players and the social planner face asymmetric information. The allocation $\varphi(\mathbb{E}_{\hat{\mu}}[v])$ is sound it term of welfare in the sense the same amount of utility is shared among the players (namely $\overline{v}(N)$) whatever the prior μ and the certifications: there is no waste.

The two-stage game admits a perfect Bayesian Nash equilibrium in pure strategies, in which a player certifies his stand-alone worth whenever it is winning, and no other coalition (Prop. 2.1). Then, I make two assumptions: the cooperative solution is strictly strongly monotonic (following Young (1985)); coalitional games that are *possible* are given the same weight by the prior belief μ. Though demanding, they are not sufficient to ensure that certifying all and only one's winning coalitions is a best response (Prop 2.2). Finally, I investigate the role of *minimal winning coalitions*–winning coalitions whose strict subsets lose. I characterise the set of perfect Bayesian Nash equilibria in three cases: when the ignorance of the social planner is restricted to minimal winning coalitions only ("horizontal asymmetry", Prop. 2.3); or to one minimal winning coalition and its subsets of cardinal *minus* 1 ("downwards asymmetry", Prop. 2.4); or to one minimal winning coalition and its supersets of cardinal *plus* 1 ("upwards asymmetry", Prop. 2.5). Furthermore, players certify all and only one's winning coalitions in strong Bayesian Nash equilibrium. I discuss the tightness of the assumptions through counterexamples of relaxations.

We have seen in the two-cities example above that Bob has an incentive to reveal his type while Adam does not. It makes sense because Bob is in a much stronger position in $\overline{v}$ (he is a "dictator": all and only his coalitions are winning) than in v (he is equally as strong as Adam) and that the Shapley value rewards that aspect. In Propositions 2.3 to 2.5, if a winning coalition S is certified to the social planner, then for every $i \in S$, every marginal contribution of i increases (against if S is not certified). This is how the social planner creates an incentive for the players to reveal their types. Obtaining this incentive is not straightforward as Proposition 2.2 shows.

Related literature. The two-stage certification game I present involves asymmetric information in a coalitional game and hence strategic aspects. It does so in a different fashion to what has been done so far in the literature. The literature about coalitional games with incomplete information is

scarce. Forges and Serrano (2013) and Forges (2017) are useful surveys for game with incomplete information for which binding agreements, that is cooperation, are possible. Salamanca (2020) defines Bayesian coalitional games (with non transferable utility). In his setting the utility of a player depends on both the allocation he gets when in a given coalition and the vector of types. What a player can get from a coalition S (the set of reachable allocations is D_S) is however commonly known; the asymmetric information is encoded in the types that are independent from the coalitions. It differs from my framework in which players do not know the worth of some coalitions (and hence what allocations are reachable by that coalition): the types are "within" the coalitional game.

My concept of *certification* relates to the literature on disclosure that begins with Grossman (1981), Milgrom (1981) and Milgrom and Roberts (1986). The three articles study the incentives for a seller to disclose the quality of their product with the possibility to withhold information but not to lie. Sellers cannot lie because their messages are verifiable at negligible cost *ex post* : sellers of good quality products therefore offer a warranty to the buyer. The cooperative structure left aside, my paper studies a multiple sender disclosure game. In Grossman (1981) there is a vector of possible qualities $D = (q_1 < q_2 < ...)$. The seller can disclose that the true quality q lies in a subset of D. The main result of the paper is that the consumer interprets the message of the seller as the most pessimistic possible one that is consistent with the affirmation. Therefore "the true quality q is greater or equal to q_i" is interpreted as "the true quality q is exactly q_i". As a consequence a seller truthfully and completely discloses the true quality of their product.

Milgrom (1981) studies the notion of "favorableness" of news: he gives a comprehensive model of the monotonicity between information and decision (e.g. the lower the cost of a firm, the lower it sets the price). In particular the model is applied to games of persuasion. In there, the seller has a vector of information regarding the quality of his good. He can decide to disclose (without lying) some characteristics to the buyer. Milgrom advocates for the use of the sequential equilibrium (Kreps and Wilson, 1982b) to make sure that the buyer cannot ignore information that is relevant to make her decision. In sequential equilibrium, the seller reveals everything. The reason, as in Grossman (1981), is that the buyer interprets any message as pessimistically

as possible. I use the concept of perfectness instead of sequentiality.

My paper differs from Grossman (1981) and Milgrom (1981)'s in two ways. First, the social planner learns the worth of a coalition only if its members unanimously decide so: a single player cannot in general disclose the worth of a coalition by his own. Second, the social planner needs not adopt a pessimistic interpretation. Indeed, the way the social planner interprets the certification of a coalition or the absence thereof is slightly more complicated. The social planner cannot ignore a certification: if S is certified she takes it into account. The way she interprets an absence of certification differs from Grossman (1981) and Milgrom (1981)'s insofar as this: if a coalition is not certified then two cases are possible, (i) it is optimal to some member of the coalition that the coalition is not certified (but his identity is unknown in general) or (ii) all the players want the coalition to be certified but do not reach unanimity.[2] Consequently, full revelation does not take place in general.

More recent papers study such disclosure games with multiple sellers, multiple buyers, unknown states of the world or dynamic disclosure (see e.g. Bhattacharya and Mukherjee, 2013, Au, 2015, Board and Lu, 2018) with a variety of results. The major difference between these papers and mine is that, in the latter, there is a cooperative structure and unanimity is required for certification (disclosure) to take place.

The other direction that is close to my paper is the literature about the strategic formation of coalitions (in complete information). There, the players can make proposals to form coalitions and solve the game. The formation is an endogenous process. Chatterjee et al. (1993) develop a protocol to solve a coalitional game. It consists of a sequential order with a discount factor where players propose the formation of a coalition and an allocation thereof. They study the stationarity of equilibria and investigate when the formation of the grand coalition occurs at the first step. Ray and Vohra (1999) add externalities, modelled by partition functions, to this framework.[3] Pérez-Castrillo and Wettstein (2001) slightly change the protocol and make it two-stage: first a bidding stage during which players bid to get the right to make a proposal; second

[2]This is what I call a "lack of coordination", see below.
[3]See also Konishi and Ray (2003).

a proposal stage in which the winner of the auction proposes the formation of a coalition and an allocation thereof. They show that in their protocol the subgame perfect equilibrium implements the Shapley value. My paper differs from theirs because, here, a social planner exogenously imposes the formation of the grand coalition and an allocation; a player's sole objective is to modify the belief of the social planner to his advantage. While his behaviour is strategic, he does not propose endogenous coalition formations to the other players.

Finally, Konishi et al. (1999) study coalition-proof Nash equilibria for agency games, that is when several principals announce incentive schemes to a single agent. A group deviation is *weakly* profitable if it makes at least one player better off and the remaining players at least as well off; it is a *strict* deviation if it makes all the players better off. A profile is a coalition-proof equilibrium when it is immune to group deviation that are themselves self-enforcing. I allow for weak deviations and make use of strong Nash equilibria, namely profiles for which no group has an interest to deviate provided the complementary strategy is fixed. Strong Nash equilibria are more demanding.

The paper goes as follows. In Section 2.2, I describe formally the two-stage certification game. Section 2.3 presents the equilibrium analysis and a discussion of the results. Section 2.4 presents possible extensions of the model. The proofs are in Section 2.5.

2.2 A Two-stage Game of Certification

Common framework. There are n players of a coalitional game and a social planner who all face asymmetric information about the coalitional game at play. In the first stage of the game, players have the ability to certify their private information under a unanimity restriction. In the second stage, the social planner updates her belief about the coalitional game, imposes the formation of the grand coalition and an allocation to the players according to a predetermined cooperative solution.

Let $N = \{1, ..., n\}$ be the set of players. They are the players of an unknown coalitional game that is simple and superadditive. They will be the (strategic) players of the information game.

A coalitional game with set of players N is a vector $(N, v) : 2^N \to \mathbb{R}$ such that $(N, v(\emptyset)) = 0$. As there is no ambiguity about the set of players, one writes v instead of (N, v). It is *simple* if it takes values in $\{0, 1\}$ with $v(N) = 1$. It is *superadditive* if for all disjoint coalitions S and $T \subseteq N, v(S \cup T) \geq v(S) + v(T)$. Let $\mathcal{G}$ be the set of such simple and superadditive games on the set of players N.

The players and the social planner share a prior μ on $\mathcal{G}$ that is common knowledge. I define the expected coalitional game at play as

$$v_\mu := \mathbb{E}_\mu[v] = \sum_{v \in \mathcal{G}} \mu(v) \times v$$

where the linear combination $v + \lambda w$ is defined by $(v + \lambda w)(S) := v(S) + \lambda w(S)$ for all v, w coalitional games, $\lambda \in \mathbb{R}, S \in 2^N$. (One extends the definition and allows a coalition to take any real value.) It is easy to check that v_μ is superadditive. Note that in general v_μ is *not* simple.

Type of $i \in N$, beliefs *ex interim*. Nature picks a coalitional game $\overline{v}$ according to μ and reveals to each player $i \in N$ the worth $\overline{v}(S)$ of the coalitions S to which i belongs, that is his type. The private belief of i over $\mathcal{G}$ is denoted $\mu_i^{\overline{v}}$ and defined for all $w \in \mathcal{G}$ by

$$\mu_i^{\overline{v}}(w) = \frac{\mu(w)}{\mu(\mathcal{G}_i^{\overline{v}})} \mathbb{1}_{\mathcal{G}_i^{\overline{v}}}(w)$$

where $\mathcal{G}_i^{\overline{v}} := \{w \in \mathcal{G} : \forall S \ni i, w(S) = \overline{v}(S)\}$ is the set of all possible coalitional games to player i conditional on his private information.[4]

Player $i \in S$ is better informed about the worth of S than the social planner if, and only if, $0 < v_\mu(S) < 1$. He can be better informed than the social planner about the worth of a coalition he does not belong to, but not necessarily: this depends on the prior μ.[5]

[4]By construction, $\mu_i^{\overline{v}}(v) > 0$ and $\cap_{i \in N} \mathcal{G}_i^{\overline{v}} = \{\overline{v}\}$. The superscript does not mean that i knows $\overline{v}$. In fact, for all $\overline{v} \in \mathcal{G}$, either there exists $w \neq \overline{v}$ such that $\mu_i^w = \mu_i^{\overline{v}}$ or $\mu_i^{\overline{v}}(\overline{v}) = 1$ (namely i knows that $\overline{v}$ is the true coalitional game).

[5]The two extreme cases are (1) the knowledge of his type does not give him additional information about the worth of the coalitions he does not belong to, namely $v_{\mu_i^{\overline{v}}}(T) = v_\mu(T)$ for all $T \not\ni i$; (2) knowing his type is sufficient to get $\overline{v}$, namely knowing one's worth implies knowing the worth of all coalitions. Everything in between is possible. In the case of the leading example, it is expected that if city $i \in \{i, j, k\}$

Strategies of the players, what the social planner observes. The social planner offers each coalition of players the opportunity to certify its worth through hard evidence: players can disclose or withhold true information but cannot lie. Unanimity is required: certification occurs if, and only if, each player of the coalition agrees to certify. Pre-play negotiation is not allowed.

Formally, call $\Delta = \{(\mu_j^v)_{j\in N} : v \in \mathcal{G}, \mu(v) > 0\}$ the set of all possible vectors of *ex interim* beliefs, $\Delta_i = \{\mu_i^v : (\mu_j^v)_{j\in N} \in \Delta\}$. A (pure) strategy for player i is a function a_i that to each $\mu_i \in \Delta_i$ associates a binary 2^{n-1}-dimensional vector $\left(a_i(S|\mu_i)\right)_{i\in S\subseteq N}$. If $a_i(S|\mu_i) = 1$ then i acts to disclose the worth of S; if $a_i(S|\mu_i) = 0$ then he does not.

Let a be a strategy profile. The social planner does not observe directly a but only the "outcome" of a. As unanimity is required, the social planner observes the worth of $S \subseteq N$ through disclosure if, and only if, $\prod_{i\in S} a_i(S) = 1$. (Equivalently, for all $i \in S, a_i(S) = 1$.) I call this *direct certification*.

***Ex post* belief and action of the social planner.** The equilibrium concept is the perfect Bayesian Nash equilibrium (hereafter "equilibrium"). Let $\hat{\mu}(a)$ be the posterior belief of the social planner given the profile a. She imposes the allocation $\varphi(v_{\hat{\mu}(a)}) \in \mathbb{R}^N$ where the cooperative solution φ is *efficient*, namely $\sum_{i\in N} \varphi_i(w) = w(N)$ for each coalitional game w. $\varphi_i(w)$ is interpreted as the share of the worth $w(N)$ of the grand coalition allocated to player i (for a game w). The players are risk-neutral and choose their strategies *ex interim*. Their sole purpose is to shape the posterior belief of the social planner to their advantage. The expected payoff of i given a is $u_i(a) = \varphi_i(v_{\hat{\mu}(a)})$. The expectation comes from $v_{\hat{\mu}(a)} = \mathbb{E}_{\hat{\mu}(a)}[v]$.

$\hat{\mu}(a)$ is obtained by Bayes' rule: given the certifications that the social planner observes and what she can deduce from the optimality of the strategies of the player, she forms her posterior belief. Note here that the social planner cannot dismiss or ignore some information as in the cheap talk literature: if S is certified (namely by hard evidence, it cannot be false) then she knows this and cannot interpret it as babbling. Hence if the members of a coalition take

knows how well i, j and k fit in together, how complementary the three of them are, then it gives insight to i about the joint ability of cities j and k. This is precisely encoded by the prior μ.

an off-path equilibrium action and a coalition S is certified as winning (while it should not in equilibrium), then the social planner associates in particular zero probability to games $v \in \mathcal{G}$ such that $v(S) = 0$.

Timing of the game. We can summarise how the game goes:

1. Nature picks a coalitional game $\overline{v}$ according to the prior distribution μ on $\mathcal{G}$;
2. For every $S \in 2^N$, Nature reveals $\overline{v}(S)$ to every $i \in S$, it yields the *interim* belief $\mu_i^{\overline{v}}$;
3. Every player $i \in N$ plays his strategy $a_i(\mu_i^{\overline{v}}) = (a_i(S|\mu_i^{\overline{v}}))_{i \in S \subseteq N}$;
4. The social planner forms a posterior belief $\hat{\mu}(a(\mu^{\overline{v}}))$ and allocates $\varphi(v_{\hat{\mu}(a(\mu^{\overline{v}}))})$ to the players.

2.3 Equilibrium Analysis

I first describe the structure of superadditive games and explain how direct certifications can entail indirect certifications. I state and define the property that the cooperative solution φ verifies: φ is strictly strongly monotonic (SSM). In words, it means that if all the marginal contributions of player i are larger in game v than in w (with at least one strict inequality), then $\varphi_i(v) > \varphi_i(w)$. Then I prove that there always exists an equilibrium in pure strategies provided the function φ is strictly strongly monotonic (SSM). Afterwards, I show that it is not optimal in general for a player to certify a winning coalition. Finally, I prove that for some types of priors, certifying all and only the winning coalitions uniformly improves i's position. In words, each strategy profile induces a posterior belief to the social planner and yields an expected coalitional game. Following this strategy ensures to player i the largest possible expected marginal contributions with respect to the posterior belief (compared to the marginal contributions resulting from any other strategy of i). I provide counterexamples to this fact for extensions of these specific priors. The proofs are given in Section 2.5.

From now on, suppose that Nature has picked a coalitional game $\overline{v}$ and that players stand *ex interim*. To ease the notation, I write $a_i(S)$ instead of $a_i(S|\mu_i^{\overline{v}})$. I shall investigate the incentive of the players to reveal their private information. Recall that a coalitional game v is *superadditive* if, for all $S, T \in 2^N$ such that $S \cap T = \emptyset, v(S \cup T) \geq v(S) + v(T)$. The belief of the social planner is consistent with the superadditive structure of $\overline{v}$: for all *interim* beliefs $\mu^{\overline{v}}$ and strategy profiles $a, v_{\hat{\mu}(a(\mu^{\overline{v}}))}$ is also a superadditive game. Indeed, superadditive games are stable by non-negative linear combinations, $\hat{\mu}(a(\mu^{\overline{v}}))$ is a vector of non-negative weights and taking the expectation boils down to computing a large linear combination. Naturally, the superadditive structure of the game yields information to the players. Player i can gather indirect information about the worth of the coalitions that are *unknown* without belonging to them.

Definition 2.1. The worth of $S \in 2^N$ is *unknown* (to the social planner) if there exists $v \in \mathcal{G}$ with $\mu(v) > 0$ and $v(S) \neq \overline{v}(S)$. Equivalently, $0 < v_\mu(S) < 1$. I denote by $\mathcal{X} = \mathcal{X}_\mu$ the set of unknown coalitions.

Here is how indirect information works. Take $S, T \in \mathcal{X}$ such that $i \in S, i \notin T$. If S is a subset of $N \setminus T$ and S is winning, then player i knows that T is losing. Indeed, out of superadditivity, $1 = \overline{v}(S \cup T) \geq \overline{v}(S) + \overline{v}(T) = 1 + \overline{v}(T)$. Superadditivity therefore encodes what a player can deduce from his type. If[6] coalition (ijk) is winning, then the superadditive structure gives information to player i about how likely (jk) itself is winning (such sharp information the social planner does not have).

Likewise, the certification of a coalition has indirect effects. The direct certification of a winning coalition S entails the *indirect certification* (as winning) of every superset of S and (as losing) of every subset of $N \setminus S$. Consequently, when choosing his strategy, i must take into account the indirect certifications it entails, not only the direct ones. Important is the fact that two disjoint coalitions can both lose at the same time. Thus, considering superadditive games gives rise to interesting second order phenomena and incentives for the players to reveal their private information.

[6] $(i_1 ... i_t)$ a shortcut for $\{i_1, ..., i_t\}$.

A winning coalition for which every strict subset loses cannot be certified indirectly as described in the previous paragraph. I call such a coalition *minimal winning*: $S \subseteq N$ is minimal winning for $v \in \mathcal{G}$ if, and only if, $v(S) = 1$ and $T \subsetneq S \implies v(T) = 0$. Call $\mathcal{M}(v)$ the set of minimal winning coalitions for v. Knowing $\mathcal{M}(v)$ is sufficient to have a complete description of v. Example 2.1 highlights how indirect certifications work:

Example 2.1. $N = (1234), \mathcal{M}(\overline{v}) = \{(12), (23)\}, \mathcal{X} = 2^N \setminus \{N, \emptyset\}$. Suppose player 1 must respond to the strategy profile a_{-1}: $\forall i \in \{2, 3, 4\}, \forall S \in 2^N$ such that $i \in S, a_i(S) = 1$

If player 1 wants (12) to be certified (surely), he must set $a_1(12) = 1$ while he is indifferent about $a_1(123)$ if he wants (123) to be certified: (123) is already certified indirectly through the direct certification of (12). The strategies $(a_1(12), a_1(123)) \in \{(1, 0), (1, 1)\}$ (all other things remaining equal) yield the same superadditive game $\mathbb{E}_{\hat{\mu}(a(\mu^{\overline{v}}))}$, hence the same payoff. They are not distinguishable from the perspective of the social planner.[7] However, no indirect certification of (12) is possible. The reason is, all the subsets of (12) may lose. (A precision is needed. $\overline{v}$ is such that (1) and (2) lose. Yet player 1 is situated *ex interim* and in his perspective (2) can win with a positive probability.)

Finally, we need to define what the properties of the solution φ are. First, φ needs be defined only on the convex hull of $\mathcal{G}$, that is, $Co(\mathcal{G}) = \{v : 2^N \to [0, 1] : v(N) = 1, v(\emptyset) = 0, v \text{ is superadditive}\}$, because the players cannot induce expected games (given the belief of the social planner) outside this set. Take then a coalitional game v. The *marginal contribution* of agent i to coalition $S \subseteq N \setminus i$ for the game v is $v(S \cup i) - v(S)$. Definition 2.2 is very similar to Young (1985)'s *strong monotonicity*. His ordering between coalitional games and his concept of monotonicity are large whereas I focus on the strict component (for details, see Footnote 9 below).

Definition 2.2.

- For all $i \in N, v, w$ coalitional games, say i prefers v to w and write $v \succeq_i w$ if for all coalition $S \subseteq N \setminus i, v(S \cup i) - v(S) \geq w(S \cup i) - w(S)$. In words, i prefers v to w if, and only if, every marginal contribution of i for v is weakly greater than the corresponding

[7] And they are payoff equivalent for the players.

marginal contribution for w. The preference is strict ($\succ_i$) if at least one of the inequalities is strict.

- φ is *strictly strongly monotonic* (SSM) if for all player $i \in N, v \succ_i w \implies \varphi_i(v) > \varphi_i(w)$.

The Shapley value,[8] defined for all coalitional game v by

$$\forall i \in N, Sh_i(v) = \sum_{S \subseteq N \setminus i} \frac{s!(n-s-1)!}{n!}\big(v(S \cup i) - v(S)\big)$$

satisfies efficiency and strictly strong monotonicity.[9,10]

Without loss of generality, we can restrict the strategies to the *unkwown* coalitions. Certifying a coalition which worth is known by the social planner does not influence her belief (neither directly, nor indirectly). Denote $\mathcal{G}(\mathcal{X}) = \{v \in \mathcal{G} : \forall S \notin \mathcal{X}, v(S) = \overline{v}(S)\}$ the set of possible games given that $\mathcal{X}$ is the set of unknown coalitions.

Recall that a strong perfect Bayesian Nash equilibrium (hereafter "strong equilibrium") is a perfect Bayesian Nash equilibrium for which no group of players has an incentive to deviate as a coalition.

[8] Without ambiguity, s stands for the cardinality of coalition S.

[9] One can wonder how large the set of compatible cooperative solutions φ is. Young (1985) shows that the Shapley value is the only cooperative solution defined on the set of all cooperative games that satisfies efficiency, strong monotonicity (replace for all $i, \succ_i$ by $\succeq_i$ and $>$ by $\geq$ in my definition) and symmetry (the allocation of a player in the permuted game is the allocation of the permuted player in the original game). Here, I do not require the solutions φ to be symmetric. Besides, I do not require the solution to be defined on the set of all coalitional games. In particular, the size of the cake is always the same in my setup: there is always $v(N) = 1$. Therefore, I allow for many more cooperative solutions than affine transformations of the Shapley value. For instance, $v \to Sh(v^\alpha)$ fulfils the criteria for all $\alpha > 0$.

[10] One can also wonder why I do not consider power indices (for a survey, see Andjiga et al. (2003); in their article, "simple" stands for "binary and monotonic for the inclusion", "proper" for "superadditive"). Power indices indeed measure the power of each player in a voting game–a framework close to my simple superadditive games. The reason is twofold. First, while the real game $\overline{v}$ is in $\{0, 1\}$, the social planner applies φ to an expected game given her belief. In general, this game is no longer binary; power indices are not defined over such games in general. Besides, many famous indices consider only minimal winning coalitions (e.g. the Banhzaf power index that measures the probability to change the outcome of a vote). However, it is impracticable in my setting because the social planner knows *not* what the minimal winning coalitions are in general.

Proposition 2.1 states that there always exists a simple equilibrium:

Proposition 2.1. [11]

Let μ be a prior on $\mathcal{G}$. The following profile is an equilibrium: for all player $i \in N$, i certifies $\{i\}$ if, and only if, $\{i\}$ is a winning coalition, he does not certify any coalition that contains two or more players (regardless of the interim *belief).*

The proof is quite straightforward: it S contains two or more members and if $j \neq i \in S$ does not certify S, then it is a best response for i not to certify it either (because of the unanimity requirement). The fact that it is a best response to certify winning $\{i\}$ comes naturally from the superadditive structure and the (SSM) property. (Note that because the equilibrium is *perfect*, the social planner knows that $\{i\}$ is losing whenever it is *not* certified. She cannot infer anything from the fact that a coalition S with more than two members is not certified because the decision not to certify the coalition depends only on its cardinal, not on its worth.)

In light of the second order phenomena described above when only superadditive games are given a positive probability, one can think that it is always a best response to certify one's winning coalitions, not one's losing coalitions. Therefore, it seems that the corresponding profile is an equilibrium. Indeed, the reasoning goes, consider $i \in S, S$ winning. If S is certified as winning, then so are all the supersets, to which i belongs; and all the subsets of $N \setminus S$, to which i does not belong, are certified as losing. It seems to improve i's every marginal contributions (compared with i not certifying S). This, however, is false in general:

Proposition 2.2. *(Impossibility to reveal information in general.) Suppose that $n \geq 4$ and that φ is (SSM). There exists a couple $(\mu, \overline{v}), i \in N$ and a strategy profile a such that (i) a is an equilibrium, (ii) $i \in S$ does not certify some winning coalition S and (iii) it would make i worse off to do so.*

[11] I thank an anonymous referee to point me the existence of this simple equilibrium in pure strategies when the simple games are superadditive and the cooperative solution is (SSM). If one of the two assumptions is lacking, then the proof falls down. A way, then, to ensure the existence of an equilibrium is to suppose that the functions φ_i are continuous and to allow the players to use mixed strategies (some definitional subtleties arise): the posterior belief of the social planner is continuous in the strategy profile and one applies a standard result of existence.

If one does not constrain φ to be (SSM) then Proposition 2.2 is easily false: consider for instance a φ that is *decreasing* with respect to marginal contributions such that players want in general to certify their losing coalitions instead.

I give conditions under which players indeed want to certify all and only their winning conditions. I always assume that $(H1) : \varphi$ is (SSM); $(H2) : \mu$ is such that $v \in \mathcal{G}(\mathcal{X}) \implies \mu(v) = 1/\#\mathcal{G}(\mathcal{X})$. Property $(H2)$ means that the prior distribution assigns the same probability to every possible game: the social planner follows a principle of insufficient reason. In Proposition 2.3, a subset of what are the minimal winning coalitions of $\overline{v}$ is unknown. A minimal winning coalition is pivotal: any subset loses and this coalition is the smallest one (with respect to monotonicity) to win. Informally, the ignorance of the social planner can be phrased in this way: all the coalitions she ignores are "on the tip" of winning, they are a limit case. What she does not know is whether they indeed win: *ex ante* $S \in \mathcal{X}$ can be either the smallest winning coalition or the largest losing coalition.

Proposition 2.3. *(Horizontal asymmetry.) Suppose that $\mathcal{X} \subseteq \mathcal{M}(\overline{v})$, namely coalition S is unknown only if it is minimal winning.*[12] *Suppose $(H1)$ and $(H2)$ hold. Then a profile a is a Nash equilibrium if, and only if, for all $S \in \mathcal{X}$, one of the two conditions is satisfied:*

- $(C1) : \forall i \in S, a_i(S) = 1$;
- $(C2) : \exists i \neq j \in S \text{ s.t. } a_i(S) = a_j(S) = 0$.

Furthermore, a is a strong equilibrium if, and only if, for all $S \in \mathcal{X}$ condition $(C1)$ is satisfied.

Under condition $(C1)$, every player of a winning coalition S decides to certify it (with full probability) and the certification takes place. It turns out to be optimal for all the players of S. However, in equilibrium, it may happen that S is not certified (with zero probability) even though every player of S would like to. Indeed, because of the unanimity requirement, if two players decide to certify S with probability zero, the deviation of one single player is not sufficient to have it certified with positive probability. I call this situation a "lack of coordination": this is precisely described by condition $(C2)$.

[12] The result holds for one up to every minimal coalitions that are unknown.

In Proposition 2.4, one minimal coalition S and its subsets of cardinal $s-1$ are unknown: the asymmetry of information is downwards. Here the social planner does not know whether the minimal coalition is S or some $S \setminus i, i \in S$,[13] or all the $S \cup k, k \notin S$. She does not know where the limit case is.

Proposition 2.4. *(Downwards asymmetry.) Take one S minimal winning for $\overline{v}$. Suppose that $\mathcal{X} = \{S, S \setminus i : i \in S\}$, $(H1)$ and $(H2)$ hold. Then a profile a is an equilibrium if, and only if,* (1) *for every losing coalition $S \setminus i$, there exists $j \in S \setminus i$ s.t. $a_j(S \setminus i) = 0$; and* (2) *coalition S satisfies one of conditions $(C1)$ or $(C2)$ as defined in Proposition 2.3. The profile is a strong equilibrium when condition $(C1)$ is satisfied.*

Finally, the asymmetry of information of Proposition 2.5 is upwards: one minimal winning coalition S and its supersets of cardinal $s+1$ are unknown. It is a converse statement to Proposition 2.4: the social planner does not know whether the minimal winning coalition (the limit case) is S or some $S \cup k, k \notin S$, or all the $S \cup (kl), k, l \notin S$.

Proposition 2.5. *(Upwards asymmetry.) Take one S minimal winning for $\overline{v}$. Suppose that $\mathcal{X} = \{S, S \cup k : k \in N \setminus S\}, (H1)$ and $(H2)$ hold. Then a profile a is a Nash equilibrium if, and only if, either (α) coalition S satisfies condition $(C1)$ as defined in Proposition 2.3; and for all $k \in N \setminus S$, one of the two following conditions holds:*

- $a_k(S \cup k) = 1$;
- $\exists i \in S$ *s.t.* $a_i(S \cup k) = 0$.

or (β) coalition S satisfies $(C2)$ while for all $k \in N \setminus S$, coalition $S \cup k$ satisfies one of conditions $(C1)$ or $(C2)$.

Furthermore, the profile a is a strong Nash equilibrium when S is certified (situation (α)).

Remark 2.1. The requirement that μ is uniform on $\mathcal{G}(\mathcal{X})$ in the propositions is a sufficient condition. It is possible to prove that a small perturbation of the uniform prior works when the φ_i are assumed continuous (see Proposition 2.6). There are priors that are far from the uniform

[13] It can be several in principle, according to μ.

for which the results do not hold. I do not know whether there exist priors "far" for which the results *do* hold.

Discussion of the results. Propositions 2.3, 2.4 and 2.5 study different situations but are very close in spirit. They hold for every φ that is (SSM). Indeed, the proofs establish that when a winning S is certified, it increases *every* marginal contribution of every player $i \in S$.[14] The same holds for i succeeding *not* to certify a losing coalition. In that sense, a player uniformly improves his position in the sense of Definition 2.2.

In the proof, I show that in each situation, a player of a winning coalition S is better off when S gets certified to the social planner. However, if the strategy profile is such that at least two players decide not to certify S and due to the requirement of unanimity, the deviation of a single player cannot make the certification of the coalition possible (condition $(C2)$). This is a "lack of coordination." In Proposition 2.4, a player is worse off when a losing coalition he belongs to is certified. The player can unilaterally prevent the certification of the coalition thanks to unanimity. Therefore, no losing coalition is certified at the equilibrium.

Finally, while there are many equilibrium profiles (because of the many possible indifferences when a coalition is not certified), one can list all the equilibrium payoffs. For instance, there are exactly $2^{|\mathcal{X}|}$ equilibrium payoffs in Proposition 2.3: for each unknown coalition S, there are two possibilities, $(C1)$ and $(C2)$. Likewise, there are exactly two equilibrium payoffs in Proposition 2.4: no losing coalition is certified, there are two possibilities for S, $(C1)$ and $(C2)$. Last, it is slightly more cumbersome to list the equilibrium payoffs in Proposition 2.5 but one can show that there are exactly $1 + 2^{n-s}$ of them.

The assumption that the prior μ gives the same probability to every possible game is very restrictive. It does not mean that the prior is symmetric (anonymous) in the players. Indeed, the social planner may have more knowledge about some players than others. For instance, in Proposition 2.4, the social planner has full knowledge about the worth of k's coalitions, $k \notin S$;

[14] Conversely, it makes the $k \notin S$ worse off. One can view the certification (or not) of S as a zero-sum game between the members of S and those of $N \setminus S$. It is zero-sum because the amount to be shared, namely $\overline{v}(N)$ is constant. Of course, one can view as many zero-sum games as there are unknown coalitions.

	(12)	(23)	μ
v_1	1	0	1/4
v_2	0	1	1/8
v_3	1	1	1/4
v_4	0	0	3/8

Table 2.1: Counterexample to Proposition 2.3 for a general prior

not about i's coalitions, $i \in S$. Rather, it means that the *ignorance* of the social planner must be symmetric in the players that are concerned. If the social planner does not know the worth of S and $S \setminus i$, then she does not know the worth of every $S \setminus j, j \in S$. In addition, what she does not know is given the same weight. It is in this sense only that μ is uniform an treats players symmetrically. It follows a sort of principle of insufficient reason.

What can we conclude from this result? In strong equilibrium a^* every winning coalition is certified (be it directly or indirectly). *A contrario*, the social planner understands that a coalition that is *not* certified is losing. Thus, the information is fully revealed: the social planner knows that the superadditive game at play is $\overline{v}$, *i.e.* $\hat{\mu}(a(\mu^{\overline{v}}))(\overline{v}) = 1$, and can allocate $\varphi(\overline{v})$.

Tightness of the result. Proposition 2.2 shows that, in general, it is not a best response to certify one's winning coalitions, not one's losing coalitions. Propositions 2.3 to 2.5 depend crucially on hypothesis $(H2) : \mu(v) = 1/\#\mathcal{G}(\mathcal{X}), v \in \mathcal{G}(\mathcal{X})$. The following examples show that there exist priors μ over $\mathcal{X}$ for which the propositions do not hold. Let a_i^* denote the strategy "i certifies all and only his winning coalitions."

First, a counterexample for the horizontal asymmetry of information:

Example 2.2. Consider $N = (123), \mathcal{M}(\overline{v}) = \{(12), (23)\}$ and suppose that players 2 and 3 certify their winning coalitions (as player 1 is situated *ex interim*, he does not know that (23) is winning). There are four possible games given with their respective prior probabilities as depicted in Table 2.1.

If player 1 faces a_{-1}^* and plays the strategy "1 does not certify (12)," then the expected worth of (23) for the social planner is $3/8 < 1/2$ which would be the expected value of (23) had player 1 played a_1^* (the expectations are *ex interim*: player 1 does not know whether (23) is a

	(12)	(13)	(23)		(12)	(13)	(23)
v_1	1	1	1	v_5	0	1	1
v_2	1	1	0	v_6	0	1	0
v_3	1	0	1	v_7	0	0	1
v_4	1	0	0	v_8	0	0	0

Table 2.2: Counterexample to Proposition 2.4 for a general prior

winning coalition). Therefore, if φ is sufficiently biased towards player 1's marginal contribution to coalition (23), then player 1 is better off when he does *not* certify his winning coalition (12). (Recall that φ is increasing with respect to marginal contributions. The weight the Shapley value gives to a marginal contribution to coalition S depends only on the cardinal of S: it is $(s!(n-s-1)!)/n!$. I say that φ is biased if it gives different weights to at least two marginal contributions to coalitions of the same cardinal. Above, φ gives a large weight to $v(123)-v(23)$ and is therefore biased towards this marginal contribution.)[15]

The weakness of this counterexample lies in the fact that it does not work for $\varphi = Sh$: one needs a biased cooperative solution. The existence of a counterexample for the horizontal asymmetry, a general prior and the Shapley value as a cooperative solution is yet to be found or denied. Second, Proposition 2.4 does not hold when $(H2)$ is not verified. The corresponding counterexample works with $\varphi = Sh$.

Example 2.3. Take $N = (123), \mathcal{M}(\overline{v}) = \{(123)\}, \mathcal{X} = \{(12),(13),(23)\}$. Suppose $\varphi = Sh$. Consider player 1 and suppose player 2 (resp. 3) certifies (12) but not (23) (resp. (13), not (23)) (strategy σ_{-1}). There are eight possible games as depicted in Table 2.2.

Call τ_1 the strategy "1 certifies (13) only." Consider

$$\mu = \left(\varepsilon^2, \varepsilon^2, \frac{1-\varepsilon}{2}, \frac{1-\varepsilon}{2}, \varepsilon^2, \varepsilon - 5\varepsilon^2, \varepsilon^2, \varepsilon^2\right)$$

This prior is such that the social planner believes that v_3 and v_4 are equally likely and that one of the two is $\overline{v}$ with probability close to one (up to an ε). It is such that, if the social planner

[15] One can check that, if $\varphi = Sh$, then whatever the positive prior μ, 1 is better off when (12) is certified. What 1 gains at not certifying (12)–his marginal contribution to (23)–is offset by his loss–his marginal contribution to (2).

knows that (12) loses (namely, $\overline{v} \in \{v_5, ..., v_8\}$), then she believes that with a small error $\overline{v} = v_6$ (as ε is small, $\varepsilon - 5\varepsilon^2$ is large before $3\varepsilon^2$).

On the one hand, if player 1 does not certify his coalitions, then $\mathbb{E}_{\hat{\mu}(a_1^*, \sigma_{-1})}[v] = v_\mu$ is close to the average $(v_3 + v_4)/2$. It entails that $Sh_1(\mathbb{E}_{\hat{\mu}(a_1^*, a_{-1})}[v]) = 1/3 + O(\varepsilon)$ with $O(\varepsilon)$ a rest of magnitude ε. On the other hand, $\mathbb{E}_{\hat{\mu}(\tau_1, a_{-1})}[v] \approx v_6$, where players 1 and 3 are (almost) symmetrical and $Sh_i(\mathbb{E}_{\hat{\mu}(\tau_1, a_{-1})}[v]) = 1/2 + O(\varepsilon)$.

Consequently, with this μ and when $\varphi = Sh$, player 1 prefers to certify (12) (against σ_{-1}). (A similar counterexample can be given for Proposition 2.5 and a prior μ that is not uniform.)

Now, the result holds for both downwards and upwards asymmetries of information, namely for $\mathcal{X} = \{S, S \setminus i\}$ and for $\mathcal{X} = \{S, S \cup k\}$ provided μ is uniform on $\mathcal{G}(\mathcal{X})$. It cannot be generalised to an asymmetry both ways, namely when $\mathcal{X} \subseteq \{S, S \cup k, S \setminus i : k \in N \setminus S, i \in S\}$.[16] The reason is straightforward. Take $k \in N \setminus S$. If k can and does certify $S \cup k$, then it improves (all other things being equal) the estimated worth of $S \setminus i$ (for $i \in S$) but leaves the belief about $S \setminus i \cup k$ unchanged (this coalition is known by assumption). Consequently, the certification of $S \cup k$ decreases k's marginal contribution to $S \setminus i$. Finally, it suffices to consider a solution φ that is (SSM) but strongly biased towards this very marginal contribution $v(S \setminus i \cup k) - v(S \setminus i)$.

Example 2.4. Take $N = (1234)$ and $\overline{v}$ defined by $\mathcal{M}(\overline{v}) = \{(12), (13), (14)\}$. Let $\mathcal{X}$ be the set $\{(1), (12), (123), (124)\}$. In particular, it is commonly known that (13) is winning. Suppose μ is uniform on $\mathcal{G}(\mathcal{X})$. One can show that if 3 certifies (123), then the estimated worth of (1) is 1/4; if 3 does not certify (123), then the estimated worth of (1) is 1/6. Therefore, certifying (123) decreases 3's marginal contribution to (1).

Interestingly, player 3 prefers to certify his winning coalitions and not his losing coalitions if $\varphi = Sh$. It means that, even though he deteriorates his marginal contribution to (1), he improves his overall position. This should not come as a surprise because μ is uniform (on $\mathcal{G}(\mathcal{X})$) and the

[16] One cannot have $\mathcal{X} = \{S, S \cup k, S \setminus i : k \in N \setminus S, i \in S\}$: this definition is ill-posed. Indeed, take $i \in S, k \in N \setminus S$. By assumption, $S \cup k \setminus i \notin \mathcal{X}$: this coalition is known. By superadditivity, this coalition reveals the worth of either $S \cup k$ or $S \setminus i$.

Shapley value is particularly balanced and rewards the players irrespective of their labels and gives the same weight to coalitions of the same cardinal.

This leads to the following open question. Say μ treats the players equally if for all $S, T \in \mathcal{X}$, $\#S = \#T$ entails that $v_\mu(S) = v_\mu(T)$: two unknown coalitions of the same cardinal have the same expected worth. In this setting, is it always optimal for a player to certify his winning coalitions and not his losing coalitions when μ treats all the players equally and the cooperative solution is the Shapley value?

2.4 Extensions

I sketch three possible extensions of my model. The first one investigates the possibility to extend the model to non simple games. The second extension enlarges the set of coalitionnal games by giving a positive probability to monotonic games instead of superadditive games only. The third extension changes the role of the social planner and allows her to have personal preferences; the proximity to the principal-agent paradigm is highlighted.

The first extension obvious: I consider only simple games, which model very well the undertaking of large-scale projects. However, there are countless situations in which it is preferable to consider coalitional games not restricted to $\{0, 1\}$. A first step would be to allow for three worth, say $\{0, 1, 2\}$. I conjecture that in many similar cases as in the article, players would have interest to certify their best coalitions (of worth 2) and not their worst (of worth 0). Yet, there could be issues for the social planner to interpret a coalition that is not certified because she might not know whether its worth is 0 or 1. This would mean that it would be even harder to reason counterfactually and to deduce full information about $\overline{v}$.

The second extension would be to drop the assumption of superadditivity and suppose that μ gives positive probabilities to monotonic simple games only.[17] In the motivation, it would mean that the required resources are not so scarcely distributed across the cities: this does not seem unrealistic. However, the issue of considering monotonic games is that it can yield highly

[17] A coalitional game v is monotonic if $S \subseteq T \implies v(S) \leq v(T)$.

Profile / Coalition	(1)	(2)	(3)	(12)	(13)	(23)	(123)
a^*	1/2	**2/5**	2/5	1	1	4/5	1
(σ_1, a^*_{-1})	5/13	**5/13**	4/13	1	10/13	10/13	1

Table 2.3: Undesirable beliefs in case of monotonic games

undesirable belief updates. In general, if S gets certified, then this can increase the expected worth of $T \subseteq N \setminus S$ (according to the new belief). Consider the following three-players example. Here, the certification of (13) increases the expected worth of (2) from 5/13 to 2/5:

Example 2.5. $N = (123)$, coalitions $(1), (12), (13)$ and (123) are winning, the others are losing. Suppose that μ gives the same (positive) probability to each monotonic simple game.[18] Suppose that player 1 responds to a^*_{-1}; call σ_1 the strategy "player 1 certifies only (12)". The results are given in Table 2.3.

One can see that under strategy a^*_1, the expected worth of coalition (2) is higher (namely 2/5) than under strategy σ_1 (namely 5/13). Therefore, certifying (13) is detrimental to player 1's marginal contribution to (2) even though (13) and (2) are independent with respect to monotonicity and superadditivity. The underlying mathematical reason is that there are proportionally more coalitional games v with $v(2) = 0$ when $v(13) = 0$. It is also worth checking that the coalitional game $\overline{v}$ is superadditive: this undesirable belief only occurs because the social planner considers that all the monotonic games are possible.

This example shows that monotonicity is a property hard to work with in this model. In the paper, I rule out this belief updating behaviour. Better understanding to what extent it is possible to consider monotonic priors is therefore of great interest to generalise this model.

The third extension is slightly different to the problem I investigated. I have supposed throughout the paper that the social planner is benevolent. There she has no personal interest and she only wants to allocate a vector that is as close as possible to the value of the real game $\overline{v}$. In particular, the social planner has neither available actions, nor well-defined preferences: it is not a player but a mere algorithm. This assumption is highly realistic in many

[18] Here, the social planner knows that $\overline{v}$ is monotonic (and that $\overline{v}(N) = 1, \overline{v}(\emptyset) = 0$), but has no further information.

economic cases–as is the State in her relationship with the different cities–but lacks credibility in other cases.

I rapidly sketch a tentative construction of a genuine player in *lieu* of a social planner. The idea is that if the social planner has mixed interests between the revelation of information and other orthogonal interests, she can choose the cooperative solution φ that best realises her orthogonal interests.

Suppose for example that the social planner wants to give the same allocation to all the players. Then pick any (SSM) solution φ that is a function of the marginal contributions only (e.g. the Shapley value). Then $\psi_\delta := \delta\varphi + (1-\delta)\mathbf{1}/n$ is (SSM), no matter how small δ is (and remains positive). There, the social planner has the best of both world for a very small cost (encoded in the δ): she gets the information about the game $\overline{v}$ and satisfies her orthogonal interests (under the hypotheses $(H1)$ and $(H2)$). Note that this example is credible in our leading example. Indeed, the wish to reward the cities exactly in the same way regardless of their strength/ability to build the treatment plant amounts to favour the weakest cities: this is a distributive policy oriented towards the less well-off.

Another example sees the social planner favouring a particular player. Suppose indeed that, for orthogonal reasons, the social planner wants player $i \in N$ to have a larger share than his peers. Then the solution φ^i_δ defined by $\varphi^i_{\delta,i}(v) = (1-\delta) + \delta\varphi_i(v), \varphi^i_{\delta,j}(v) = \delta\varphi_j(v), j \neq i$, fulfills both interests (up to a small δ). Obviously, this example no longer fits with our leading example. Indeed, the planner is here no longer benevolent and favours player i due to his label only, not due to his objective characteristics. There, the proximity with the principal-agent model arises. One can indeed view the $\delta\varphi_j(v)$ payment as satisfying an incentive constraint of the players (for them to reveal their winning coalitions–there is no participation constraint here, the players are compelled to participate to the two-stage game). Then, among all the mechanisms/cooperative solutions that permit the principal/planner to obtain the information she wants, she selects the best mechanism to satisfy orthogonal preferences.[19]

[19]In my setting, the players must be strictly better off to certify a coalition. It would not make sense, here, to assume that a player wants to certify his coalition by default–that is, unless it strictly harms him. It means that the inequalities must be strict. Therefore, there is not best mechanism for the principal (the

This idea of giving preferences to the planner, while not perfectly suitable to the leading example of this paper and to some extent not completely aligned to the question that is addressed, seems to be an interesting extension to be pursued.

2.5 Proofs

In this Section, I shall prove the results stated in Section 2.3. To that end, I shall need some intermediary results here and there.

Proof. (Of Prop. 2.1.) Consider any *interim* belief. Take $S \subseteq N$ a coalition of two or more players. The profile is such that no player of the coalition certifies its worth. As unanimity is required, no single deviation (of a player switching from "no certification" to "certification") can change the fact that no direct certification will take place. Not certifying S is therefore a best response for each player of S.

Let $i \in N$. For every *interim* belief, i knows whether $\{i\}$ is a winning coalition or not and can decide whether its worth will be certified or not. If $\{i\}$ is winning then by superadditivity of the prior, i knows that $\overline{v}$ is such that $\overline{v}(S) = 1$ if, and only if, $i \in S$. If $\{i\}$ is certified then the social planner knows that all the marginal contributions of i to $T \subseteq N \setminus i$ are equal to 1: this is the highest possible. As φ is (SSM) it yields the largest utility to i. It is therefore (strictly) optimal for i to certify $\{i\}$ when the latter is winning. By the same token, one can prove that if $\{i\}$ is losing then it is optimal for i not to certify it. □

Proof. (Of Prop. 2.2.) It is sufficient to provide an example. Take $N = (1234), \mathcal{M}(v_1) = \{(12)\}$, $\mathcal{M}(v_2) = \{(13), (14)\}, \mu(v_1) = \mu(v_2) = 1/2$. Then it makes player 1 worse off to certify (12) when $\varphi = Sh$ (*ex ante*).

Suppose first that the game that Nature picks is $v_1 = \overline{v}$. The set of unknown coalitions is $\mathcal{X} = \{(12), (13), (14), (134)\}$. First, coalition (13) (resp. (14)) is not certified in equilibrium because it would inform the social planner that player 3 (resp. 4) is a null player. The same holds

δ must be positive); it is at best ε-optimal. Mathematically, we lack closedness.

for (134). Second, player 2 wants (12) to be certified because he wants the social planner to know that $\overline{v} = v_1$. Yet, it is detrimental for player 1 when the social planner knows the worth of (12) because 1 is "stronger" overall in v_2 than in v_1 in the sense that $Sh_1(v_1) = \frac{1}{2} < Sh_1(v_2) = \frac{2}{3}$. To sum up, only player 2 wants some information to be transmitted whereas players $1, 3$ and 4 want the social planner to remain uninformed. Therefore, the following strategy profile (restricted to the unknown coalitions) is a Nash equilibrium: $a_2(12) = 1, a_i(S) = 0$ otherwise.

If Nature picks $v_2 = \overline{v}$, on the other hand, there is indeed revelation of information in equilibrium. All in all, whether winning or not (depending on the draw of Nature), coalition (12) is not certified.

The extension to $n \geq 4$ arbitrary is immediate. It suffices to use the same coalitional games v_1 and v_2 and to extend them to players $5, 6, ..., n$ by assuming that these players are null players in both games: it writes (again) $\mathcal{M}(v_1) = \{(12)\}, \mathcal{M}(v_2) = \{(13), (14)\}$ and the argument applies *mutatis mutandis*. □

What happens in the example of the proof is that the heuristics–certify all and only one's winning coalitions–is true only all other things being equal. Yet, certifying a winning coalition can be detrimental to other orthogonal marginal contributions. Here, certifying (12) indeed improves player 1's marginal contribution to (2) but worsens his marginal contributions to (3) and (4).

Proofs of Propositions 2.3, 2.4 and 2.5. In order to prove the three propositions, I first need to define what independent coalitions are. Let $S, T \in 2^N$. Say that S and T are *independent* (with respect to superadditivity) when: $S \cap T \neq \emptyset, S \not\subseteq T$ and $T \not\subseteq S$. Two independent coalitions can both win at the same time, both lose at the same time, can have different worth. Two minimal winning coalitions are independent. So are $S \setminus i, S \setminus j$ (provided $\#S \geq 3$) and $S \cup k, S \cup k'$.

In order to prove Propositions 2.3, 2.4 and 2.5, we need an intermediary lemma to understand what is at work.

Lemma 2.1. *Suppose that the coalitions of $\mathcal{X}$ are pairwise independent. Then $|\mathcal{G}(\mathcal{X})| = 2^{|\mathcal{X}|}$ and for each $S \in \mathcal{X}$, there are exactly $2^{|\mathcal{X}|-1}$ possible games for which S is a winning coalition.*

Proof. Let us prove the lemma by induction over the cardinal of $\mathcal{X}$. If $\mathcal{X} = \{S\}$, it is easy: all the coalitions but S are known and S can take two values. Then, $\mathcal{G}(\mathcal{X}) = \{v_1, v_2\}$ with $v_1 = v_2 = \overline{v}$ on $2^N \setminus \{S\}, v_1(S) = 0, v_2(S) = 1$.

Let p be an integer. Suppose that the result holds for $|\mathcal{X}| = p$. Let us prove that it holds for $|\mathcal{X}| = p + 1$. To that end, write $\mathcal{X} = \{S_1, ..., S_{p+1}\}$. By definition, $\mathcal{G}(\mathcal{X}) = \{v \in \mathcal{G} : \forall T \notin \mathcal{X}, v(T) = \overline{v}(T)\}$. We can partition $\mathcal{G}(\mathcal{X})$ in two sets $\mathcal{G}(\mathcal{X})^+$ and $\mathcal{G}(\mathcal{X})^-$ where

$$v \in \mathcal{G}(\mathcal{X})^+ \text{ (resp. } v \in \mathcal{G}(\mathcal{X})^-) \iff v(S_{p+1}) = 1 \text{ (resp. } v(S_{p+1}) = 0) \tag{2.1}$$

Pose $\psi : \mathcal{G}(\mathcal{X})^- \to \mathcal{G}(\mathcal{X})^+, v \to \psi(v)$ with $v = \psi(v)$ on $2^N \setminus \{S_{p+1}\}$. As S_{p+1} is independent of each $S_t, t \in [p]$, S_{p+1} and S_t can be winning at the same time: $\psi(v)$ is indeed a superadditive game. ψ is thus a well-defined function. It is easy to see that ψ is injective. By noting with a slight abuse of notation $\psi^- : \mathcal{G}(\mathcal{X})^+ \to \mathcal{G}(\mathcal{X})^-$ the function that makes the reverse operation, we observe that $\psi \circ \psi^{-1} = Id$.

Hence $\mathcal{G}(\mathcal{X})^-$ and $\mathcal{G}(\mathcal{X})^+$ are in bijection and have the same cardinal.

Now, call $\mathcal{Y} = \{S_1, ..., S_p\}$. If, without loss of generality, $\overline{v}(S_{p+1}) = 1$, then $\mathcal{G}(\mathcal{Y}) = \mathcal{G}(\mathcal{X})^+$. By assumption, the $S_t, t \in [p]$ are pairwise independent and we can apply the induction hypothesis, $|\mathcal{G}(\mathcal{X})^+| = 2^p$. Finally, $\mathcal{G}(\mathcal{X})^-$ and $\mathcal{G}(\mathcal{X})^+$ partition $\mathcal{G}(\mathcal{X})$ so $|\mathcal{G}(\mathcal{X})| = 2^p + 2^p = 2^{p+1}$. There are half of the coalitional games in $\mathcal{G}(\mathcal{X})^+$ for which $S_t, 1 \leq t \leq p$, wins by the induction hypothesis. The same holds in $\mathcal{G}(\mathcal{X})^-$ in virtue of ψ. Therefore, we have the same result in $\mathcal{G}(\mathcal{X})$. This obviously holds for S_{p+1} in virtue of ψ. □

Proof. (Of Prop. 2.3.)

It is now fairly easy. The coalitions $S_1, S_2, ... \in \mathcal{X}$ are pairwise independent. Hence certifying or not S_t does not have any influence on the estimated value of $S_{t'}, t \neq t'$, by the social planner. We can isolate the strategy regarding each coalition from the others.

By Lemma 2.1, without certification of S_t, its expected value by the social planner is 1/2. Hence players of S_t strictly improve its expected value by certifying it–and their marginal contribution accordingly. This means that a player is uniformly better off (in the sense of Definition 2.2)

when a winning coalition he belongs to gets known by the social planner. This said and recalling that φ is (SSM), it is easy to characterise the set of equilibria, and the strict subset of strong equilibria. □

Remark 2.2. Take $S_1, ..., S_p \in \mathcal{M}(\overline{v})$. It is not possible to extend this result to

$$\mathcal{X} = \{S_t, N \setminus S_t : 1 \leq t \leq p\}$$

. Indeed, the fact that two coalitions S and T are independent does not entail that the property of independence holds for the complements. In detail,

$$S \text{ and } T \text{ independent } \implies S \cap N(\setminus T) \neq \emptyset, S \not\subseteq (N \setminus T) \text{ but } (N \setminus T) \subseteq S \text{ is possible}$$

The argument of the proof of Proposition 2.3 falls down and a counterexample exists.

We have all the necessary tools to prove Propositions 2.4 and 2.5. Recall that a_i^* is the strategy "i certifies all and only his winning coalitions"; call σ_{-i} the following adverse strategy: $\sigma_j(S) = 1$ if S wins *or* $i \in S, \sigma_j(S) = 0$ else (for all $j \neq i$). σ_{-i} is so defined that, on every $S \ni i, i$'s response "decides" whether S is certified or not. I shall prove that i's *strict* best response precisely is a_i^*.

Proof. (Of Prop. 2.4.)

Notice that only the $i \in S$ are strategic players. Indeed, $k \notin S$ does not belong to any unknown coalition. He thus plays no role and can be skipped in the proof

I shall consider three cases: S is a singleton, a pair, none of the two first cases. If $S = \{i\}$, then i is a dictator and it is trivial to prove that it is a strict best response to certify he is a dictator to the social planner. If $S = (ij)$ and j certifies (ij) but not (j), $\mathcal{G}(\mathcal{X}) = \{v_1, v_2, v_3, v_4\}$ where $v_k = \overline{v}$ on $2^N \setminus \mathcal{X}$. The v_k's are detailed in Table 2.4. For player i, $v_2 \succ_i v_k, k \in \{1, 3, 4\}$. Any strategy yields a convex combination of these four games. Playing a_i^* ensures the best vector of

	$v(S)$	$v(i)$	$v(j)$
v_1	1	0	0
v_2	1	1	0
v_3	1	0	1
v_4	0	0	0

Table 2.4: Case $|S| = 2$

coefficients for the v_k's. It is therefore a strict best response. The spirit is the same for player j.[20]

Let us finally suppose that $|S| \geq 3$. Take $i \in N$ and suppose he faces a_{-i}. Notice that $i \in S$ knows that the $S \setminus j, j \neq i$ lose, but he does not know whether $S \setminus i$ loses. Call τ_i the following strategy: i does certify S but does also certify the $S \setminus i_l, (i_l) \subseteq S \setminus i$ (l such coalitions are not certified). Let us prove that

$$V^* := \mathbb{E}_{\hat{\mu}(a_i^*, a_{-i})(\mu^{\overline{v}})}[v] \succ_i \mathbb{E}_{\hat{\mu}(\tau_i, a_{-i})(\mu^{\overline{v}})}[v] =: V \tag{2.2}$$

By assumption, $|S| \geq 3$. It entails that the $S \setminus i_l$ are pairwise independent. Besides, the $S \setminus i_l$ are unknown. It means that the $N \setminus (S \setminus i_l)$ lose. In the same spirit as Lemma 2.1, there are therefore 2^l possible coalitional games. *Ex interim*, $V^*(S \setminus i) = V(S \setminus i) = 1/2$ out of pairwise independence and uniformity of μ: whatever i does, does not have an influence on the belief of the social planner regarding the worth of $S \setminus i$.

If $S \setminus i_j$ is certified (losing), then $V(S \setminus i_j) = 0 < V^*(S \setminus i_j)$. Consequently, V and V^* are equal on $2^N \setminus \mathcal{X}$ and on $\{S \setminus i\}$ (to 1/2, due the uniformity of μ), $1/2 = V^*(S \setminus i_j) > V(S \setminus i_j)$. This shows Equation 2.2 for this class of strategies τ_i.

Finally, there is a single possible game v_0 where S loses. It is easy to prove that any $v \in \mathcal{G}(\mathcal{X}), v \neq v_0$ is uniformly strictly preferred to v_0. Consequently, not certifying S is a strictly dominated response. This proves Equation 2.2 in general.

This shows that every player $i \in S$ is uniformly better off when S is certified, when the $S \setminus j, j \neq i$ are not. It is then easy to characterise the set of equilibria and to identify which equilibria are strong. □

[20] μ uniform is not necessary for $|S| = 1, 2$.

Proof. (Of Prop. 2.5.)

If $|S| = n-1$, it is trivial. Suppose that $|S| \leq n-2$. I shall prove a slightly stronger result: S and some $S \cup k, k \in N \setminus S$ are unknown (but not necessarily for all k). There are two steps: (1) prove that playing a_i^* is a strict better response to a_{-i}^* for $i \in S$ than any strategy τ_i that does not certify S; (2) prove that playing a_j^* is a strict best response to a_{-j}^* for $j \in N \setminus S$, always in the sense of Definition 2.2.

(1). Let $i \in S$. Let τ_i be the following (pure) strategy: i does not certify S, does not certify the $S \cup k_l, (k_l) \subseteq N \setminus S$ (the inclusion can be strict, l such coalitions are *not* certified). Call $V^* := \mathbb{E}_{\hat{\mu}(a_i^*, a_{-i}^*)}[v], V = \mathbb{E}_{\hat{\mu}(\tau_i, a_{-i}^*)}[v]$. Let us show that $V^* \succ_i V$. First, $V^* = \overline{v}$.

There are $1+2^l$ games compatible with the strategy profile (τ_i, a_{-i}^*) , $1+2^{l-1}$ for which $S \cup k_l$ wins. Indeed, there is $\overline{v}$, the only one for which S wins. Then, the $S \cup k_l, k_l \in (k_l)$ are pairwise independent. We can apply Lemma 2.1. Therefore, $V(S) < 1 = \overline{v}(S), V(S \cup k_l) < 1 = \overline{v}(S \cup k_l)$. V and $\overline{v}$ are equal elsewhere. Thus $\overline{v} = V^* \succ_i V$. (We do not need the uniformity of μ for Point (1).)

(2). Take $j \in N \setminus S$, now. j does not know whether S and the $S \cup k_l, k_l \neq j$ win. For j, if S loses, then $S \cup j$ cannot be certified indirectly. Certifying this coalition does not change the belief of the social planner regarding the worth of $S \cup k_l, k_l \neq j$ *because* the prior μ is uniform.[21] It is therefore sufficient to prove that the certification of $S \cup j$ strictly improves the marginal contribution of player j to coalition S.

Recall j is situated *ex interim*. He believes this: for the social planner, $1 + 2^{n-s}$ coalitional games are possible. These games are partitionable in three sets: $\{\overline{v}\}, A^+, A^-$. A^+ (resp. A^-) is the set of coalitional games for which $S \cup j$ wins (resp. loses). A^- is non empty. Call τ_j the strategy "j does not certify $S \cup j$."

Then one can verify that *ex interim* for j

$$V^*(S \cup j) - V^*(S) = \frac{\#A^+}{1 + \#A^+} > \frac{\#A^+}{1 + \#A^+ + \#A^-} = V(S \cup j) - V(S)$$

[21] It would be false in general without the assumption of uniformity.

hence the desired result.

So far, we known that certifying a winning coalitions increases or leaves unchanged every marginal contribution, certifying a losing coalition does the opposite. It is therefore easy to characterise the set of equilibria and to identify which of them are strong. □

Remark 2.3. The proofs show that it is possible for players to play mixed strategies at the equilibrium, but *only* when they are indifferent and they cannot change the belief of the social planner. That is, when at least one other player certifies a coalition with probability zero (whatever the others do, the coalition is not certified directly), or when a player knows that he needs not certify a winning coalition $S \cup k$ because it is certified indirectly by the strategy profile (this holds for $i \in S$).

Therefore, the possibility to mix at the equilibrium is illusory. In particular, if it is a best response for player i to certify a coalition with probability $p \in]0,1[$, then it is a best response to certify this coalition with probability $p' \in [0,1], p' \neq p$. In other words, it cannot be a *strict* best response to certify one coalition with probability p.

Finally, the condition "μ is uniform on $\mathcal{G}(\mathcal{X})$" can be slightly relaxed, provided the φ_i are continuous in addition.

Proposition 2.6. *Suppose the φ_i are continuous. Proposition 2.4 holds in an open neighbourhood of the uniform prior $\overline{\mu}$.*

Proof. Let us formalise this. One can see the prior μ as a vector of $\mathbb{R}^{|\mathcal{G}(\mathcal{X})|}$ with non-negative entries that sum to 1.

Take a pure strategy profile b. Denote

$$\mathcal{G}(b) = \{v \in \mathcal{G} : \forall S \text{ certified by } b, v(S) = \overline{v}(S)\}$$

and with a slight abuse of notations

$$\mathbb{E}_\mu[\mathcal{G}(b)] := \sum_{v \in \mathcal{G}(b)} \frac{\mu(v)}{\mu(\mathcal{G}(b))} \times v$$

Now, for all $i \in S$, for all pure strategy b_i,

$$F_{b_i} : \begin{cases} \mathbb{R}^{|\mathcal{G}(\mathcal{X})|} & \to \mathbb{R} \\ \mu & \to \varphi_i(\mathbb{E}_\mu[\mathcal{G}(a^*)]) - \varphi_i(\mathbb{E}_\mu[\mathcal{G}(b_i, a^*_{-i})]) \end{cases}$$

is a continuous function. Indeed, for every $\mathcal{F} \subseteq \mathcal{G}, \mu \to \mathbb{E}_\mu[\mathcal{F}]$ is continuous and so is φ_i by assumption. Therefore, the set $\{F_{b_i} > 0\}$ is an open set (as the reciprocal image of $\mathbb{R}^*_+$). A *finite* intersection of open sets is open:

$$O := \bigcap_{i \in S} \bigcap_{\substack{b_i \text{ pure} \\ \text{strategy}}} \{F_{b_i} > 0\}$$

is open.[22] By construction, a prior $\mu \in \mathbb{R}^{|\mathcal{G}(\mathcal{X})|}$ makes Proposition 2.4 true if $\mu \in O$. This proposition proves that $\overline{\mu} \in O$. Therefore an open neighbourhood of $\overline{\mu}$ lies in O. Such elements of the neighbourhood are *not* uniform over $\mathcal{G}(\mathcal{X})$.

(The same proof works for $\mathcal{X} = \{S, S \cup k\}$.) □

[22] We can restrict ourselves to pure strategies as only pure strategies are played at the equilibrium provided a player is not indifferent. Therefore, there are finitely many pure strategies and the finite intersection of open sets is indeed open.

3 Ideological Competition and the Taming of Experts

3.1 Introduction

The ones who decide need not be the ones who know. This is particularly the case for the members of Parliament (she) who lack precise knowledge that would be relevant to vote for a bill. To remedy this, they seek for advice. Summoned experts (he) are compelled to show up and are dismissed at will by members of Parliament. Both the majority and the opposition may summon experts. Consequently, experts represent the whole political spectrum: there is *ideological competition*. Conservative or liberal experts are well identified.

While their preferences are common knowledge, their sincerity should not be taken for granted. An expert before the Parliament may defend a policy that he thinks best responds to his ideological preferences, whether or not it is suited to the current situation. For instance, a conservative expert could advocate for tax cuts, be it relevant in the economic situation or not.

The competence of an expert, namely his ability to provide good advice, is private information. Only by interviewing an expert can the members of Parliament learn his competence. Experts are legally bound to tell the truth. However, intentional deception is difficult to identify. First, experts are unequally competent (hidden competence); second, even the best expert turns out to be wrong here and there (imperfect competence). Suppose an advice turns out to be wrong *ex post*. Is the advice wrong because the expert in place provides bad advice in general? Is it

because the good expert had bad luck? Or is it because the expert pushed his own agenda to the detriment of the decision maker? Is the error sincere or opportunistic?

In general, the nature of the advice relationship is repeated. So long as the members of Parliament are satisfied with the advice they receive from an expert, they tend to keep him. Only from time to time do they switch to a new one. Whenever the members of Parliament are unsatisfied with the expert in place, they are free to dismiss him and to summon a fresh one.

A competent expert faces the following trade-off. On the one hand, he wants to recommend his preferred action regardless of his signal to get short term payoffs. On the other hand, he must provide sufficiently good advice for the decision maker to prolong the relationship. He must balance his payoff today with anticipated future payments. To make the threat credible, firing an expert must yield at least the same expected payoff to the decision maker than keeping him.

Competing with another expert may change the trade-off faced by each expert and relax the condition of credibility for the decision maker. Indeed, a competent expert faces not only the threat to be fired and replaced by a fresh expert who shares his bias, he also faces the present advice of an expert with an opposed bias. The decision maker, the argument goes, could have more freedom to fire one expert because her continuation payoff is given not by his substitute but by his competitor: the threat to fire gets more credible. In turn the expert is efficiently tamed by competition to give good advice.

However, this line of thought rests on the assumption that an expert with an established reputation gives better advice than a fresh one. When there is very little uncertainty about the competence of a fresh expert, or when competent expert perfectly observe the state of the world, this does not necessarily hold. The effect of adding competition is therefore not trivial.

In light of identified and transparent preferences but private competence of experts, this chapter tackles the question: Is ideological competition beneficial to the members of Parliament? More specifically, should experts be hired for fixed contracts instead of being freely dismissed? Is permanent sincerity, *i.e.* truthful revelation of the signal, sustainable in equilibrium? To what extent is competition useful to the decision maker when she faces both imperfect and hidden competence? Is it still the case when one of the latter conditions falls?

I define a repeated game of cheap talk between one decision maker (the member of Parliament) and two pools of experts–each pool represents an ideological preference. At every period, the decision maker faces at most one expert of each pool. Each expert cheap talks and recommends one of two actions to the decision maker. The former prefers that the action chosen fits his (transparent) ideology irrespective of the state of the world. The latter wants the action to match with the state of the world (which is drawn at every period). The interests of the experts and the decision maker for the current period are sometimes aligned, sometimes opposed. Under ideological competition, at every period exactly one expert has aligned interests with the decision maker; yet she does not know who he is.

The expert may be competent or incompetent. Only competent experts receive an informative signal about the state of the world. The recommendation of the expert is cheap talk: it is costless, non-binding and unverifiable *ex interim*. An expert is sincere if he truthfully reports his signal. After the decision maker takes an action, everyone learns whether its matches with the state of the world and she decides whether she wants to dismiss or to keep the experts in place. The game resumes. Over time the decision maker forms a probabilistic estimate that the expert is competent, his *reputation*.

No communication can be informative in the static game because the expert wants to push his own agenda in the short term and does not fear to lose future payments. By backward induction, the same holds if the game is repeated a finite number of time. Fixed term contracts prove unsuccessful to promote good advice.

There can be meaningful communication in equilibrium when the game is infinitely repeated. I show that it is a dominant strategy for the decision maker always to follow the recommendation of the expert; and to fire immediately those who are likely incompetent. In any Pareto-optimal equilibrium, an expert recommend his favourite action when he believes it matches with the state of the world–that is, when the short term interests of the expert and the decision maker are likely to coincide.

I distinguish two by two cases. Whether there is hidden competence or not (namely, there

are only good experts); whether there is imperfect competence or not (namely, good experts perfectly observe the state of the world). When there is no hidden competence, the decision maker is not better off with competing experts. Indeed, when all the experts have the same competence, there exists an equilibrium where experts are permanently sincere and the decision maker fires an expert whenever he proved insincere. This is the best the decision maker can hope for payoff-wise. Likewise, when competent experts are perfect, them being permanently sincere can be sustained in equilibrium–be there ideological competition or not.

When there is both hidden competence and imperfect competence, the decision maker fares strictly better when she faces two competing experts: ideological competition *does* tame the experts, they lie less when they have a competitor. Competition increases the payment of the decision maker through two channels. Directly, two competing experts with the same reputation are sincere no matter how high their reputations (this cannot be sustained without competition). Indirectly, a good expert with higher reputation must lie less than would he face no competitor. Indeed he knows that, due to the direct channel, the "outside option" of the decision maker (after she fires him) is larger under competition. Permanent sincerity does not happen in equilibrium.

Related literature. This chapter relates to two streams of literature: the cheap talk literature that begins with Crawford and Sobel (1982) and the literature about reputation that begins with Kreps and Wilson (1982a) and Milgrom and Roberts (1982). The closest article to mine is Schottmüller (2019). In the cheap talk literature, it is often assumed that the bias preference-wise between the expert and the decision maker is time invariant; in this article, the short term interests of the two parties may or may not coincide at every period.

My chapter is very close to Schottmüller (2019) but with small changes in the model and different questions of interest. In his article, the state of the world (namely the preferred option of the decision maker) changes at every period and so does the preferred action of the expert (independently). Importantly, the decision maker does not know what the current preferred action of the expert is. She can fire the expert at any time, in which case she earns an outside option and the game ends. In this chapter, I consider transparent motives: the decision maker

always knows what the experts prefer her to do. Also, I consider a cheap talk framework whereas recommendations are directly payoff relevant in Schottmüller (2019) (he explains however how one can reframe the model as a cheap talk setting, see his supplementary material). Besides, there is no outside option in this chapter and future payments of the decision maker are endogenous to the strategies because new experts step in and the game resumes.

More importantly, the questions differ. Schottmüller (2019) builds his article on the idea that an expert can be fired not in spite of having a high reputation, but because so. The reason is that, the higher the reputation of an expert, the more he has room to push his interests. The moral hazard becomes too problematic. Technically, it holds because reputation updates are infinitely slow in the neighbourhood of 1. I want to highlight the interest of ideological competition without an outside option. I do not always assume imperfect experts or two possible types of experts: I vary these parameters.

Sobel (1985) develops a theory of credibility in a finitely repeated game. The sender is either a friend or an enemy of the receiver, once and for all. The reputation refers to the honesty of the sender, not his competence or ability. Benabou and Laroque (1992) extend Sobel (1985) to imperfect private information. There, a sender can repeatedly frame the receivers because an error can be both sincere or opportunistic in principle–the same idea significantly holds in this article.

Mailath and Samuelson (2001) and Jullien and Park (2014) stage long-lived senders (sellers, companies) and short-lived receivers (buyers). The reputation of the sender is the probability of being competent. In Mailath and Samuelson (2001), the companies coexist and their managers come and go exogenously. New managers come in and may buy the reputations while consumers do not observe the change. It gets too difficult to incentivize good agents with high reputations to exert high effort. My paper differs to these papers because they focus on "honest Markov equilibria"–good agents exert high efforts–whereas I focus on quasi-markovian strategies and explain how sincerity cannot be sustained in all circumstances.

Besides, the competition between companies takes a different aspect: in Mailath and Samuelson (2001), each company produces a (personal) good and wants the consumer to buy its good.

In my setting, the two experts compete on the "same" state of the world (underlying) that is relevant to the decision maker. The same difference about how competition works holds with Park (2005) (but the receiver is also long-lived; also the sender knows perfectly the state of the world, I do not assume this).

Ottaviani and Sørensen (2006a) and Ottaviani and Sørensen (2006b) study several experts who want to maximise their reputations–which can be done by intentionally misrepresenting their signals. Here, experts want to maximise their payments in terms of fit between the action chosen and the state of the world. Having a good reputation is only a mean to an end.

More specifically related to the relationship between experts and members of Parliament is the literature about committees. It begins with Gilligan and Krehbiel (1989) who first consider a committee with partial and asymmetric information. The committee is better informed than the legislator and advises the latter (in a static framework, hence no reputational issue). In Austen-Smith (1990), three players play a three-stage voting procedure: cheap talk, decision of the agenda, vote of one bill of the agenda. Players want to influence the belief of others about the consequences of a public policy. My paper is different because experts and the decision maker are clearly separated (first), because the voting procedure is simpler but infinitely repeated (second).

Finally, Minozzi (2011) considers two politicians who communicate to the public by cheap talk (static). They perfectly observe the state of the world. Minozzi (2011) conceptualizes "jamming": at every period, some politician tells the truth while the other tells exactly the opposite (in equilibrium). The public is not able to distinguish the deceiver and the truth-teller. In this article, should experts send discording messages, it needs not be because one tries to jam. Indeed, the experts may be truthful but one received a wrong signal. Also, when I assume that the experts are perfect, incentives to jam are lower (up to inexistent) because the experts have dynamic incentives to tell the truth in order to survive.[1]

The rest of the paper goes as follows. Section 3.2 defines the model. Section 3.3 presents

[1]On the other side, Schnakenberg (2015, 2017) studies expert relationships with one expert and several voters (static). Farrell and Gibbons (1989) study abstract cheap talk with multiple receivers.

dominant strategies for the players. Section 3.4 characterises when ideological competition is beneficiary. Section 3.5 studies reputation-invariant equilibrium payoffs. Section 3.6 develops two extensions–experts receive a payment for advising *per se*, bad experts become strategic players– and concludes. All the proofs are relegated to Section 3.7. Section 3.8 sketches an alernative model with one state of the world per expert.

3.2 The model

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space on which all the random variables are defined.

The players. The experts (he) distinguish themselves along two parameters: the ideological preference and the competence. Call $E \in \mathcal{E}$ a liberal expert and write $E = A$ if the expert is competent/good, $E = B$ otherwise. The same notations with a prime denote conservative experts (e.g. A' is a competent conservative expert). A good expert is a genuine strategic player. A bad expert is pure noise. D is the decision maker (she). The sets of experts $\mathcal{E}$ and $\mathcal{E}'$ are infinite.

The ideological preference of every expert is common knowledge. The competence of each expert is his private information. $\mu \in [1/2, 1]$ is the prior probability that an expert is good (whichever his ideology).

At each period of time, D is advised by a single expert E in the non-competing case; by one expert E and one expert E' in the competing case.

The states of the world, the information structure. At every period of time, the state of the world is $\omega \in \{0, 1\}$ with $\mathbb{P}(\omega = 1) = 1/2$. Expert E privately receives a signal $s \in \{0, 1\}$ conditional on ω. If $E = A, \mathbb{P}(s = 1|\omega = 1) = \mathbb{P}(s = 0|\omega = 0) = p > 1/2$; if $E = B, \mathbb{P}(s = 1|\omega = 1) = \mathbb{P}(s = 0|\omega = 0) = 1/2$ (a bad expert is non informative). The value p is common knowledge. The signals of experts E and E' conditional on the state of the world are independent. The states of the worlds and the signals are independent across periods.

Reputation of the experts. The prior probability for an expert E to be of type A is exactly μ. As the relation of expertise goes on, the decision maker updates her belief about the competence of the expert (beginning at μ). I call reputation and denote r the probability, according to the decision maker, that the expert she faces is competent. It depends on the strategies (in equilibrium) and on whether the message fits the state of the world.

The strategies. All the players use quasi-markovian strategies. At every period, E receives a signal $s \in \{0,1\}$ and publicly sends a message $m \in \{0,1\}$ to D. B is noisy and sends $m = 1$ with probability $1/2$ at each period–irrespective of the signals. Assume without loss of generality that the experts directly recommend actions.[2] The probability for A to send message $m = 1$ at the current period depends only on the current signal and the reputations r and r'. So does E' simultaneously and independently. (Their strategies are indeed markovian, (r, r') being the state variable.)

The strategy of D has two aspects. First, the probability for D to choose action $a = 1$ given the messages and the reputations (markovian). After the action is taken, the payoffs are distributed and the state of the world is revealed. Then, the probability for D to fire expert E (or E') given the true state of the world, the messages and the reputations. The decision to fire an expert is therefore a function of the reputations *and* whether the messages fit the state of the world. It is not markovian, hence the term "quasi-markovian."

In particular, the decision to fire or not an expert is *not* contingent on whether the action chosen by D turns out to match with the state of the world. If E is fired, then a new liberal expert (also labelled E) steps in for the next period. His initial reputation is μ. The stage game resumes.

The payoffs of the stage game. At every period, the only payoff-relevant action is a. D receives 1 if, and only if, $a = \omega$ (that is, the action and the state match). Expert A is biased towards action 1 and receives 1 when $a = 1, 0$ otherwise. Conversely, expert A' is biased towards

[2] In principle, any message is possible in cheap talk. But every message sent entails the choice of an action by the decision maker in equilibrium. Up to a relabelling of the message, it can be phrased as a direct recommendation. Let us focus directly on this.

action 0 and receives 1 when $a = 0$.[3] After experts are fired, they receive a payment of 0 in all future periods.

All the players share the same discount rate $\delta \in]0, 1[$. (This is without loss of generality.) The players publicly observe whether the messages match with the state of the world. Consequently, at any time of the game, the reputation of an expert is common knowledge. Hereafter, ω is "high" when its value is 1, "low" otherwise–the same goes for the signals and the messages. Note that for A' a high signal is bad news while a low signal is good news.

3.3 First results

I focus on equilibria that are (i) in symmetric strategies and (ii) informative. Symmetry means that experts A and A' behave alike given their opposed preferences: the probability for A to send m given s and (r, r') equals the probability for A' to send $m' = 1 - m$ given $s' = 1 - s$ and (r', r). Symmetry means also that the strategy of D does not depend on the label of the expert but only on the quality of his advice and the reputations. D treats E and E' equally. Informativeness means that good experts give at least as good advice in average as bad experts. Consequently, if experts with reputation μ are perfectly sincere, it is a best response for D to fire any expert with reputation lower than μ. The condition that experts are perfectly sincere when their reputation is μ is innocuous: it is very easy for D to induce this behviour–all the equilibria below are such.

Lemma 3.1. *Suppose that $r < \mu$. Then it is a best response for D to fire E (whatever r' and whether the message fits the state of the world).*

Heuristics. When the state of the world is high, the interests (for the current period) of A and D are aligned whereas when it is low, the interests of A and D are opposed. Consequently, whenever A receives a high signal, he thinks that it is likely that his and D's short term interests

[3] Players need only observe their payments to guess what the state of the world was (indeed, the choice of action a is public). It is in this sense that they observe publicly the state of the world after D takes an action.

are aligned (because $p > 1/2$). It entails that whenever $s = 1$, we can expect A to truthfully reveal his signal. Contrariwise, when $s = 0, A$ may be tempted to lie and send $m = 1$ to ensure a short term payment.

Hence there are two possible significations to $m \neq \omega$. If $m = 1$ and $\omega = 0$, it is an "opportunistic error"; if $m = 0$ and $\omega = 1$, it is a "sincere error" due to incompetence. Note that, even though $m = 1$ and $\omega = 0$, it is possible that the expert was sincere: even good experts get wrong from times to times (with probability $1 - p$). It is also possible, in principle, that an expert A lies about a signal $s = 1$ and sends $m = 0$ because–under the equilibrium strategies–it is worth it to hide a good signal today to ensure higher payoffs tomorrow.

Nevertheless, if the expert is good, it is natural to interpret a sincere error as really sincere. I put a restriction on the firing strategy of the decision maker and I assume that the decision maker never punishes an expert by firing upon sincere error or when the message and the state match:

Assumption 3.1. Provided $r \geq \mu$, D never fires an expert upon sincere error or when the message fits the state of the world.

The following proposition says that, under condition of Pareto-optimality, it is always optimal for A to reveal a high signal (and for A' to reveal a low signal). It means, contrariwise, that a sincere error is *genuinely* sincere in equilibrium: for expert $A, m = 0$ implies that $s = 0$.

Proposition 3.1. *Suppose there are only E experts.*

- *In equilibrium, there is no $r \in [\mu, 1]$ such that when A has reputation r, he strictly prefers to send $m = 1$ when $s = 0$ and $m = 0$ when $s = 1$;*
- *Suppose that there exists $r \in [\mu, 1]$ such that when A has reputation r, he strictly prefers to send $m = 0$ regardless of the state of the world. Then there exists a Pareto improvement of this equilibrium in which A sends $m = 1$ in r, regardless of the state of the world.*

It is optimal for D to play the action that likely fits the state of world. It means that it cannot be an optimal strategy for D to make a short term sacrifice (and play an action that will likely fail) in order to control the future behaviour of the experts. The reason is that the

decision to fire or to keep an expert is contingent on whether the message and the state of the world match; it does not depend on whether the action chosen by D matches with the state of the world. The action chosen does not bear any consequence on future payments. Consequently, choosing the action that gives the highest expected payment today is (obviously) good for today and is immaterial with respect to the future.

Due to informativeness, in the one expert case, it means that it is optimal for D to follow the recommendation of the expert. The following result is an adaptation of the supplementary material of Schottmüller (2019), Section 2.

Proposition 3.2. *For any (r, r', m, m'), it is strictly optimal for D to play action high if, and only if, $\mathbb{P}(\omega = 1|m, m', r, r') > 1/2$. D is indifferent between the two actions whenever the former probability equals $1/2$.*

Assume that D breaks ties equally when she is indifferent between the two actions. In the following, we can therefore forget the choice of a by D and assume that the recommendations of the experts are directly payoff-relevant.[4] The only strategy for D is hence the probability $y(r, r')$ to *keep* an expert of reputation larger than $r \geq \mu$ given his opportunistic error and the reputation r' of the competing expert. (And to fire when the reputation is lower than μ.)

3.4 Is competition beneficial among experts?

This section addresses the following questions: Does competition among experts make them more sincere? And does it permit the decision maker to achieve a larger expected payment? The answers depend on p and μ in the following sense. First, if either there are only good experts ($\mu = 1$) or good experts perfectly observe the state of the world ($p = 1$), there exist equilibria where good experts are perfectly sincere–be there ideological competition or not–, which is the best D can have. Second, if there is both uncertainty about the type of the expert ($\mu < 1$) *and* good experts can receive wrong signals ($p < 1$), then for any equilibrium with only E experts,

[4] And that if two experts send different messages, D automatically follows the most credible recommendation.

$p \backslash \mu$	$\mu = 1$	$\mu < 1$
$p = 1$	Competition is useless payoff-wise	Competition is useless strategy-wise
$p < 1$	Competition is useless payoff-wise	Competition is beneficial strategy- and payoff-wise

Table 3.1: Summary of the benefice of competition among experts.

there exists an equilibrium with E and E' experts that yields a higher payoff to D (and good experts tend to be more sincere in the latter: competition *does* tame the experts).

In plain English, competition does not benefit the decision maker when either there are only good experts and no noisy ones, or when good experts are perfect. Competition does benefit the decision maker when she faces the two issues: hidden competence and imperfect competence.

As will be clear later, when $p = 1$, D receives a higher payment when there is competition, even though good experts adopt the same strategy whether there is competition or not. The reason is quite simple: in each case, D samples the experts until she faces a good one (or two good ones). The sampling procedure is faster in the competing case because two experts are tested at a time. Table 3.1 sums up the results.

No hidden competence: $\mu = 1$.

Proposition 3.3. *Suppose that* $\delta \geq \dfrac{2}{2p+1}$. *Then the following profile is a subgame perfect equilibrium: expert A truthfully reports his signals, decision maker D keeps expert A with probability* $y \leq \dfrac{\delta(2p+1)-2}{\delta(2p-1)}$ *in case of an* opportunistic *error (type* $m = 1, \omega = 0$*) and with probability* 1 *otherwise. Notice that* $y = 0$ *works. The expected payment of D under this profile is* $\dfrac{p}{1-\delta}$.

If $y = 0$, *then the expected payment of A is* $V_A = \dfrac{1}{2-\delta(1+p)}$; *if* $y = \dfrac{\delta(2p+1)-2}{\delta(2p-1)}$, *then* $V_A = \dfrac{1-1/2p}{1-\delta}$.

The same profile works in the same way when there are competing experts:[5]

[5]The decision maker could also ensure the same payment with two competing experts by only listening to one of them and ignoring the other one. This is *not* symmetric however.

Corollary 3.1. *Suppose that* $\delta \geq \frac{2}{2p+1}$. *The profile (each experts* A *and* A' *report truthfully their signals, they are kept with probability* y *given by Proposition 3.3 upon opportunistic mistake) is an equilibrium. This equilibrium provides the same expected payments for* D *and the experts.*

All in all, competition does not play any role when there are only good experts.

Perfect competence: $p = 1$**.** Perfect experts directly observe the state of the world. They never receive false signals and never get wrong. Errors cannot be sincere. Assume that the strategy of a perfect expert is to be always sincere. Then, no matter how close to 1 reputation r is, it falls to 0 when a single mismatch between the message and the state of the world occurs. Consequently, under the sincere strategy, a perfect expert does not have any incentive to lie here and there because he gets fired immediately. Conversely, it is a best strategy for D to fire him because she believes he is a noisy expert (and so yields her low payments). This explains why sincerity is a sustainable equilibrium strategy:

Proposition 3.4. *Suppose that* $\delta \geq 2/3$. *Suppose there are only* E *experts. The following profile is an equilibrium:* A *truthfully reports the state of the world,* D *fires an expert whenever the reputation of the latter falls below* μ, y *can be chosen freely. The* ex ante *expected payment of* D *is*

$$V_D(\mu) = \frac{\frac{\mu}{1-\delta} + \frac{1-\mu}{2-\delta}}{1-(1-\mu)\frac{\delta}{2-\delta}} \tag{3.1}$$

Suppose an expert sent a message that fit the state of the world in the first t *periods. Then his reputation after the* t^{th} *message is* $r^t = \frac{2^t \times \mu}{1+(2^t-1)\mu}$. *His reputation converges at exponential rate to* 1.

The decision maker simply samples the experts E until she faces a perfect expert A. It never happens that, after an error (opportunistic or sincere), the reputation of the expert is larger than μ. Dismissing due to incompetence is automatic and we need not specify y.

Unsurprisingly, there exists an equilibrium with competing experts where both A and A' are perfectly sincere:

Proposition 3.5. *Suppose that $\delta \geq 2/3$. Suppose there are both E and E' experts. The following profile is an equilibrium: A and A' truthfully report the states of the world; D fires an expert whenever his reputation is below μ. With the notations defined above, $r^t = \dfrac{2^t \times \mu}{1 + (2^t - 1)\mu}$.*

By the same token, the decision maker samples experts E and E' until she faces perfect experts A and A'. If $r < r'$, then D follows the recommendation of E' (in virtue of Proposition 3.2). All in all, experts A need not face E' to be perfectly sincere: they are tamed by the threat of being fired and replaced (by their fellows), not by their counterparts. Strategically, the behaviour of a perfect expert is the same whether there is competition or not. Interestingly, however, the expected payment of D is higher when there is competition than when there is none. The reason is straightforward: D "tests" two experts at a time and the sampling is faster. Computing $V_D(\mu)$ in Proposition 3.5 proves very nasty.

Hidden competence ($\mu < 1$) and imperfect competence ($p < 1$). Theorem 3.1 states that competition proves beneficial when there are both hidden competence and imperfect competence.

Theorem 3.1. *Suppose that $p < 1$ and $\mu < 1$. Suppose that $\delta \times p$ is large enough. Then for every equilibrium with only E experts, there exists an equilibrium with E and E' experts for which the expected payment of D is (strictly) higher.*

There are two channels through which the payment of D increases under competition. The proof of Theorem 3.1 goes along the following lines. Fix some equilibrium in the single expert case. Suppose that A and A' have the same reputation (no matter how high). Then it is possible to provide them the right incentives such that they are sincere. Schottmüller (2019, Theorem 1) proves on the other hand that, in the single expert case, sustaining perfect sincerity when the reputation is high is not possible. This is a *direct* channel through which the payment of D increases.

Assume that A' babbles whenever $r' < r$ and that D fires expert E' if $r' < r$.[6] Then, an expert A of higher reputation (than his peer E') is "alone." He and D cannot play the single expert case

[6]One needs here to define what *babbling* means. Does A' given (r, r') send a random message? (Tra-

equilibrium. Indeed the "outside option" of D, should he fire E under r, is higher when there are two experts of the same reputation (that is because of the direct channel). A therefore must lie less than prescribed by the former equilibrium in order not to be fired immediately. This is an *indirect* channel through which the payment of D increases.

Remark 3.1. In the proof, I assume that whenever $r' < r, A'$ babbles and D fires E'. It allows me to define an equilibrium with competing experts that outperforms the equilibrium with E experts only. One can wonder whether there are equilibria in which A' does not babble–and whether they are better for D. Suppose for instance that a good expert A' with lower reputation (than his peer) is always sincere. A good expert A with higher reputation, in turn, must lie "rarely" because he faces the immediate competition of A'. This would, the argument goes, be even more beneficial to the decision maker.

I do not have a definitive answer to this. Let me just point to the following fact. If A' is sincere, then D can expect $1/2 + r'(p - 1/2)$ at the current period should he listen to E'. If A lies with probability $1 - r'/r$ when he receives signal $s = 0$, then he exactly matches the payment given by E' to D. *Ex ante*, then, D is indifferent between following the recommendation of E or E'.

Ex interim, however, should $m \neq m'$, then the conditional probability that $\omega = m'$ is strictly larger and D always finds it optimal to follow the recommendation of E'. Consequently, A finds it under-optimal to lie even in the short term. He should lie less (than $1 - r'/r$) and increase the current payment of D. But then D is not indifferent between keeping and firing E. I do not completely understand this *ex interim* behaviour when A' does not babble and what an equilibrium would look like.

---

ditional definition of babbling.) Or does he always send message $m' = 0$ irrespective of the signal? (In order to get high short term payments.) The distinction is immaterial here because of Proposition 3.2: A is more credible than A' under (r, r') and D simply never listens to A'. In any case, firing E' is optimal for D given that A' babbles.

3.5 Equilibrium with constant payments for D

Let us get back to the case $p = 1$ where there are only E experts. In Proposition 3.4, the expected payment of D is an increasing function of the reputations. This section investigates the case where the current payment of D is a constant. Assume that $\mu = 1/2$. Then, the maximal current payment of D when $r = \mu = 1/2$ is $3/4$ (when a good expert is sincere). The maximal discounted payment of D where current expected payments are constant (in r) is hence $\frac{3/4}{1-\delta}$.

Theorem 3.2 states that this discounted payment is an equilibrium (for $p = 1$). It also states that the result does not hold for $p < 1$.[7] For all integer n, write $r_n = 1 - 1/n$. The condition $\mu = 1/2$ is necessary.

Theorem 3.2. *Fix $\mu = 1/2$ and suppose there are only E experts. Write $x(r) = \mathbb{P}(m = 1|s = 0, r)$.*

1. *Assume that the experts are perfect: $p = 1$. Suppose that $\delta \geq 2/3$. The following profile is an equilibrium: $x(r) = 1 - \frac{1}{2r}, y(r) = 3 - 2/\delta$.*

   For every $r \in [1/2, 1], D$ receives exactly $3/4$ per period in expectation, $V_A(r) = \frac{1/2}{1-\delta}$.

   Under any history, $r \in \{r_n : n \geq 1\}$. Before he is fired, the reputation of a noisy expert converges to 1 at an exponential rate but not monotonically so. The reputation of a perfect expert never falls below $1/2$ and he is fired only through the strategy y. He is fired in finite time almost surely (the rate of convergence to 1 is the same as the one of the noisy expert).

2. *Assume V_A is non-decreasing in r. If the experts are not perfect ($p < 1$), there exists no equilibrium for which the decision maker receives $\frac{p+1/2}{2}$ at every period (in expectation).*

The equilibrium described in Point 1 is Pareto-dominated (by the equilibrium of Proposition 3.3). Indeed, A might as well always tell the truth. He would earn $\frac{1/2}{1-\delta}$ whereas D would earn more than $\frac{3/4}{1-\delta}$ (her expected payment would there be increasing in the reputation r). Here, A earns more per period (up to $3/4$ in expectation when $r = 1$) but gets fired almost surely in finite time. He is greedy and his short term greediness is exactly compensated by he

[7] The corresponding current expected payment of D being $(p + 1/2)/2$.

being fired. Notice that the quality of the message sent by a perfect expert is decreasing in his reputation.

The result cannot be generalized to $p < 1$. The reason is that the reputation r of an expert will be sufficiently close to $1/2$ (in histories with positive probabilities), such that A will strictly prefer to be sincere in r. Consequently, in r, the current expected payment of D will be higher than $(p+1/2)/2$. I conjecture that the assumption that V_A is non-decreasing in r can be disposed of but this requires further developments.

3.6 Extensions and concluding remarks

In this section, I want to develop two possible extensions of the model. The first extension is to change the payment of the good expert. Expert A could earn $1 + l$ (when $a = 1$) or l (when $a = 0$). The quantity $l > 0$ is here the intrinsic payment of advising, irrespective of the actions chosen. The second extension is to get rid of the assumption of noisiness of the bad expert and to allow him to have preferences and strategies. Throughout the discussion of the two extensions, I will explain when Assumption 3.1 remains relevant, when it must be removed.

A receives a payment $l > 0$ for the sake of advising. So far, we supposed that an expert who has ideological preferences receives a payment only insofar as the decision maker takes decisions that he likes. A natural extension of this is to allow the expert to have two orthogonal payments: the ideological payment given above, and a payment $l > 0$ per period for the sake of advising. Indeed, advising the members of Parliament is a good signal for an expert. This payment may be expressed in terms of outside concrete opportunities, say.[8]

Heuristically, the larger l is, the more a good expert values advising *per se* over having his ideological preferences fulfilled. Consequently, the larger l is, the more the interests of the expert and the decision maker are aligned; hence the less he wants to lie upon a bad signal. Importantly,

[8]Recall that the Parliament does not pay the experts for advising (for legal reasons). Therefore, D does not bear the cost l. In a more general framework, one can assume that l is a periodical transfer from D to the expert in place. There, the opportunity for D to interview experts should be questioned from the start.

the concepts of "sincere" and "opportunistic" errors disappear with l as good experts want more and more to survive. Experts may prefer, because $l > 0$, to lie about favourable signals.

Which results fall down? And which results remain true? Proposition 3.3 and Corollary 3.1 do not hold when $l > 0$. The reason is straightforward. Good experts go wrong here and there (with probability $1-p$). If future payments are high, then a good expert A wants to send message $m = 0$ when $s = 1$ because he cannot be fired upon sincere mistake. He values his survival a great deal, which is detrimental to the decision maker. No transmission of information is possible any more (no matter how small l). To drive this issue out, one must get rid of Assumption 3.1: it must be possible to fire experts upon sincere mistakes for them to have incentives to reveal favourable signals.

Propositions 3.4 and 3.5, however, still hold. Again, this is straightforward. Any mistake entails a zero reputation and a replacement. Good experts, want to reveal truthfully unfavourable states of the world but also favourable ones (the concepts of "sincere" and "opportunistic" errors make little sense, but regardless experts are fired *because* their reputations fall to 0 and nothing else). I conjecture that Theorem 3.1 holds when $l > 0$.

Finally, Point 1 of Theorem 3.2 holds with the following modification: $y(r) = (3+2l-2/\delta)/(1+2l)$ (it is increasing in l and converges to 1 when l gets large); $\delta \geq 2/(3+2l)$. $x(r) = 1 - 1/2r$ is unchanged. Then it is an equilibrium with $V_A(r) = (1/2 + l)/(1 - \delta)$. Notice that the larger l, the larger y: the more a good expert values advising *per se* over his ideological preferences, the less he wants to lie. So the less "dangerous" a lie must be to make him indifferent, hence the larger the probability to be *kept* upon an opportunistic error must be.

A bad expert is strategic (instead of noisy). This extension is immaterial if we suppose that $\mu = 1$. We assume that a bad expert does not have any relevant information: his signals are not informative. This does not entail necessarily that he is noisy. Indeed, even bad experts may enjoy advising because it is a good signal. Likewise, even bad experts may have ideological preferences. (One needs not be Karl Marx to adhere to communism; or Friedrich Hayek to adhere to *laissez-faire*.)

Prescribing preferences to a bad expert is therefore a choice. Assuming that $l > 0$, I develop two possibilities. The payment of a bad expert is independent of the action taken by the decision maker (or aligned to her) because his lack of information entails a lack of conception of what would best fit him (i). The bad expert, even though he lacks information, has identified ideological preferences–hence A and B have exactly the same preferences, so have A' and B' (ii). Either way, the strategy of a bad expert depends only on the reputations, not on the signals. These statements must be made precise.

Under possibility (i), a bad expert is only interested in lasting as long as possible. He cannot correlate the message he sends to the state of the world. He can, however, try to "mimic" the strategy a good expert would adopt to fake his identity. The decision maker, in turn, is aware that this is what a bad expert wants to do. Under possibility (ii), we can expect that in many instances a bad expert wants to gain short term payments, no matter what: the cost of mimicking should be too high compared to the short term payment.

To be more precise, Propositions 3.4 and 3.5 still hold: perfect experts are always sincere, any error entails an expert to be fired. Under (i), a bad expert randomizes: indeed, he survives if, and only if, he sends the fitting message; and he has no information about which message is best. Under (ii), a bad expert B always sends $m = 1$ (and B' sends $m' = 0$) because it yields the highest short term payment and bears no consequence on whether he will survive or not.[9]

The question of whether Theorems 3.1 and 3.2 still hold, with which changes, and whether B always sends $m = 1$ (under (ii)) is unanswered. I conjecture that the decision maker must be allowed to fire experts upon sincere mistakes in order to mitigate the incentives for a bad expert to "play safe" and gain l at all time.

Finally, notice that this framework can be extended to other situations, like saver/financial adviser or patient/physician relationships. There, a patient must choose between two medicine

[9]It means that, if $m = 0$, automatically $E = A$. Given this, B might have an incentive to send $m = 0$: with probability $1/2$ he is fired, with probability $1/2$ he is identified as a perfect expert. We need to impose that if $r = 1$ but an error occurs, then the reputation drops to 0. With this belief prescription off-path, it is indeed optimal for B to always send $m = 1$ and only perfect experts attain reputation 1.

at every period, one from company X or Y. She knows that the physician wants her to pick the medicine from company X because he is known to have shares (say). But only the physician knows whether X's or Y's medicine is best for her today. Does it help, then, to ask another physician who has shares in company Y?

3.7 Proofs

To ease the notations and make the expressions shorter, I write q_i instead of $q = i$ for all $q \in \{\omega, s, m, a\}$ and $i = 0, 1$. I also write "$A(r, r')$ plays..." instead of "A plays... under reputations (r, r')". The same goes with primes.

Proof. (Of Lemma 3.1.)

A noisy expert sends a message at random. Listening to and following the recommendation of such an expert gives in average $1/2$ per period. Due to the assumption of informativeness, a good expert gives at least as good advice as a noisy expert so listening to and following the recommendation of a good expert yields a current payment of $v_A(r) \geq 1/2$ (for all $r \in [0, 1]$). Besides, μ-experts are perfectly sincere by assumption and therefore yield higher expected payoffs than any good expert with lower reputation (no matter what the strategy of the latter is). Consequently it is optimal to fire an expert if his reputation is below μ. □

Before we continue the proofs, it is necessary to introduce notations and to investigate (i) how the reputation updates work, (ii) what indifference equations expert A must fulfil. Suppose that $\mu < 1, p \in]1/2, 1]$ and assume that there are only E experts (the same will hold whenever there are only E' experts). Recall that $x(r) = \mathbb{P}(m_1|s_0, r)$ is the probability to lie for $A(r)$ when the signal is low.

Proposition-Definition 3.1. Given r, there are *four* possible updates:

- True positive: $f_1(r) = \mathbb{P}(E = A|m_1, \omega_1, r) = \dfrac{r(p + (1-p)x(r))}{r(p + (1-p)x(r)) + (1-r)/2}$
- True negative: $f_2(r) = \mathbb{P}(E = A|m_0, \omega_0, r) = \dfrac{rp(1-x(r))}{rp(1-x(r)) + (1-r)/2}$
- Sincere error: $f_3(r) = \mathbb{P}(E = A|m_0, \omega_1, r) = \dfrac{r(1-p)(1-x(r))}{r(1-p)(1-x(r)) + (1-r)/2}$

- Opportunistic error: $f_4(r) = \mathbb{P}(E = A|m_1, \omega_0, r) = \dfrac{r(1-p+px(r))}{r(1-p+px(r)) + (1-r)/2}$

For all $r \neq 0, 1$, $f_1(r), f_2(r) > r > f_3(r), f_4(r)$.

Proof. I shall write the proof only for f_1, the other cases are similar. First, by Bayes' rule, we have $\mathbb{P}(m_1, \omega_1|E = A, r) = \mathbb{P}(\omega_1) \times \mathbb{P}(m_1|\omega_1, E = A, r) = \dfrac{1}{2}(p + (1-p)x(r))$. Likewise, $\mathbb{P}(m_1, \omega_1|E = B, r) = \dfrac{1}{2} \times \dfrac{1}{2}$.

Therefore, by the law of total probabilities and Bayes' rule,

$$\begin{aligned} f_1(r) = \mathbb{P}(E = A|m_1, \omega_1, r) &= \frac{\mathbb{P}(m_1, \omega_1|E = A, r)}{\mathbb{P}(m_1, \omega_1|E = A, r) + \mathbb{P}(m_1, \omega_1|E = B, r)} \\ &= \frac{r(p + (1-p)x(r))}{r(p + (1-p)x(r)) + (1-r)/2} \end{aligned}$$

□

$A(r)$, conditional on the signal he receives, weighs the prospects of lying and telling the truth. These prospects depend on the current payment and future ones. Conditional on s_0, $A(r)$ prefers (or is indifferent between) to send m_0 if and only if

$$0 + \delta p V_A(f_2(r)) + \delta(1-p)V_A(f_3(r)) \geq 1 + \delta(1-p)V_A(f_1(r)) + \delta p \times y(r) \times V_A(f_4(r)) \tag{3.2}$$

Conditional on s_1, $A(r)$ prefers to send m_1 if and only if

$$1 + \delta p V_A(f_1(r)) + \delta(1-p)y(r)V_A(f_4(r)) \geq \delta(1-p)V_A(f_2(r)) + \delta p V_A(f_3(r)) \tag{3.3}$$

Proof. (Of Prop. 3.1.)

Recall that A must verify Equations 3.2 and 3.3 when he sends a message (or the converse inequality, or equality when there is indifference). Fix some $r \geq \mu$. Write V_i instead of $V_A(f_i(r))$.

First step: $\{m_i, s_{-i}\}$ is impossible. Assume first that in r, A is indifferent between being truthful and lying when s_0. Consequently there is equality in Equation 3.2. Therefore

$$y^* = \frac{pV_2 + (1-p)V_3 - (1-p)V_1}{pV_4} - \frac{1}{\delta p V_4} \tag{3.4}$$

Inject y^* in Equation 3.3. The left term writes

$$\begin{aligned} &1 + \delta p V_1 + \delta(1-p)V_4 \frac{1}{pV_4}\Big(pV_2 + (1-p)V_3 - (1-p)V_1 - 1/\delta\Big) \\ &= \frac{2p-1}{p} + \delta\frac{2p-1}{p}V_1 + \delta(1-p)V_2 + \delta\frac{(1-p)^2}{p}V_3 \end{aligned} \tag{3.5}$$

After some simplifications, A prefers to send message $m = 1$ conditional on $s = 1$ if, and only if, $1 + \delta V_1 \geq \delta V_3$, which is true. It means that if A is indifferent between the two messages when $s = 0$, then he prefers to send $m = 1$ when $s = 1$. It is easy to check that this is all the more true when A strictly prefers to send message $m = 1$ conditional on $s = 0$.

To sum up

$$\text{If } \Big(s = 0 \implies m = 1 \text{ or } A \text{ indifferent when } s = 0\Big), \text{ then } \Big(s = 1 \implies m = 1\Big) \tag{3.6}$$

or by contraposition

$$\text{If } \Big(s = 1 \text{ and } m = 0\Big), \text{ then } \Big(s = 0 \implies m = 0\Big) \tag{3.7}$$

Step 2: m_0 whatever the signal is Pareto dominated. Suppose that there exists r such that $A(r)$ sends message m_0 whatever the signal in equilibrium. Under this strategy, only B can send message m_1 in r and $f_1(r) = f_4(r) = 0, f_2(r) = f_3(r) = f(r)$. The payoff of $A(r)$ is therefore $\delta(1-p)V_2 + \delta p V_3 = V_A(r)$ (which, given the preference, is larger than 1). The payoff of D is $V_D(r) = 1/2 + V_D(f(r))$.

Let us define the following strategy: A keeps his strategy fixed except in r where he sends message m_1 regardless of the signal. Under this strategy, only B can send message m_0 in r and $f_2(r) = f_3(r) = 0, f_1(r) = f_4(r) = g(r)$. It is straightforward to see by symmetry that $f = g$.

The payoff of A is therefore

$$V'_A(r) = 1 + \delta p V_1 + \delta(1-p)V_4 \times y(g(r)) \tag{3.8}$$

and the payoff of D is

$$V'_D(r) = 1/2 + \big(\delta p + (1-p)y(g(r))\big)V_D(g(r)) + \delta(1-p)\big(1-y(g(r))\big)V_D(\mu) \tag{3.9}$$

Set $y(g(r)) = 1$, this yields $V'_A(r) > V_A(r)$ and $V_D(r) = V'_D(r)$: this is a Pareto improvement. This adapted strategy forms an equilibrium because the inversion yields reverted belief updates– $f = g$–and so the comparison is static: the current payments at r form a Pareto improvement; this does not change what happens in $r' \neq r$. □

Proof. (Of Prop. 3.2.)

The firing/keeping strategy of D is contingent on whether the message fitted the state of the world; not on whether the chosen action fitted the state of the world. Therefore, the choice of the action at the current period does not impact the continuation value of D. Consequently, it is optimal for D to choose the action that maximises the current expected payment, namely a_1 (reps. a_0) whenever $\mathbb{P}(\omega_1|m, m', r, r') > 1/2$ (resp. the probability that ω_0 given...).

If $\mathbb{P}(\omega_1|m, m', r, r') = 1/2$, then D is indifferent between the two actions. It can happen in the competing case: if E and E' have the same reputation, that A and A' have symmetric strategies, and that $m \neq m'$, then D does not have valuable information. Out of symmetry, it is assumed that D breaks ties evenly. □

The benefice of having competition. In the following paragraph, I prove the main result of the article: the decision maker benefits from having ideological competition only when $\mu < 1,\ p < 1$.

Let us prove first that competition is useless in the case $\mu = 1$. Proposition 3.3 states that D can extract all the information when there are only A experts.

Proof. (Of Prop. 3.3.) I show that each strategy best responds to the other. It is clear that D cannot expect a payment greater than p in expectancy per period. It is therefore sufficient to show that y is so chosen that it is a best response, for A, to truthfully report his type.

Step 1: conditional on s_1 it is a best response for A to tell the truth. If A tells the truth, his expected payoff is $1+\delta\big(1-(1-p)(1-y)\big)V_A$ while if he lies, it is $0+\delta V_A$. One therefore needs to have $V_A \leq \dfrac{1}{\delta(1-p)(1-y)}$. Yet one can show that $V_A = \dfrac{1}{2-\delta(1+p+y(1-p))}$.

Indeed,

$$\begin{aligned} V_A &= \mathbb{P}(s_0)V_A(s_0) + \mathbb{P}(s_1)V_A(s_1) \\ &= \frac{1}{2}\Big(0+\delta V_A + 1 + \delta(pV_A + (1-p)yV_A)\Big) \\ &= 1/2 + \delta\big(1+p+(1-p)y\big)V_A/2 \\ &\iff V_A\big(1-\delta(1+p+(1-p)y)/2\big) = 1/2 \\ &\iff V_A = \frac{1}{2-\delta\big(1+p+(1-p)y\big)} \end{aligned} \tag{3.10}$$

The latter is smaller than the former if, and only if, $\delta \leq 1$.

Step 2: conditional on s_0, A best responds by telling the truth. If A tells the truth, his expected payoff is $0+\delta V_A$ while if he lies, it is $1+\delta(1-p+py)V_A$. One therefore needs to have $V_A \geq \dfrac{1}{\delta p(1-y)}$. Yet,

$$\begin{aligned} V_A \geq \frac{1}{\delta p(1-y)} &\iff \frac{1}{2-\delta\big(1+p+(1-p)y\big)} \geq \frac{1}{\delta p(1-y)} \\ &\iff \delta p(1-y) \geq 2-\delta\big(1+p+(1-p)y\big) \\ &\iff \delta(2p+1) - y\delta(2p-1) \geq 2 \\ &\iff y\delta(2p-1) \leq \delta(2p+1) - 2 \iff y \leq \frac{\delta(2p+1)-2}{\delta(2p-1)} \end{aligned} \tag{3.11}$$

The condition $\delta \geq \dfrac{2}{2p+1}$ ensures that $\dfrac{2-\delta(2p+1)}{\delta(2p-1)} \geq 0$. Computing the expected payment of A for the extreme values of y is straightforward. □

Proof. (Of Corollary 3.1.) The same proof works *mutatis mutandis*. The current payment for each expert can be determined by symmetry: A gets a payment of 1 at the current period if, and only if, A' gets 0. By symmetry, it must then be that the expected current payment is 1/2. □

Let us now prove that competition is immaterial–strategy-wise–when good experts are perfect ($p = 1$). As $p = 1$, a good expert directly observes ω and he conditions his message directly thereupon (I skip the notations with s).

Proof. (Of Prop. 3.4.) Under this strategy profile and because $p = 1$ mistakes immediately unveil noisy experts (be errors sincere or opportunistic): $f_3(r) = f_4(r) = 0$ for all $r \geq \mu$. Write $f = f_1 = f_2$.

Step 1: best response for A.

Conditionally on ω_1, the expected payment of A if he sends m_1 is $1 + \delta V_A(\frac{2r}{1+r})$; it is 0 if he sends m_0. It is optimal to reveal a high state of the world.

Conditionally on ω_0, the expected payment of A if he sends m_0 is $\delta V_A(f(r))$; it is 1 if he lies and sends m_0.

It is easy to see that, should he abide to the strategy profile,[10] a good expert will never be fired and will earn 1/2 per period in expectation: $V_A(r) = \dfrac{1/2}{1-\delta}$ for all $r \geq \mu$. For A to be sincere upon $\omega = 0$, it needs be $\delta \times \dfrac{1/2}{1-\delta} \geq 1$, namely $\delta \geq 2/3$.

Step 2: best response for D.

The value function $r \to V_D(r)$ is strictly increasing under this strategy profile. Indeed, $V_D(r) = \dfrac{1+r}{2} + \delta\times$ future payment. The future payment of D should she fire an expert is $V_D(\mu)$. Consequently, it is optimal to fire experts whose reputation is below μ (in any history) and to keep the others.

Step 3: computation of $V_D(r)$.

Fix $r \geq \mu$. I shall reason *ex ante*. Write $V_1^*(r)$ the expected payment of D if $E = A, V_2^*(r)$ the expected payment of D if $E = B$. By construction, $V_D(r) = rV_1^*(r) + (1-r)V_2^*(r)$ and $V_1^*(r) = \dfrac{1}{1-\delta}$.

[10] This rests on Assumption 3.1, in particular on the assumption that an expert is never fired after a fitting message.

By construction, in the event of $E = B, D$ earns $1/2$ per period in expectation; with probability $1/2, B$ sends the fitting message and survives, with probability $1/2$ he is wrong and fired. It writes

$$V_2^*(r) = 1/2 + \Big(\delta V_2^*(f(r)) + \delta V_D(\mu)\Big)/2 \tag{3.12}$$

As we reason conditionally on $E = B, V_2^*$ is a constant with respect to the reputation r. Consequently $V_2^*(r) = \dfrac{1 + \delta V_D(\mu)}{2 - \delta}$ and

$$V_D(r) = r \times \frac{1}{1-\delta} + (1-r) \times \frac{1 + \delta V_D(\mu)}{2 - \delta} \tag{3.13}$$

For $r = \mu$, it finally writes

$$V_D(\mu) = \frac{\dfrac{\mu}{1-\delta} + \dfrac{1-\mu}{2-\delta}}{1 - (1-\mu)\dfrac{\delta}{2-\delta}} \tag{3.14}$$

To find the expression of $V_D(r)$, plug $V_D(\mu)$ in Equation 3.13. (Notice that $V_D(r)$ is indeed increasing in r.) □

The same ideas are at play when there are both experts E and E'.

Proof. (Of Prop. 3.5.)

The strategies of the experts are the same be there competition or not. The reputation updates and the payments of the good experts are the same. The strategy of D is likewise optimal: never fire an expert unless his reputation falls below μ.

The payment of D, however, is larger than in Proposition 3.4. She indeed samples experts until she faces a good one. When there is competition, she samples two experts at a time. Consequently, the length of time before she faces a good expert is, in expectation, lower[11] and (accordingly) her expected payment is larger. (To compute her payment is very cumbersome.) □

Finally, Theorem 3.1 states that the decision maker benefits from competition when $\mu < 1$ and $p < 1$. The proof requires an intermediary result. Lemma 3.2 states that in equilibrium, good

[11] For instance, the probability that she faces a good expert from he start is $1 - (1-\mu)^2 = \mu(2-\mu) > \mu$ when $\mu < 1$.

experts cannot be perfectly sincere when their reputations get too high and without competition. It is a straightforward adaptation of Schottmüller (2019, Theorem 1).

Lemma 3.2. *A good expert cannot be perfectly sincere at all reputations in equilibrium and he is fired with positive probability when he is wrong.*

Notice that Lemma 3.2 proves in particular that Proposition 3.4 cannot hold for $p < 1$. I shall keep the proof informal as it can be retrieved from the original article.

Proof. Suppose that A is not fired when his reputation is sufficiently large. Because $p < 1$ if r is sufficiently close to 1, then the belief update becomes arbitrarily slow. For several periods, A is therefore sure not to be fired regardless of the message he sends. Let T be the number of such periods. His optimal strategy is then to send message $m = 1$ and to earn 1 in the next T periods. Because there is discounting, this payment $1 + \delta + ... + \delta^T$ is higher than his continuation value $V_A(r)$. This shows that A cannot be always sincere.

It follows that he must be fired upon opportunistic error with positive probability. □

Proof. (Of Theorem 3.1.) The strategy of proof goes as follows. Lemma 3.2 shows that sincerity is not sustainable when there are only E experts. I prove here that so long as A faces an expert E' of equal reputation, he may be sincere in equilibrium. I then prove that if $A(r)$ faces a babbling $E'(1/2)$, he must be *more* sincere than in the case without competition should he want not to be fired. This gives two channels through which the expected payment of D is higher under competition.

Consider the following strategy profile:

1. If $r < \mu$ (resp. $r' < \mu$), then A is fired (resp. A');
2. $r = r' \implies A$ and A' are perfectly sincere;
3. $r' < r \implies A'$ babbles and D fires E';
4. If $r' < r < r_{\min}$, then D fires E;
5. If $r > r'$ and $r \geq r_{\min}$, then A and D play $(\hat{x}, \hat{y})$.

The question is: Can we find a suitable $(\hat{x}, \hat{y})$ that makes the profile an equilibrium (in quasi-markovian strategies)? And is it better than (x, y) for D? The answers are twice positive. The cut-off reputation $r_{\min}$ below which a "single" expert is immediately fired will be determined endogenously by the equilibrium payoff of D.

I shall prove that each bullet consists in a best response. Some argument will be provisory and will only be conclusive when we reach the final bullet.

Step 1. I shall first prove that when $r = r'$, it is a best response for each expert. Focus on A (who prefers ω_1). As the experts are sincere under this strategy profile, it entails that $r_1 := f_1(r, r) = \dfrac{rp}{rp + (1-r)/2} = f_2(r, r) =: r_2$. Likewise $r_3 := f_3(r, r) = \dfrac{r(1-p)}{r(1-p) + (1-r)/2} = f_4(r, r) =: r_4$.

Conditional on s_1, it is clear that A prefers to send m_1. Suppose then that s_0. His beliefs about the value of ω and s' do not depend on the message he sends.

In detail,

- $\mathbb{P}(\omega_0, s'_0|s_0) = p^2$; $\mathbb{P}(\omega_0, s'_1|s_0) = p(1-p)$;
- $\mathbb{P}(\omega_1, s'_0|s_0) = (1-p)^2$; $\mathbb{P}(\omega_1, s'_1|s_0) = p(1-p)$.

Indeed,

$$\begin{aligned} \mathbb{P}(\omega_0, s'_0|s_0) = \mathbb{P}(\omega_0, s'_0, s_0)/\mathbb{P}(s_0) &= \frac{\mathbb{P}(\omega_0)\mathbb{P}(s_0, s'_0|\omega_0)}{\mathbb{P}(s_0)} \\ &= \mathbb{P}(s_0|\omega_0) \times \mathbb{P}(s'_0|\omega_0) = p^2 \end{aligned} \tag{3.15}$$

The second-to-last equality comes from the fact that (i) $\mathbb{P}(\omega_0) = \mathbb{P}(s_0) = 1/2$ and (ii) conditional on the state of the world, the signals of the two experts are independent. The three remaining equalities are alike.

Suppose that A sends message m_0 (conditional on s_0). Recall that A' is sincere in this profile when $r = r'$ and sends $m' = s'$.

- If the event $\{\omega_0, s'_0\}$ is realised, then the payment of A is $0 + \delta V_A(r_2, r_2)$;
- If $\{\omega_0, s'_1\}$, then $1/2 + \delta V_A(r_2, \mu)$;
- If $\{\omega_1, s'_1\}$, then $1/2 + \delta \times 0$;

- If $\{\omega_1, s_0'\}$, then $\delta V_A(r_3, r_3)$.

All in all, the expected payment of A should he be truthful at (r, r) when s_0 is

$$\begin{aligned} &p^2\delta V_A(r_2, r_2) + p(1-p)\Big(1/2 + \delta V_A(r_2, \mu)\Big) + p(1-p) \times 1/2 + (1-p)^2\delta V_A(r_3, r_3) \\ &= p(1-p) + \delta \times \Big(p^2 V_A(r_2, r_2) + p(1-p)V_A(r_2, \mu) + (1-p)^2 V_A(r_3, r_3)\Big) \end{aligned} \tag{3.16}$$

If A decides to lie, the conditional probabilities are the same and the expected payment of sending m_1 is

$$\begin{aligned} &p^2 \times 1/2 + p(1-p)\Big(1 + \delta V_A(r_4, r_4)\Big) + p(1-p)\Big(1 + \delta V_A(r_1, r_1)\Big) + (1-p)^2\Big(1/2 + \delta V_A(r_1, \mu)\Big) \\ &= \Big(p^2/2 + 2p(1-p) + (1-p)^2/2\Big) + \delta\Big(p(1-p)V_A(r_4, r_4) + p(1-p)V_A(r_1, r_1) + (1-p)^2 V_A(r_1, \mu)\Big) \end{aligned} \tag{3.17}$$

Recall that $r_1 = r_2, r_3 = r_4$. Consequently, careful simplifications boil down to: A prefers to be sincere when s_0 if, and only if,

$$\delta \times (2p-1) \times \Big(pV_A(r_2, r_2) - (1-p)V_A(r_4, r_4) + (1-p)V_A(r_2, \mu) \Big) \geq \frac{1}{2} \tag{3.18}$$

Besides, it needs be that, for all $(r, r'), V_A(r, r') \geq \dfrac{1/2}{2 - \delta(1+p)}$. Indeed, this corresponds to the strategy (for A) to always be truthful and when any opportunistic mistake entails firing (worst case scenario). Provided that the value of A is non-decreasing in his reputation, Equation 3.18 is satisfied whenever $\delta \times (4p^2 - p + 1) \geq 2$.[12]

So long as the competing experts are sincere and provide good advice, it clearly is a best response for the decision maker to keep them.

Step 2. Suppose that $r' < r$. Under the strategy profile, $A'(r', r)$ babbles so it is a best response for D to fire him. Contrariwise, $A'(r', r)$ is fired so it is a best response to babble. (Note in particular that it is subgame perfect.)

Suppose that $r = r', A$ has it right, A' has it wrong. Call $f(r)$ the updated reputation of $A, g(r)$ the updated reputation of A'. By construction $g(r) < f(r)$. By the argument described above, it is a best response for D to fire A' (and for A' to babble).

Step 3. With suitably chosen $(\hat{x}, \hat{y})$, we have an equilibrium. Write $A(r)$ instead of $A(r, \mu)$. In general, A and D cannot play the sub-profile (x, y) when E' has reputation μ and babbles (whichever the type of the expert). Indeed, compared to the one-expert case (x, y), two competing experts of the same reputation are sincere when they have the same reputation, no matter how high the reputation. We know that sincerity cannot happen in equilibrium at some point under (x, y) because of Lemma 3.2. This means that two competing experts are more sincere than one expert. In turn, this means that the outside option of D is higher when the experts compete. Consequently, $A(r)$ must tell the truth (when s_0) more often than under strategy x.

The question is: Can $A(r)$ adopt a strategy $\hat{x}$ that offers sufficiently much to D? Can D in turn chose a strategy $\hat{y}$ such that A is indifferent between the two messages when he lies?

[12] This condition is sufficient. Lower $\delta \times p$s would work.

Call $W_D(r,\mu)$ the expected of D when $A(r)$ is perfectly sincere when he is alone (note that it means that the experts that have a strategic role are always perfectly sincere, then). This "always sincere" strategy cannot be an equilibrium (apply Lemma 3.2 to see this). Yet for any equilibrium payoff function V_D, it must be that $W_D(1,\mu) > V_D(\mu,\mu)$. Also $W_D(r,\mu)$ is increasing in r. If $W_D(\mu,\mu) \geq V_D(\mu,\mu)$, set $r_{\min} = \mu$. Otherwise, there exists by continuity $r_{\min}$ such that $W_D(r_{\min},\mu) = V_D(\mu,\mu)$.

It means that if $r \geq r_{\min}$ and if he decides not to lie too often, $A(r)$ can make sure that D is *at least* indifferent between keeping and firing $A(r)$. If $r < r_{\min}, D$ fires $A(r)$ immediately, on the other hand.[13] Suppose thereafter that $r \geq r_{\min}$.

Conditional on s_0, the indifference equation of A between the two messages is

$$\delta p V_A(r_2) + \delta(1-p)V_A(r_3) = 1 + \delta(1-p)V_A(r_1) + \delta p \times \hat{y}(r_4) \times V_A(r_4) \tag{3.19}$$

It is quite straightforward to see that if p and δ are large enough (namely the probability to be right and the weight of the future) and if $\hat{y}(r_4) = 0$, then the left-hand side of the equation is strictly larger than the right-hand side.

By continuity, A can tailor his strategy $\hat{x}$ such that D is either indifferent between keeping and firing A or prefers to keep him at any period: $V_D(r,\mu) \geq V_D(\mu,\mu)$–then when D is indifferent, she can choose $\hat{y}$ as she likes (to make A indifferent when hie lies conditional on s_0). If D is "too rarely" indifferent, $A(r)$ can lie slightly more such that D will be indifferent and fire him with positive probability.

□

Equilibrium with constant payments for D. The proofs of each bullet will require several intermediary results. The quantity $(p+1/2)/2$ is exactly the expected payment of D for the current period if A is sincere and has reputation $1/2$ (for instance, when he just steps in).

Lemma 3.3. *Suppose $A(r)$ wants to yield an expected payment of $(p+1/2)/2$ to D per period.*

[13] The necessity to define $r_{\min}$ comes from the "sampling effect" we encountered above. If $r = \mu + \varepsilon$, then it may be possible that D prefers to fire both experts when one gets wrong and to start from scratch.

The strategy of A associated to this payoff is $x(r) = 1 - 1/2r$ (in particular, $x(1/2) = 0$).

Proof. The expected payoff at the current period for D is by definition $rv_A(r) + (1-r)/2$ where $v_A(r)$ is what D gets from $A(r)$. By indifference, it needs be $rv_A(r) + (1-r)/2 = (1/2+p)/2$, hence $v_A(r) = 1/2 + \frac{1}{2r}(p - 1/2)$.

By construction,

$$\mathbb{E}[\text{current payment of } D | E = A, r, x] = \frac{1}{2}(p + (1-p)x(r)) + \frac{1}{2}p(1 - x(r))$$

Solving the equation for $x(r)$ yields $x(r) = 1 - 1/2r$. □

Thereafter, we have $x(r) = 1 - 1/2r$. The reputation updates of Proposition-Definition 3.1 write (second term when $p = 1$)

- $f_1(r) = \dfrac{2r - (1-p)}{r+p}; = \dfrac{2r}{1+r};$
- $f_2(r) = \dfrac{p}{p+1-r}; = \dfrac{1}{2-r};$
- $f_3(r) = \dfrac{1-p}{2-p-r}; = 0;$
- $f_4(r) = \dfrac{2r-p}{r+(1-p)}; = \dfrac{2r-1}{r}.$

The proof of Point 1 of Theorem 3.2. In particular, when $p = 1$, $f_2 \circ f_4 = f_4 \circ f_2 = Id$. For all integer $n \geq 1$, set $r_n = 1 - 1/n$.

One can check that for all $n \geq 2$, $f_4(r_{n+1}) = r_n$, $f_2(r_n) = r_{n+1}$ and $f_1(r_n) = r_{2n-1}$. This rests on the assumption that $\mu = 1/2$. Then whatever the history of the game, the set of reputations an expert can have is exactly $\{r_n : n \geq 1\}$. (The proof is a simple induction. In particular, the reputation $r = 1$ is never reached.)

To simplify the notations, write $y_n = y(r_n), x_n = x(r_n), V_n = V_A(r_n)$. Look at the indifference equations A must satisfy. Equation 3.3 is void (automatically satisfied). Equation 3.2 now writes

$$\delta V_{n+1} = 1 + \delta y_n V_{n-1} \tag{3.20}$$

In addition, one can compute V_{n+1} unconditionally:

Lemma 3.4. *It writes*

$$V_{n+1} = \frac{1}{2} \times \frac{3n-4}{2n-2} + \frac{1}{2}\delta\Big(V_{2n+1} + \frac{n}{2n-2}V_{n+2} + y_{n+1}\frac{n-2}{2n-2}V_n\Big) \tag{3.21}$$

Proof. This result is *ex ante*, namely before expert A learns the state of the world of the next period. Assume that the reputation of A is r_{n+1}. There are two possibilities, depending on the state of the world. If ω_1, then the payment of A is 1 and his reputation moves from r_{n+1} to r_{2n+1}. This happens with probability 1/2. If ω_0, the payment of A is (in expectation) x_n (he lies) and with probability x_n his reputation lowers to r_n and he is fired with probability y_{n+1}, with probability $(1-x_n)$ his reputation reaches r_{n+2}. This happens with probability 1/2.

Putting everything together, we have that

$$V_{n+1} = \frac{1}{2}(1+x_n) + \frac{1}{2}\delta\Big(V_{2n+1} + (1-x_n)V_{n+2} + x_n y_{n+1}V_n\Big)$$

One has $x_n = \dfrac{n-2}{2n-2}, 1-x_n = \dfrac{n}{2n-2}$. Hence the desired result. □

One can get rid of the y_n, it is a bit puzzling :

Corollary 3.2. *One has*

$$V_{n+1} = \frac{1}{2} + \frac{1}{2}\delta\Big(V_{2n+1} + V_{n+2}\Big) \tag{3.22}$$

Proof. It suffices to combine Equations 3.20 and 3.21. Recall that, by indifference, $\delta V_{n+2} = 1 + \delta y_{n+1}V_n$. Equivalently $y_{n+1}V_n = V_{n+2} - 1/\delta$.

Therefore Equation 3.21 writes

$$\begin{aligned} V_{n+1} &= \frac{1}{2} \times \frac{3n-4}{2n-2} + \frac{1}{2}\delta\Big(V_{2n+1} + \frac{n}{2n-2}V_{n+2} + \frac{n-2}{2n-2}(V_{n+2} - 1/\delta)\Big) \\ &= \frac{1}{2} \times \frac{3n-4}{2n-2} - \frac{1}{2} \times \frac{n-2}{2n-2} + \frac{1}{2}\delta\Big(V_{2n+1} + V_{n+2}\Big) \end{aligned}$$

The observation that $\dfrac{1}{2} \times \dfrac{3n-4}{2n-2} - \dfrac{1}{2} \times \dfrac{n-2}{2n-2} = 1/2$ leads to the desired result. □

Finally, one can compute $V_A(1)$ and $y(1)$.

Lemma 3.5. *One has* $V_\infty = V_A(1) = \dfrac{1/2}{1-\delta}$ *and* $y_\infty = y(1) = 3 - 2/\delta$ *(which makes sense only insofar as* $\delta \geq 2/3$*).*

Proof. Notice that when $r = 1$, no update of the reputation takes place any more. $x(1) = 1/2$ so $V_\infty = 3/4 + \delta(1 - \mathbb{P}(\text{being fired at a current period}))V_\infty$. This *ex ante* probability of being fired is precisely $\mathbb{P}(\omega_0) \times \mathbb{P}(m_1|\omega_0, r = 1) \times (1 - y_\infty) = (1 - y_\infty)/4$.

Consequently $V_\infty = 3/4 + \dfrac{\delta}{4}(3 + y_\infty)V_\infty$. This yields $V_\infty = \dfrac{3}{4 - \delta(3 + y_\infty)}$.

Besides, Equation 3.20 writes $\delta V_\infty = 1 + \delta y_\infty V_\infty$. Solving this for y_∞ yields $y_\infty = 3 - 2/\delta$. Finally, plugging this into $V_\infty = \dfrac{3}{4 - \delta(3 + y_\infty)}$ entails $V_\infty = \dfrac{1/2}{1-\delta}$. □

Now we can prove the first statement of Theorem 3.2.

Proof. (Of Point 1 of Theorem 3.2.) The previous intermediary results unfold. The expected payment of the decision maker at every period and for every reputation r is $v_D(r) = rv_A(r) + (1-r)/2 = r(1/2 + \dfrac{1}{2r}(1-1/2) + (1-r)/2 = 3/4$. Therefore, at any period such that $r \geq 1/2$ the decision maker is indifferent between firing and keeping the expert in place; this holds in particular when the message is m_1 while the state of the world is ω_0. This means that $0 < y(r) < 1$ is a best response for D. If $r < 1/2$ it is a strict best response for D to fire the expert.

The quantity $3 - 2/\delta$ exactly solves Equation 3.20 and is associated with the continuation value $V_A = \dfrac{1/2}{1-\delta}$. The couple $(y(r), V_A(r)) = (3 - 2/\delta, \dfrac{1/2}{1-\delta})$ solves the Bellman equation. It means that A is indeed indifferent, upon observing ω_1 between the two messages he can send. He can therefore use a mixed strategy $x(r)$.

The fact that at any stage $r \in \{r_n : n \geq 1\}$ was proved above. □

The proof of Point 2 of Theorem 3.2.

Proof. It is fairly mechanistic. The idea is that if r is sufficiently close to $1/2$, then $f_4(r) < 1/2$ and $A(r)$ does not want to lie in r because he would be fired *for sure* upon opportunistic error. Such "sufficiently close to $1/2$" reputations happen with positive probability (whereas when $p = 1$, the second-lowest reputation a perfect expert can have under any history is $2/3$ and $f_4(2/3) = 1/2$).

One can check that $f_1(1/2) = f_2(1/2) = \dfrac{2p}{2p+1}$. Then $f_4(2p/(2p+1)) = \dfrac{3p-2p^2}{3p-2p^2+1} > \dfrac{1}{2}$ but $f_4\Big(\dfrac{3p-2p^2}{3p-2p^2+1}\Big) < 1/2$.

It means that, if $r = \dfrac{3p-2p^2}{3p-2p^2+1}$, then $V_A(f_4(r)) = 0$ (so is $V_A(f_3(r)) = 0$) and Equation 3.2 now writes

$$\delta p V_A(f_2(r)) = 1 + \delta(1-p)V_A(f_1(r)) \tag{3.23}$$

As $V_A(t) \geq \dfrac{1}{2-\delta(1+p)}$ for all t and V_A is non-decreasing in the reputation, Equation 3.23 can be written as $\delta(2p-1) \times \dfrac{1}{2-\delta(1+p)} > 1$, which is true whenever $\delta \times p > 2/3$. Consequently, A cannot be indifferent in $r = \dfrac{3p-2p^2}{3p-2p^2+1}$ (namely he prefers to be truthful) and cannot play mixed.[14] Conversely, A is sincere and D gets in expectation $rp + (1-r)/2 > \dfrac{p+1/2}{2}$: this is a contradiction. □

3.8 An alternative model: Multiple states of the world

This section describes an alternative model that I first considered to tackle the political question. It is not part of the article. Contrary to the model as defined in Section 3.2 where there is one common state of the world $\omega \in \{0,1\}$, it is conceivable that several state of the world coexist, one per expert. The rationale behind is that each expert "comes" with a policy which may be good or not to the decision maker, each independent with the others. (The model I consider in the body can be reinterpreted as two states of the world that are perfectly negatively correlated. Hence two experts come each with a policy at every period; but the two policies are perfectly incompatible instead of independent.)

I believe that the framework "multiple states of the world" is not as realistic as the framework "common state of the world." Indeed, left- and right-wing policies are supposed to be incompatible quite often. Regardless, I claim it is interesting for two reasons. First, it casts exactly the opposite

[14] Of course, if δ or p is too low–that is if experts value the future too little, or are not good enough–, then A favours his short term interest and always sends m_1.

result: the more competition there is (as defined below), the easier it is to lie to the decision maker. This sheds light on how the rationale that competition helps the decision maker is model-dependent. Besides, the "multiple states of the world" framework can be applied to sellers/buyer relationships–e.g. several sellers compete with each their own good, if the buyer opts for the good of the first company, she cannot know (counterfactually) whether it would have been a good idea to choose the good of the second company instead.

In the following, I describe how the "multiple states of the world" model differs to the model of Section 3.2, then I present the main results. Finally I discuss the weakness of this toy model.

The "multiple states of the world" model. There are n senders $N = \{1, ..., n\}$ and a decision maker D. It is again a repeated game of cheap talk. At each period, a state of the world $\omega = (\omega_j)_{j \in N} \in \{0, 1\}^n$ is drawn. The ω_j are independent and $\mathbb{P}(\omega_j = 1) = q_j$. Each expert $j \in N$ receives a signal s_j conditional on ω_j. His signal is of quality $p_j \geq 1/2$. Given his signal s_j, expert j sends a (cheap talk) message m_j.

The set of actions of D is $A = \{a_1, ..., a_n, a^*\}$ at every period with the interpretation that choosing action a_j boils down to opting for j's policy. D earns 1 if the underlying state of the world is high, that is if $a = a_j$ and $\omega_j = 1$, and 0 otherwise. The action a^* is the "outside option" of D that she chooses only if she believes that all the ω_j are low.[15] Expert j earns 1 if, and only if, his policy is chosen. This entails that at each period at most one expert receives 1 while the others receive 0. Quite importantly, if D chooses a_j, she cannot know whether expert $i \neq j$ was sincere (this is due to independence).

The decision maker is allowed to "punish" an expert by not choosing his preferred action for as many periods as she likes. Firing an experts amounts to punishing forever after. Importantly, there is no replacement by a peer expert of the same ideological bias.

The capacities q_j are known, so there is no notion of hidden competence, hence no reputation either.

[15] This could not happen above as states of the world were perfectly negatively correlated.

Heuristics of the results. Let us assume to simplify that for all $j \in N, q_j = q \geq 1/2, p_j = 1$: each policy is good (for D) with probability q, experts have perfect signals. Whenever $s_j = 1$, j knows that $\omega_j = 1$, namely that the state of the world is favourable to him. It is then optimal to disclose such information to the decision maker (under hypothesis that an expert is never punished when he tells the truth). The incentive to lie comes when an expert receives a bad signal and knows that "his" state of the world is low. The trade-off is the following when an expert faces competing experts that sincerely report their signals. On the one hand, the probability that a lie pays off in the short run (that is, his policy is chosen today against concurrent good policies) is lower than should he stand alone: deviations pay less in the short run. On the other hand, a lie is less easily detected (as, again, his policy may not be chosen in spite of the deviation): deviations are less "dangerous".

The latter effect takes precedence over the former. In other words, experts prefer low "profitability" of deviations and rare punishments to high profitability and frequent punishments. It means that when there are several competing experts, there are more incentives to deviate and lie. This makes credible communication more difficult to sustain up to the asymptotic case on n where δ must be larger than $1/(1+n)$ for credible communication to happen.

These results depend crucially on two hypotheses. First, to each expert their independent state of the world and, henceforth, their proposed policy. Second, when an expert's policy is not chosen by the decision maker, then it is not possible to know whether this expert was sincere or not–hence lies may not be detected. All in all, the results depicted here represent converse statements to Section 3.4: ideological competition makes the decision maker *worse off*.

Formal results. I focus on the following strategies for the decision maker: when an expert lied, he can be punished for one period (that is, next period his policy cannot be chosen, it resumes thereafter); for two periods; forever after (that is, as if he were fired). I describe the best responses of the experts, I do not consider whether her punishing strategy actually consists in a best response for the decision maker.

Due to discounting ($\delta < 1$), one euro is better today than tomorrow. If the punishment runs

over one period only, then experts always want to lie upon a bad state of the world, no matter how many they are. There cannot therefore be credible communiaction through cheap talk.

Proposition 3.6 treats the two-periods case. It shows that sincerity is a best response (for sufficiently high δ) when there is a single expert and with two competing experts (but with a more stringent condition on δ). It also shows that for large n, telling the truth is not a best response to others being sincere.

Proposition 3.6. *(Punishments for two periods.) Suppose that the decision maker punishes an expert over two periods when he lied.*

- *One expert case: it is a best response for the single expert to (always) tell the truth if and only if* $\delta \geq -\frac{1}{2} + \frac{\sqrt{q^2+4q}}{2q}$;[16]
- *Two experts case: against a sincere competing expert, it is a best response to tell the truth if and only if* $\delta \geq -\frac{1}{2} + \frac{\sqrt{(f(q)^2+4f(q)}}{2f(q)}$ *where* $f(q) = q \times (1-q/2)$ *(this condition is more stringent);*
- *Asymptotic case: it is never a best response to tell the truth against sincere experts (it needs be that asymptotically* $\delta \geq \sqrt{n}$*).*

Proposition 3.7 states that, asymptotically, experts should be infitely patient to have incentives to tell the truth. This highlights the idea that when deviations are both less dangerous and less worth it in the short run, the former effect takes precedence. Credible communication cannot be sustained, no matter how severe the punition.

Proposition 3.7. *Suppose that the decision maker punishes an expert forever after a lie. Let* $i \in N$ *and suppose that the experts* $j \neq i$ *are sincere. Then, upon* $\omega_i = 1$ *expert* i *has his policy chosen by* D *with probability* $(1-(1-q)^n)/(nq)$*. Conditional on* $\omega_i = 0$*,* i *prefers to tell the truth if and only if* $\delta \geq 1 - 1/(n+1)$*. (I rule out lower-order terms.)*

Limits of the alternative. This short presentation of an alternative toy model is plagued by two weaknesses. First, I impose that the signals are perfect for everyone and the "qualities"

[16] This makes sense only insofar as $q \geq 1/2$.

of the policies are the same (that is, $q_j = q$). It remains to verify that the increasing incentives to lie as a function of the number of experts holds when there is heterogeneity between the experts' capacities. I suggest that this is innocuous to the main result. Second, I do not study the behaviour of the decision maker. Instead I prescribe her fixed strategies without questioning whether these are best responses for her. Consequently, I am able neither to establish Nash equilibria, nor test the property of subgame perfection. In a sense, only half of the work is done.

There is room for more work. In particular, one can wish to assess the effect of ideological competition when states of the world are partially correlated (instead of either independent, or perfectly negatively correlated). I discuss this research program in the conclusion of this thesis.

Proofs of the propositions.

Proof. (Of Proposition 3.6.) It is quite straightforward. First, experts never lie when their state of the world is high. Second, let us fix a period $t = 0$. As the punishment, if any, runs over two periods, the payment from $t = 3$ onwards of an expert is the same should he deviate or not (only) at $t = 0$. When can therefore study only the incentives to lie at $t = 0$ with respect to $t = 0, 1, 2$.

- *One expert case.* Conditionnal on $\omega_1 = 0$, if the expert sends $m_1 = 1$ then he earns 1 today (for sure) and his deviation is identified (for sure also). He therefore gets $1 + 0 \times \delta + 0 \times \delta^2$ (plus his payment from $t = 3$ onwards). If he is truthful, then he gets 0 today and his "equilibrium payoff" from $t = 1$. By construction, he gets q per period in average. All in all, being truthful outperforms a deviation today if and only if $1 \leq q \times \delta + q \times \delta^2$. It is a polynomial inequation in δ which positive solution is given by the statement of the Proposition.
- *Two experts case.* If both experts are truthful, then the *ex ante* expected payment for today is $q \times (1 - q/2) =: f(q)$ for each. Indeed, take the point of view of the first expert. He gets zero if $\omega_1 = 0$ (with probability $1 - q$). With probability q his state of the world is high. Then with probability $1 - q, \omega_2 = 0$, the second expert sends $m_2 = 0, D$ chooses action a_1 and the first expert earns 1. But with probability $q, \omega_2 = 1, m_2 = 1$ and D breaks ties in favour of the first expert with probability $1/2$. Hence the (expected *ex ante*) current

payment of the first expert is $(1-q)\times 0 + q\times\big((1-q)\times 1 + q\times 1/2\big)$ as claimed.

Now by the same token, conditionnal on $\omega_1 = 0$, expert 1 earns $(1-q/2)$ if he lies and he is detected with probability $(1-q/2)$, in which case he is punished at periods $t = 1, 2$. With complementary probability $q/2 = \mathbb{P}(\omega_2 = 1)\times\mathbb{P}(\text{expert 1 looses tie})$ he is not detected and the at $t = 1, 2$ is as if he had been truthful. All in all, being truthful outperforms a deviation today if and only if $f(q)\times\delta + f(q)\times\delta^2 \geq 1$ which is the same polynomial inequation with $f(q)$ instead of q. As $f(q) < q$, the condition on δ is indeed more stringent.

- *Asymptotic case.* More generally, if the experts are truthful, the current expected *ex ante* payment of each expert is $(1-(1-q)^n)/n \sim_{n\to+\infty} 1/n$: the numerator is the probability that at least one state of the world $\omega_j = 1$ (for some $j \in N$), namely that at least one expert sees his policy chosen. As the experts are symmetrical, one must divide by n to have the payment for each expert. Consequently, conditionnal on $\omega_1 = 1$, the current (expected *ex interim*) payment of expert 1 is $(1-(1-q)^n)/(nq)$. The same arguments unfold and the inequation becomes (using the equivalent) $\delta/n + \delta^2/n \geq 1$. Two-steps punitions do not incentivize experts to be truthful when the number of experts grows arbitrarily large.

□

Proof. (Of Proposition 3.7.) I work with the equivalent $1/n$ of $(1-(1-q)^n)/n$. If all the experts are truthful, then the *ex ante* expected discounted payment of each is $\dfrac{1/n}{1-\delta}$. One needs only to add up the arguments presented in the proof of Proposition 3.6. Indeed, conditional on $\omega_1 = 1$, telling the truth (and sending $m_1 = 0$) is better than deviating if and only if

$$0 + \delta\times\frac{1/n}{1-\delta} \geq \frac{1}{nq} + \Big(1-\frac{1}{nq}\Big)\times\delta\times\frac{1/n}{1-\delta}$$

which is equivalent to $\delta \geq \dfrac{1}{n+1}$. □

4 Strategic Experimentation and Protection

4.1 Introduction

Floods, thunder and fires used to be considered signs of the wrath of God. Sacrifices and processions had very little power against fate. This is no longer the case, we human beings believe that we can master or at least mitigate the effects: this depends on our doing. We know how to react to, say, the possibility of a flood: we build dams. Yet this does not prevent floods to occur (we are not precautious enough) and to be natural disasters indeed. For instance, in July 2021, a flood caused the death of 269 people in Germany, the Netherlands and Belgium. In October 2024, a flood caused the death of 225 people in the region of Valencia (Spain).

Contrary to our ancestors, we no longer speculate on why God(s) has decided to punish us and what our misbehaviour was. Rather, confident in our ability to tame Nature, we ask: Why were the public authorities not able to prevent this disaster to happen?[1] Specifically, the public authorities failed along two parameters: Why could it not forecast the event (deficit of information)? And why could it not resist the event (deficit of protection)?

Being confident that one could have acted some way is different to genuinely acting. In the face of a conceivable disaster, adopting the "proper" course of action is difficult as one faces two

[1]The word "disaster" is here ambiguous. Does it refer to the flood itself? In this case, preventing it would amount to build dams and retention ponds upstream. Does it refer to the death casualties directly? Then, preventing it could go through sheer evacuation. In the article, I focus on the first interpretation.

types of uncertainty. First, one does not know if the disaster may happen in the future–how suitable for a flood is the valley? Second, if a disaster may happen, one does not know when–is the valley in danger tomorrow or in fifty years?

In light of this two-fold uncertainty, one can adopt a precautionary behaviour and build protection today to prevent future casualties. However, if the disaster will not happen, this proves waste *ex post*. On the other hand, one can adopt a wait-and-see attitude. While cheap and possibly good-looking *ex ante*, this may turn out disastrous *ex post*. A middle ground policy tries to lift the uncertainty by commissioning research (about the hydrography and geology of the valley, say).

The framework can be extended to several public authorities. Indeed, Western Germany, the Netherlands and Belgium are equally concerned by the Rhine watershed. They therefore face similar issues. To a large extent, protection is private: the dams Belgium builds benefit her. Conversely, research is public: if Belgium carries on research, the results become common knowledge. The asymmetry between private protection and public knowledge creates incentives to free ride on research expenditures of other parties.

This article aims at answering the following questions: What is the optimal course of action of one public official against the two-fold uncertainty of a natural disaster (decision problem)? And what are the magnitude, if any, of the free riding effect and the encouragement effect when several players strategically interact?

To tackle this problem, I define a game of exponential strategic experimentation. An "event" may or may not happen in the future. If yes (say it is active), its arrival rate is exponentially distributed. Once it arrives, final payments are distributed and the game ends. Each player has three possible actions: protecting, searching and waiting. The mere fact that the event does not occur diminishes the subjective probability that it will occur in the future. Research speeds up the learning process. Protection mitigates the damage caused by the event conditional on its occurence. Protection and research are costly, waiting is costless.

I focus on stationary Markovian strategies where the state variable is the tuple (subjective probability that the event is active, stock of protection accumulated).

The project is at an early stage. I currently struggle with the decision problem, namely the problem of a benevolent social planner. There is no notion of a myopic threshold because both protection and research pay off in the future. An agent that maximises their present payment always prefers to wait, namely irrespective of their belief.

I determine the Bellman equation of the social planner (Equation 4.3) and compute the three utility functions of a player should she commit to one particular action forever after. I am not able yet to solve the boudary conditions (on the belief-stock plane) that the player must satisfy. I present reasons why the smooth pasting condition (that is, the utility function is continuously differentiable at the boundary) might not hold. The strategic problem has not been dealt yet.

Literature review.

This article is related to the litterature on strategic experimentation (Hörner and Skrzypacz (2017) survey the literature). Experimentation aims at capturing the dilemma between exploiting an action of known reward and exploring a risky (but possibly more rewarding) action. Rothschild (1974), Mandelbaum (1987), and Presman (1991) study the single-player problem. The N-players case ($N \geq 2$) starts with Bolton and Harris (1999) for the Brownian model and Keller et al. (2005) for the exponential model.

This article slightly differs to the usual literature in its interpretation of the arms. While in the two-arms literature, one arm is *safe* (providing a known mediocre flow of payment) and one arm is *risky* (providing, if good, high payoffs at random time), the distinction is not as clear cut in this model: building protection is the precautious course of action (in terms of outcomes, not of uncertainty) while waiting can be interpreted as risky; searching endorses no interpretation of the like. Anyhow, a player assumes a flow of deterministic costs, the uncertainty lies in the occurence of the event only. (It is not a stopping game because the actions are not irreversible before the event happens.)

Time is continuous and both the event and research are conclusive at an exponential rate (see e.g. Keller et al. (2005)). The state variable is the tuple (subjective probability of the type, stock of protection) contrary to the literature where only the (common) subjective probability

matters.[2] Also, there is no notion of myopic threshold here because research and protection prove useful only when the future is taken into consideration; it is always better to wait in the short run. Finally, waiting does not freeze the Bayesian update (contrary to the literature), the mere fact that the event does not occur decreases the subjective probability. The evolution of beliefs happens without experimentation (which speeds up the process only).

The model is a "bad news" model: whether the event occurs or research is conclusive, it is bad news–and no news is good news. This makes Keller and Rady (2015) the closest article to mine. Keller and Rady (2015) argue that there is a difference of nature between lump-sum costs and payments. They find that smooth pasting fails at the boundaries between playing safe and experimenting. While it is too early to be affirmative, I present reasons why smooth pasting should fail in my model too.

It is not possible, at this stage, to compare the results I do not have with the literature. Especially, the magnitude of the free riding effect and the presence of the encouragement effect described in Bolton and Harris (1999) cannot be assessed yet. Such results help determine how underoptimal an equilibrium fares compared with the social planner problem.

Observation of both the actions taken and the payoffs received of one player by the other is perfect. It is to differentiate with models where the actions are public but the payoffs are private (Rosenberg et al. (2007)), observation is costly (Kultti and Miettinen (2007) for sequential players, Marlats and Ménager (2021) for long-lived experimenting players), cheap talk communication bypasses private information (Heidhues et al. (2015)), etc.

Finally, at this stage, externalities are purely informational–that is, a player benefits only from their peers searching. However, it would make sense in the framework of a flood to assume that protection is a public benefice for the players downstream, not upstream. Some article study strategic experimentation with externalities that are not purely informational. For instance, Strulovici (2010) assumes that the players face independent one-armed bandits, and the same action is played by everyone: the decision is taken by the majority.

[2]For one probability per player, see however Klein and Rady (2011) where players have imperfect negatively correlated bandits.

4.2 Model

Time $t \in \mathbb{R}_+$ is continuous and is discounted at rate $r > 0$ by two risk-neutral players. An *event* is of two types $\theta \in \{0, 1\}$. If $\theta = 0$, then the event never occurs in the future. If $\theta = 1$, then the event occurs at an exponential rate λ; the prior probability that θ is high is p_0. Hence the probability that the event happens in $[t, t + dt[$ is $\theta \times \lambda \times dt$.

At each period $[t, t + dt[$, the players have one unit of resource to split between protection, research and wait. The resource is a flow. Let[3] k_t denote the investment in protection, q_t the investment in research such that $k_t + q_t \leq 1$. $1 - k_t - q_t$ denotes the budget spared (to wait). Playing (k_t, q_t) costs $(\alpha_k \times k_t + \alpha_q \times q_t)dt$ with $\alpha_k, \alpha_q > 0$: waiting is costless.

Upon playing q_t on $[t, t + dt[$, the probability to learn the type θ is $\theta \times \mu \times q_t \times dt$. Learning can happen conclusively only when $\theta = 1$. It means that no news is good news.[4] The rate of arrival is exponential of parameter μ.

Let $l > 0$ denote the damage caused by the event. Upon the occurence of the event at time $t \in \mathbb{R}_+$, a player receives a lump-sum cost min $\{m_t - l, 0\}$ where $m_t = \int_0^t k_s ds$ stands for the stock of protection built by the player. The minimum materialises the fact that overprotection is waste.

Let Y_t denote the event "no event and no learning of θ until time t." The subjective probability that $\theta = 1$ at time t is therefore $p_t = \mathbb{P}(\theta = 1|Y_t)$. As actions and outcomes are public, the probability p_t is the same for all the players and is common knowledge.

Players use stationary Markovian strategies. The state variable is the tuple (p, m) where $p \in [0, 1]$ is the subjective probability that $\theta = 1$ and $m \in \mathbb{R}_+$ denotes the stock of protection of the player. A *strategy* or *policy* by a player is therefore a mapping $\sigma : [0, 1] \times \mathbb{R}_+ \to \{(k, q) : k+q \leq 1\}$ that prescribes a choice of actions $\sigma(p, m) = (k, q)$ at any (p, m) (with the additional restriction that σ is Borel-measurable on its domain). Note that at any time p_t is the same for the two players;

[3] In principle, one should write $k_t^i, i \in \{1, 2\}$ labels the players. As I will consider only the decision problem, I do not embarass myself with this notational constraint.

[4] Another possibility would be to set the probability at $\mu \times q_t \times dt$, hence no news is genuinely no news. Choosing the latter means that it is possible (through research) to prove that something does not exist (or equivalently, will never happen). Choosing the former assumes that one is able prove only the *existence* of something.

the stocks $m_t = \int_0^t k_s ds$ could in principle differ; they do not if the strategies are symmetric.

It is never optimal to have a stock m exceeding the possible damage l. In the following, it is therefore assumed that σ is defined on $[0,1] \times [0,l]$.

The law of motion. When the event occurs or learning is succesful, the belief jumps to 1. So long as this does not happen, the belief update on $[t, t+dt[$ obeys to the equation

$$p_t + dp_t = \frac{p_t(1-(\lambda+\mu q_t)dt)}{1-p_t+p_t(1-(\lambda+\mu q_t)dt)} \tag{4.1}$$

Indeed, if $\theta = 1$ the instantaneous probability that neither the event nor learning occurs is $1-(\lambda+\mu q_t)dt$; if $\theta = 0$, it is 0. Bayes rule gives then the formula. Having $dt \to 0$ and ommiting terms of order 2 and higher yields the differential equation

$$p_t' = -p_t(1-p_t)(\lambda+\mu q_t)$$

which solution is $p_t = \dfrac{p_0 e^{-\lambda t - \mu \int_0^t q_s ds}}{1-p_0+p_0 e^{-\lambda t - \mu \int_0^t q_s ds}}$.

Parametric assumption and payoff function. When $p=0$, it is optimal to wait (namely to set $k=q=0$). To make the problem non trivial, it must be that when $p=1$ (or sufficiently close to 1), a player prefers to build protection rather than to wait (independent of $m<l$). To that end, it needs be that the cost α_k is small against the time of arrival of the event and that the agent does not discount future too much. Assumption 4.1 ensures this:

Assumption 4.1. $\alpha_k \leq \dfrac{\lambda}{\lambda+r}$

Call u the bivariate payoff function of the player. Because when $p=1$, it is optimal to build protection only until the stock reaches l, $u(1,m)$ writes (in expectancy)

$$\begin{aligned} u(1,m) &= \int_0^{l-m} re^{-rt}e^{-\lambda t}\Big(\lambda\big(t-(l-m)\big)-\alpha_k \times 1\Big)dt \\ &= r\frac{\lambda-\alpha_k(\lambda+r)}{(\lambda+r)^2}\underbrace{\Big(1-e^{-(l-m)(\lambda+r)}\Big)}_{>0 \text{ because } m<l} - \frac{\lambda r}{(\lambda+r)}(l-m) \end{aligned}$$

The payoff function starting at $(p_0, 0)$ and playing $(k_t, q_t)_{t\geq 0}$ is finally

$$\begin{aligned} u(p_0,0) &= \int_0^{+\infty} re^{-rt}\mathbb{P}(Y_t)\Big(\lambda p_t(m_t-l)-\alpha_k k_t-\alpha_q q_t+\mu p_t q_t u(1,m_t)\Big)dt \\ &= \int_0^{+\infty} re^{-rt}\Big(-(\alpha_k k_t+\alpha_q q_t)+p_0 e^{-\lambda t-\mu\int_0^t q_s ds}(\lambda(m_t-l)+\mu q_t u(1,m_t))\Big)dt \end{aligned} \tag{4.2}$$

4.3 The Bellman equation

Let $t \in \mathbb{R}_+$ and suppose that on $[t, t+dt[$, the subjective probability is p and the stock is m. Assume left differentiability for the first variable, right differentiability for the second variable of the payoff function. By the principle of optimality it needs be that

$$u(p,m) = \max_{k+q\leq 1}\Big\{(1-e^{-rdt})\Big(\lambda p(m-l)-\alpha_k k-\alpha_q q\Big)+e^{-rdt}(1-\lambda p dt)\mathbb{E}[u(p+dp,m+dm)|p,q,k]\Big\}$$

Yet $e^{-rdt} = 1 - rdt + o(dt)$ and the terms in $o(dt)$ vanish. By linearity $dm = k \times dt$, and $dp = -p(1-p)(\lambda + \mu q)dt$. To compute the expectancy, notice that with probability $p \times \mu q dt$, one learns that $\theta = 1$ and the future (expected) payment is $u(1, m + kdt)$. With complementary probability the state of the world remains unknown and the future payment is given by the Taylor expansion

$$u(p,m) - (\lambda+\mu q)p(1-p)\frac{\partial u}{\partial p}(p,m)dt + \frac{\partial u}{\partial m}(u,m)kdt$$

All in all, the Bellman equation is

$$u(p,m) = \frac{1}{r}\max_{k+q\leq 1}\Big\{A(p,m)+q\times B(p,m)+k\times D(p,m)\Big\} \tag{4.3}$$

where

- $A(p,m) = \lambda r(m-l)p - \lambda p(1-p)\dfrac{\partial u}{\partial p}(p,m) - \lambda p \times u(p,m)$
- $B(p,m) = \mu p \times u(1,m) - r\alpha_q - \mu p(1-p)\dfrac{\partial u}{\partial p}(p,m) - \mu p \times u(p,m)$
- $D(p,m) = -r\alpha_k + \dfrac{\partial u}{\partial m}(p,m)$

Equation 4.3 is linear in the control variable (k, q). Therefore it is without loss of generality to restrict oneself to three possible behaviours at the optimum: $(k, q) \in \{(0,0), (0,1), (1,0)\}$.

Write $u_{(1,0)}$ the payoff function of the player if she commits to the strategy $(k,q)=(1,0)$. Define similarly $u_{(0,0)}$ and $u_{(0,1)}$. I consider the three possible cases alternatively.

Waiting is optimal: $(k,q)=(0,0)$**.** This corresponds[5] to $B(p,m)<0$ and $D(p,m)<0$. With $k=q=0$, Equation 4.3 writes

$$(r+\lambda p)\times u(p,m)+\lambda p(1-p)\frac{\partial u}{\partial p}(p,m)=\lambda r p(m-l) \tag{4.4}$$

which unique solution is $u_{(0,0)}(p,m)=(m-l)\dfrac{\lambda r}{\lambda+r}p+C_1(m)(1-p)\Big(\dfrac{1-p}{p}\Big)^{r/\lambda}$. As a player optimally waits in the neighbourhood of $\{0\}\times[0,l]$, it needs be that $C_1(m)=0$ for all $m\in[0,l]$. Consequently, $u_{(0,0)}(p,m)=\dfrac{\lambda r}{\lambda+r}p\times(m-l)$.

For waiting to be optimal, it needs be that both $B(p,m)<0$ and $D(p,m)<0$. We can plug $u_{(0,0)}$ in these equations. The former condition amounts to

$$p<\frac{\alpha_q}{\mu}\times\frac{(\lambda+r)^2}{\lambda-\alpha_k(\lambda+r)}\times\frac{1}{1-\exp\big(-(l-m)(\lambda+r)\big)}$$

There are two facts to notice. (i) The higher α_q (namely, the higher the cost of research) and the lower μ (namely, the lower the efficiency of the research), then the easier it is to fulfill condition $B(p,m)<0$. Indeed, the more costly and the less efficient research is, the better waiting fares compared to research. (ii) The closer to l the stock of protection m is, the less demanding the condition on p it must be for waiting to be optimal. Indeed, when m is small, research can be very valuable and this value is decreasing with m (up to the point where $m=l$ and research is useless).

The latter condition $D(p,m)<0$ does not depends on m and amounts to $p<p^*=\alpha_k\times\dfrac{\lambda+r}{\lambda}$. In words, when protection is expensive (large α_k) or when discounting is large (small r), the area where waiting dominates protection is large. Figure 4.1 shows that waiting is optimal on the left of the two curves in the (p,m)-plane.

[5] I consider here strict optimality, not indifference.

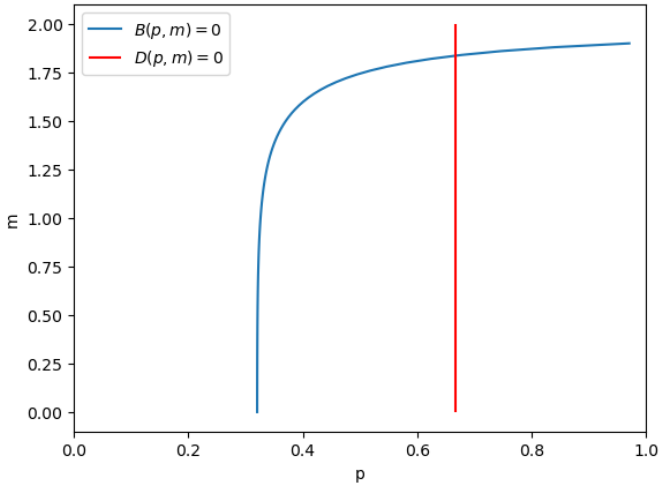


Figure 4.1: Optimality of waiting in the (p, m)-plane. Waiting is better than research on the left of the blue curve given by equation $B(p, m) = 0$. Waiting is better than protection on the left of the red line given by equation $D(p, m) = 0$. Specification of the parameters: $r = 1, l = 2, \lambda = 3, \mu = 5, \alpha_k = 1/2, \alpha_q = 1/10$.

Searching is optimal: $(k, q) = (0, 1)$**.** This corresponds to $B(p, m) > \max\{0, D(p, m)\}$. Computations show that

$$u_{(0,1)}(p,m) = -\alpha_q + \frac{1}{\lambda+\mu+r}\Big(r\lambda(m-l)+\mu u(1,m)+(\lambda+\mu)\alpha_q\Big)p + C_2(m)(1-p)\Big(\frac{1-p}{p}\Big)^{r/(\lambda+\mu)}$$

where $C_2(m)$ is the constant of integration. As m is a parameter of the differential equation, the constant is an implicit function thereof.

Protecting is optimal: $(k, q) = (1, 0)$**.** This corresponds to $D(p, m) > \max\{0, B(p, m)\}$. Equation 4.3 is then a first-order linear partial differential equation (PDE) that writes

$$\frac{\partial u}{\partial m}(p,m) - \lambda p(1-p)\frac{\partial u}{\partial p}(p,m) = (r+\lambda p) \times u(p,m) + r\big(\alpha_k - \lambda p(m-l)\big)$$

the solution of which is

$$\begin{aligned} u_{(1,0)}(p,m) = &-\alpha_k + \frac{\lambda r}{\lambda+r} \times p(m-l) + \frac{\lambda}{(\lambda+r)^2}\Big(\alpha_k(\lambda+r)+r\Big) \times p \\ &+ C_3 \times (1-p)\Big(\frac{1-p}{p}\Big)^{r/\lambda}\Big(m + \frac{1}{\lambda}\log\big(\frac{p}{1-p}\big)\Big) \end{aligned} \tag{4.5}$$

As the differential equation depends not only on p but also on m, the constant of integration C_3 is a "true" constant, not an implicit function of the variable m.

4.4 Contradictions and the absence of smooth pasting.

Let us assume in the section that a player can either protect or wait, research is not possible. The question is therefore whether one should build or wait. There is indifference between the two actions when the cost of building exactly equals its future benefits. We work in the (p, m)-plane.

Then, in (p, m), one is indifferent between building protection and waiting if and only if the marginal cost equals the incremental protection, namely $\frac{\partial u}{\partial m}(p, m) = r\alpha_k$. I start with this equation, find that it does depend on p only, deduce a p^* that is a cut-off between waiting on its left and building protection on its right. This allows me to show that the constant C_3 is equal to 0. I finally derive a contradiction.

The difficulty I have is to find *where* the contradiction lies, namely which hypothesis or argument is false. Let us start with the indifference condition $(E_1) : \frac{\partial u}{\partial m}(p, m) = r\alpha_k$ (namely $D(p, m) = 0$). The set of solutions $\mathcal{S}$ of Equation (E_1) consists in the indifference curve between waiting and protecting in the (p, m)-plane.

Cut-off in p, independent of m. The first argument is that both $u_{(0,0)}$ and $u_{(1,0)}$ are *linear* in the variable m. Consequently, u is linear in the neighbourhood of m, whether below or above. Condition (E_1) is therefore an equation on p only and there exists a $p^* \in]0, 1[$ such that for all $(p, m) \in [0, 1] \times [0, l[$, it is stricly better to wait if and only if $p < p^*$. Geometrically, the indifference curve $\mathcal{S}$ is a vertical line in the (p, m) plane.

Computation: $C_3 = 0$. Let us assume that $C_3 > 0$ and derive a contradiction. Denote by $u_h(p, m)$ the solution to the homogeneous PDE (for $k = 1, q = 0$). For any $(p, m) \in \mathcal{S}$, value

matching hence writes

$$0 = \underbrace{-\alpha_k + \frac{\lambda}{(\lambda+r)^2}\Big(\alpha_k(\lambda+r)+r\Big) \times p}_{f(p)} + C_3 \times u_h(p,m) \tag{4.6}$$

Let us take (p,m) and (p',m') in $\mathcal{S}$ such that $m < m'$. They both satisfy Equation 4.6. Yet u_h is strictly increasing in the second variable. Consequently one has

$$\begin{cases} 0 = f(p') + C_3 u_h(p',m') > f(p') + C_3 \times u_h(p',m) \\ 0 = f(p) + C_3 \times u_h(p,m) \end{cases}$$

Consequently, it needs be that $p \neq p'$. This contradicts the existence of a unique p^* established above. Therefore the claim that $C_3 > 0$ is false. By symmetry it cannot be that $C_3 < 0$ either, hence $C_3 = 0$.

Computation of p^*. The difficulties come here. Value matching requires that $u_{(0,0)}(p,m) = u_{(1,0)}(p,m)$ for all $(p,m) \in \mathcal{S}$. This yields $\hat{p} = \alpha_k \dfrac{(\lambda+r)^2}{\lambda(\alpha_k(\lambda+r)+r)}$. Yet Condition (E_1) requires that in the neighbourhood of every $(p,m) \in \mathcal{S}$, $\dfrac{\partial u_{(0,0)}}{\partial m}(p,m) = \dfrac{\partial u_{(1,0)}}{\partial m}(p,m) = r\alpha_k$ This is *a* smooth pasting property (*the* smooth pasting condition would have equality of gradients, not only of the second partial derivative). This yields $p^* = \alpha_k \times \dfrac{\lambda+r}{\lambda}$. There cannot be two values. $\hat{p} \geq p^*$ whenever Assumption 4.1 holds. (Notice also that both values are smaller than 1 if and only if $\alpha_k \leq \lambda/(\lambda+r)$, some consistency remains.)

Where is the mistake? In the sequence of arguments that is presented, there must be a mistake somewhere. My guess goes to the smooth pasting condition. Maybe the utility function does not verify smooth pasting on the boudary between waiting and building protection. But then, on what ground is it legitimate to study condition (E_1) in the first place?

At this point, many questions are left open. In the Conclusion, I discuss a research program that builds on the partial results that are presented.

5 Conclusion

This thesis manuscript represents a three-years long effort. Instead of briefly summing up the main results, I will take advantage of this conclusion to try and draw some perspectives. Such perspectives are meant to sketch a research program.

Chapter 3 stages a game of repeated cheap talk. It investigates whether ideological competition between experts may be profitable to a decision maker (or Member of Parliament). The answer is positive as stated by Theorem 3.1 granted that at each period there is a single binary and unidimensional state of the world that determines the payment possibilities of all the players. In other words, experts compete not only on the different policies they propose to the decision maker, but also on the very state of the world. The rationale behind is that policies are incompatible with one another and that the state of the world determines which of them is good for the decision maker, which is not. Indeed, tax cuts or tax raises may be good to the economy but hardly the two at a time.

However, this incompatibility rules out credible possibilities. Some bills are transpartisan. Sometimes both progressive and conservative bills may "work" so long as they are consistently carried on. There is consequently a need to enrich the model and to allow preferred policies of competing experts be somewhat compatible and equally good for the decision maker. Section 3.8 takes the opposite stand to incompatibility and assumes that the quality (with respect to the decision maker) of the policies are independent. I briefly explain there why ideological competition may turn out harmful: incentives to lie increase with the number of experts because they entail lower consequences (lying is less detected).

A fruitful extension of Chapter 3 is hence the following. First, one should state the results

of Section 3.8 more precisely. Then, the purpose is to fill the gap between full independence and full incompatibility between experts' preferred policies. Indeed, the political (or ideological) spectrum should allow all the nuances of correlations between policies. "Socialist" and "green" policies may indeed differ. Yet it is conceivable that, due to their proximity, the fact that the green bill turns out today to be good for the decision maker gives insight as to how good the socialist bill can be. As a consequence, the correlation between the quality of the policies of two experts precisely codes the ideological distance. (Presumably Marx and Hayek have incompatible *Weltanschauungen*, so they are negatively correlated.)

The second perspective concerns Chapter 4. This chapter stages a game of strategic experimentation with three exponential bandits. I still struggle with the decision problem. I wish to study the strategic problem with two players. It is natural in the literature to assess the magnitude of the free riding effect: to what extend do players free ride on the research expenditures of others? Do they have more incentives to free ride when the alternative is to build protection? or to wait?

Afterwards, I wish to extend the model along two orthogonal directions that are peculiar to the political problem: how best to protect society against natural disasters. The policy about natural disasters is saturated by the precautionary principle. The intellectual birth of the precaution principle is hard to date. One can argue that it starts with *The Imperative of Responsability* in 1979 (Jonas, 2008) which first states extensively that the doings of the current generation should not impede the possibilities to have a decent life for the future generations. One can argue that it dates from its legal ascertainment that consists in the Rio Declaration in 1992. Both formulations insist on the dangers of irreversibility: the character of irreversibility of an event should precisely refrain one from acting. Irreversibility should be understood with respect to outcomes of a disaster. For instance, casualties are irreversible whether or not the disaster precludes us actions in the future.

In economics, the principle of irreversibility is given a slightly different meaning. An action contains an element of irreversibility when, formally, it shrinkens one's action space in the future: one ties their hands. Hence Arrow and Fisher (1974) explain that, even though they are risk

neutral, players may have a perceived risk aversion because they intend to have flexibility and open possibilities in the future. Likewise, Epstein (1980) proves (under the assumptions of his model) that the more accurate the information one gains tomorrow, the less inclined they are to act today. The precautionary principle, in economics, is an incentive not to act today to have actions available tomorrow.

The concepts of precaution and irreversibility are used somewhat differently in the political and philosophical literature and in the economics literature. Both literature can help to shed light on the situation depicted in Chapter 4. I wish to delve deep in these concepts.

Furthermore, natural disasters are rare but dramatic events. It is well know since Kahneman and Tversky (1979) that small probabilities are over-evaluated (see also Prelec, 1998): this suggests that standard expected utility theory cannot handle this problem. Also, it is not clear at all that politicians who manage the possible forthcoming occurence of a natural disaster are risk neutral. Assuming risk aversion seems in many cases more accurate. In such case, natural disasters may also be modelled as background risks, namely risks that are not hedgeable. The literature on background risk (e.g. Gollier and Pratt, 1996, Quiggin, 2003) explains how a background risk (here the natural disaster) affects orthogonal risks–there is substituability that risk aversion only cannot account for.

A line of research is consequently to drop the assumption that players are risk neutral and that they use expected utility theory altogether (but a modification thereof, as in Yaari, 1987). One should justify carefully what the hypotheses are and how it gives a more accurate description of the problem of Chapter 4. Of course, a modified model of strategic experimentation must be developed accordingly. This program fits with my growing interest in the theory of utility.

List of Figures

4.1 Optimality of waiting in the (p, m)-plane. Waiting is better than research on the left of the blue curve given by equation $B(p, m) = 0$. Waiting is better than protection on the left of the red line given by equation $D(p, m) = 0$. Specification of the parameters: $r = 1, l = 2, \lambda = 3, \mu = 5, \alpha_k = 1/2, \alpha_q = 1/10$. 117

List of Tables

1.1 Matrix of payoffs of the Prisoner's Dilemma. 25

2.1 Counterexample to Proposition 2.3 for a general prior 54
2.2 Counterexample to Proposition 2.4 for a general prior 55
2.3 Undesirable beliefs in case of monotonic games 58
2.4 Case $|S| = 2$ 64

3.1 Summary of the benefice of competition among experts. 80

A Résumé en français

A.1 Introduction

La théorie des jeux est l'étude des décisions interactives entre agents. Une décision est interactive dès que le paiement d'un agent dépend des décisions prises par les autres (et non uniquement des siennes ou du hasard). Une telle situation peut être modélisée comme un *jeu* entre *joueurs* qui sont instrumentalement rationnels, c'est-à-dire qu'ils optimisent leurs choix sous contraintes.

La théorie des jeux se situe à l'intersection des mathématiques appliquées et des sciences sociales. Elle a émergé dans les années 1940 avec les travaux de Von Neumann and Morgenstern (1944) et de Nash (1950b) ; elle s'est épanouie depuis. Elle est désormais autant un champ de recherche en soi qu'une boîte à outils pour les autres disciplines (par exemple, les sciences politiques, la biologie). La théorie des jeux prend une place particulièrement importante en économie. En effet, le cadre de joueurs en interaction qui sont motivés par leurs intérêts (et l'individualisme méthodologique) correspond aux fondements intellectuels de la discipline.

Le concept-clé de la théorie des jeux est la notion d'*équilibre* de Nash. Un équilibre est un plan d'actions tel qu'aucun joueur préférerait dévier et changer de décision unilatéralement. Il est aisé de prouver qu'un équilibre existe sous des hypothèses et conditions faibles. Il indique aux décideurs quels plans d'actions sont possibles, stables et dont on peut s'attendre à ce qu'ils émergent de joueurs stratégiques.

Depuis les jeux ponctuels (*one-shot*) en information totale étudiés par Nash, la théorie des jeux a investi de nombreux problèmes – information partielle, cadres répétés ou dynamiques, notions de coopération, communication stratégique... – et a défini de nouveaux concepts au-delà

du simple équilibre de Nash pour y répondre.

Cette thèse est motivée par deux domaines d'intérêts distincts. D'une part, elle s'intéresse à des questions de sciences politiques : schématiquement, elle étudie les décisions d'un décideur public mal informé et les interactions qu'il entretient avec des individus mieux informés. D'autre part, cette thèse offre des réponses issues de la théorie des jeux à ces questions politiques : elle repose sur l'économie des asymétries d'information et sur la transmission stratégique d'information.

Plus spécifiquement, cette thèse comporte trois chapitres indépendants qui sont autant d'articles.

- Le chapitre 2, intitulé *Certification of coalitions in games of imperfect information*, définit et analyse un modèle de jeux coalitionnels pour lesquels chaque joueur connaît en principe uniquement les valeurs des coalitions auxquelles il appartient. J'y étudie sous quelles conditions un planificateur social peut inciter les joueurs à dévoiler leur information privée. Le cadre politique de ce modèle de théorie des jeux est la provision d'eau aux particuliers en France : les communes ont tendance à se regrouper (en coalitions) pour réduire les coûts sans pour autant savoir parfaitement quelles sont les capacités de leurs pairs ;
- Le chapitre 3, intitulé *Ideological competition and the Taming of Experts*, définit et analyse un jeu de « communication facile » (*cheap talk*) répétée entre un décideur (ici, une commission parlementaire ou son député représentatif) et des experts aux préférences politiques transparentes. La règle est double : (i) sont auditionnés des experts aux biais politiques opposés de sorte qu'il y ait une concurrence idéologique ; (ii) chaque expert peut être licencié à tout moment pour être remplacé par un pair de même biais. Ce chapitre apprécie ce qu'apporte la règle (i) à la règle (ii) pour des députés qui veulent que les experts soient aussi sincères que possible ;
- Le chapitre 4, intitulé provisoirement *Strategic Experimentation and Protection*, est à l'état d'ébauche. Il définit un modèle d'expérimentation stratégique à deux joueurs qui font face à la même possibilité d'une catastrophe naturelle (comme l'ont été les inondations récentes en Europe de l'Ouest). Cette incertitude a deux têtes : la catastrophe aura-t-elle lieu un jour ? et si oui, quand ? Chaque joueur dispose de trois actions : dresser de la

protection, diligenter de la recherche pour lever la première incertitude ou adopter une attitude attentiste (elle, sans coût). La protection peut s'avérer être un gâchis *ex post* mais obéit à un principe de précaution.

A.2 La certification de coalitions dans les jeux à information imparfaite

La question de recherche. Concevoir et mener à bien des projets publics à grande échelle est notoirement difficile. Ils subissent fréquemment retards, coûts additionnels voire même annulations pures et simples. Ces revers ont deux explications possibles : difficultés matérielles (ingénierie, technologie, etc.) et difficultés institutionnelles (désaccords entre parties prenantes, multiplicité des sources de financement, dilution des responsabilités, etc.).

Une multitude d'échecs de projets impliquant plusieurs structures hétérogènes tombent sous la seconde explication. Ainsi, telle structure peut disposer de la priorité politique tandis que telle autre sait mieux ce qui est souhaitable d'entreprendre. Par exemple, l'approvisionnement d'eau en France est sous la responsabilité des communes. Elles peuvent se regrouper si elles le désirent – souvent pour baisser les coûts. L'État (représenté par le préfet) est un facilitateur et peut prendre la main si une commune fait défaut à ses obligations : il joue un rôle de planification.

Au-delà de la délimitation légale des responsabilités subsiste une incertitude : qui, parmi les communes, souhaite quoi ? et surtout, qui sait quoi ? Chaque commune sait ce que ses administrés souhaitent et ce dont elle est capable – son budget, ses capacités d'ingénierie, la qualité de son sol, etc. ; elle peut ignorer quelles sont les capacités d'une commune voisine. L'État, à plus forte raison, est lointain : il dispose d'une information granulaire seulement.

Ainsi, prenons l'exemple de la construction d'une usine de traitement des eaux usées. Quand une commune s'associe avec deux de ses pairs, elle ne connaît pas nécessairement leurs aptitudes. Elle peut savoir, disons, qu'à elles trois les communes sont en mesure de construire l'usine alors qu'associée à l'une des deux seulement, elle n'y parviendra pas. Mais elle peut ignorer si ses deux pairs, à eux deux, peuvent construire l'usine. En d'autres termes, chaque commune a

une information seulement partielle sur les possibilités de coopération de ses pairs. L'État, en principe, est d'autant plus ignorant.

Cela soulève la question abstraite suivante : comment des entités qui coopèrent peuvent-elles partager les dividendes de la coopération lorsqu'il subsiste de l'incertitude sur les capacités de chacune ? Comment le planificateur social peut-il arbitrer au mieux entre les besoins et les ressources des entités ? Comment l'information peut-elle être dévoilée au cours du processus de coopération – en particulier au planificateur social ? Quel serait un partage équitable entre parties mal informées ?

La théorie des jeux coalitionnels offre le bon cadre pour répondre à ces questions. Un jeu coalitionnel représente ce que chaque *coalition* (sous-groupe de joueurs) peut obtenir si tous les membres de la coalition agissent ensemble (sa *valeur, worth*). Un jeu étant donné, deux questions : quelle partition des joueurs va se former ? Comment ensuite partager les fruits de la coopération ?

Cependant, la théorie des jeux coalitionnels suppose que la connaissance des valeurs des coalitions est commune. L'ignorance de chaque municipalité à propos des autres montre que cette hypothèse ne tient pas systématiquement. Le chapitre 2 suppose que chaque joueur connaît les valeurs des coalitions auxquelles il appartient, pas nécessairement celles des autres.

Je définis alors un jeu bayésien stratégique où (i) les joueurs disposent de cette information et (ii) le jeu coalitionnel sous-jacent est *simple*, c'est-à-dire qu'une coalition est soit gagnante, soit perdante : les joueurs certifient alors leurs capacités à leur avantage.

Le modèle. Le jeu stratégique est à deux étapes. Tout d'abord, les joueurs peuvent *certifier* les valeurs des coalitions auxquelles ils appartiennent, si les membres desdites coalition le décident à l'unanimité. Ensuite, un planificateur social met à jour ses croyances à propos du jeu coalitionnel sous-jacent, impose la formation de la *grande coalition* et alloue à chaque joueur son paiement selon une solution coopérative prédéterminée.

Soit $N = \{1, ..., n\}$ l'ensemble des joueurs d'un jeu coalitionnel $\overline{v} : 2^N \to \{0, 1\}$. $\overline{v}$ est superadditif, c'est-à-dire que $\overline{v}(S \cup T) \geq \overline{v}(S) + \overline{v}(T)$ pour toutes coalitions disjointes S et

$T \subseteq N$[1].

Les joueurs et le planificateur social partagent une croyance *a priori* (un *prior*) μ sur l'ensemble $\mathcal{G}$ des jeux simples et superadditifs. Puisque les joueurs connaissent les valeurs des coalitions auxquelles ils appartiennent, leurs croyances *ex interim* sont données par les vecteurs de croyances conditionnelles $\mu_i^{\overline{v}}, i \in N$. Spécifiquement, pour tout $w \in \mathcal{G}$,

$$\mu_i^{\overline{v}}(w) = \frac{\mu(w)}{\mu(\mathcal{G}_i^{\overline{v}})} \times \mathbb{1}_{\mathcal{G}_i^{\overline{v}}}(w)$$

où $\mathcal{G}_i^{\overline{v}} := \{w \in \mathcal{G} : \forall S \ni i, w(S) = \overline{v}(S)\}$ est l'ensemble des jeux coalitionnels possibles.

Les stratégies. Les membres d'une coalition peuvent certifier sa valeur s'ils le décident à l'unanimité. Ils peuvent dévoiler ou faire de la rétention d'information mais ils ne peuvent pas mentir. Formellement, soit $\Delta = \{(\mu_j^v)_{j \in N} : v \in \mathcal{G}, \mu(v) > 0\}$ l'ensemble des croyances *ex interim* possible, $\Delta_i = \{\mu_i^v : (\mu_j^v)_{j \in N} \in \Delta\}$. Une stratégie (pure) du joueur $i \in N$ est une fonction a_i qui à chaque $\mu_i \in \Delta_i$ associe un vecteur $\left(a_i(S|\mu_i)\right)_{i \in S \subseteq N} \in \{0,1\}^{2^{n-1}}$. Si $a_i(S) = 1$, alors i souhaite dévoiler la valeur de S ; si $a_i(S) = 0$, alors non (pour simplifier les notations, j'omets le conditionnement par μ_i). Puisque l'unanimité est requise, S est certifiée si, et seulement si, $\prod_{i \in S} a_i(S) = 1$.

Qu'observe le planificateur social ? Il n'observe ni les stratégies des joueurs, ni même leurs choix (conditionnellement à leurs types) : il observe uniquement si chaque coalition a été certifiée (et si oui, sa valeur). Ainsi, si une coalition n'est pas certifiée, le planificateur social sait qu'au moins un des membres ne l'a pas souhaité – mais il en ignore l'identité. Le concept d'équilibre est l'équilibre bayésien parfait. Soit $\hat{\mu}(a)$ la croyance *ex post* du planificateur social et $v_{\hat{\mu}(a)} = \mathbb{E}_{\hat{\mu}(a)}[v]$ le jeu coalitionnel espéré associé. Alors le planificateur social impose l'allocation $\varphi(v_{\hat{\mu}(a)})$ où φ est une solution coopérative connue des joueurs, supposée efficiente et croissante par rapport aux contributions marginales[2].

Les joueurs sont neutres au risque et leur unique objectif est de modeler les croyances du

[1] En particulier, comme $\overline{v}$ est simple, si $\overline{v}(S) = 1$ (S est gagnante), alors nécessairement $\overline{v}(T) = 0$ (T est perdante).

[2] Un exemple est la valeur de Shapley. Ce n'est pas le seul.

planificateur à leur avantage. La croyance *ex post* $\hat{\mu}(a)$ est déterminée par la règle de Bayes, le planificateur postulant correctement les stratégies choisies par les joueurs à l'équilibre.

Résultats principaux. Un résultat simple d'existence d'équilibre :

Proposition A.1. *Le profil d'actions suivant est un équilibre : chaque joueur $i \in N$ certifie $\{i\}$ si et seulement si il s'agit d'une coalition gagnante et ne certifie aucune autre coalition (quel que soit le* prior *μ).*

Cela tient à l'exigence d'unanimité : si $a_i(\{i,j\}) = 0$ alors j n'a aucune incidence sur la certification de $\{i,j\}$ et $a_j(\{i,j\}) = 0$ est une meilleure réponse.

Structure des jeux superadditifs et certification indirecte. Soit $S \subseteq N$. Supposons que S ait été certifiée (directement) gagnante. Alors par superadditivité, le planificateur apprend également que tout sur-ensemble de S est gagnant (que ses membres aient décidé de le certifier ou non) et il en va de même de tous les sous-ensembles de $N \setminus S$ qui sont certifiés perdants. Ce sont des certifications *indirectes*. Les membres de S doivent prendre en compte, lorsqu'ils formulent leurs stratégies, ce phénomène de second ordre.

Heuristique des coalitions gagnantes.

Au vu des certifications indirectes, il semblerait que certifier les coalitions gagnantes soit optimal, ne pas certifier les coalitions perdantes soit optimal. En effet, si $i \in S$ parvient à certifier S gagnante, alors tous les sur-ensembles de S auquel i appartient sont certifiés gagnants tandis que tous les sous-ensembles de $N \setminus S$ auxquels i n'appartient *pas* sont certifiés perdants. Cependant, cette heuristique est fausse en général car certifier comme gagnante une coalition peut entraîner la dépréciation de la valeur (en espérance, aux yeux du planificateur) d'autres coalitions auxquelles le joueur appartient, qui sont « orthogonales » :

Proposition A.2. *Supposons que $n \geq 4$ et que φ est strictement croissante par rapport aux contributions marginales. Alors il existe $(\overline{v}, \mu)$ et un profil a tels que (i) a est un équilibre, (ii) $i \in S \subseteq N$ ne certifie pas la coalition gagnante S et (iii) le faire diminuerait strictement son paiement.*

La preuve est donnée par un exemple.

Quand l'heuristique se vérifie. Cela étant, l'heuristique est souvent vérifiée. Ce paragraphe donne des conditions suffisantes pour lesquelles il est effectivement optimal de certifier toutes et uniquement les coalitions gagnantes. Une coalition gagnante dont tous les sous-ensembles perdent ne peut pas être certifiée indirectement. Elle est dite *minimale gagnante* : notons $\mathcal{M}(\overline{v})$ l'ensemble de ces coalitions pour $\overline{v}$. Appelons enfin *inconnue* une coalition $S \subseteq N$ telle que $v_\mu(S) \in]0,1[$, c'est-à-dire dont la valeur est incertaine *a priori* et notons $\mathcal{X}(\mu)$ l'ensemble de ces coalitions.

Deux coalitions minimales gagnantes sont indépendantes au sens où la certification de l'une n'apporte aucune information indirecte sur la valeur de l'autre. Partant, elles sont le bon objet à étudier. Dans les trois propositions suivantes, chaque joueur préfère strictement qu'une coalition gagnante soit certifiée (fût-ce indirectement) et qu'une coalition perdante ne le soit pas. Que chaque membre d'une coalition préfère qu'elle soit certifiée ne suffit pas à ce que ce que la coalition le soit. Cela tient à l'unanimité : il s'agit d'un « manque de coordination ».

Les résultats reposent sur un principe de raison insuffisante pour le planificateur : le *prior* donne la même probabilité (uniforme) à tout jeu coalitionnel qui est possible. Rappelons enfin qu'un équilibre est *fort* quand aucune coalition n'a intérêt à dévier.

Proposition A.3. *(Asymétrie horizontale.) Supposons que $\mathcal{X}(\mu) \subseteq \mathcal{M}(\overline{v})$, c'est-à-dire que seules les coalitions minimales gagnantes sont inconnues. Les équilibres sont les suivants. Pour tout S, ou bien tous les joueurs certifient avec probabilité 1, ou bien au moins deux joueurs certifient avec probabilité nulle (et les autres membres sont indifférents). L'unique équilibre fort voit toutes les coalitions certifiées ; sinon, il existe un manque de coordination.*

L'asymétrie vers le haut et l'asymétrie vers le bas suivent le même principe.

Proposition A.4. *(Asymétrie vers le bas.) Soit $S \in \mathcal{M}(\overline{v})$. Supposons que $\mathcal{X}(\mu) = \{S, S \setminus \{i\} : i \in S\}$. À l'équilibre, aucune coalition $S \setminus \{i\}$ n'est jamais certifiée. Dans l'équilibre fort, S est certifiée. Sinon, S n'est pas certifiée par manque de coordination.*

Proposition A.5. *(Asymétrie vers le haut.) Soit $S \in \mathcal{M}(\overline{v})$. Supposons que $\mathcal{X}(\mu) = \{S, S \cup k :$*

$k \in N \setminus S\}$. *Alors le principe est le même : chaque joueur appartenant à une coalition gagnante inconnue souhaite qu'elle soit certifiée (directement ou indirectement) ; ce n'est pas systématique à l'équilibre du fait des manques de coordination. Enfin, si* $i \in S$ *sait que* S *sera certifiée directement, alors* $a_i(S \cup k)$ *peut être choisi librement – cela n'a aucune incidence sur la croyance* ex post *du planificateur.*

Dans les trois cas, ce sont des conditions suffisantes. L'article discute ensuite de la robustesse des résultats. Par exemple, l'hypothèse de *prior* uniforme n'est pas nécessaire : une perturbation de ce *prior* fonctionnerait. Toutefois n'importe quel *prior* qui laisserait les coalitions inconnues inchangées mais en distordant les poids donnés aux différents jeux inconnus rendrait les propriétés fausses en général.

Aussi, il est tentant d'étudier le cas d'une asymétrie vers le haut et vers le bas ($\mathcal{X}(\mu) = \{S, S \setminus i, S \cup k\}$). Les résultats ne tiennent pas.

A.3 Compétition idéologique et domestication des experts

La question de recherche. Nos sociétés reposent largement sur la connaissance technique et scientifique. Avant de faire et voter la loi, les députés doivent donc obtenir de telles connaissances. Cela les met face à une difficulté : d'une part, leurs décisions devraient aussi fondées sur la science que possible ; d'autre part, acquérir une expertise suffisante sur tous les sujets est impossible. Encourager la spécialisation de chaque député sur un sujet précis ne saurait être une solution valable. En effet, seuls quelques députés seraient pertinents pour chaque loi ; cela affaiblit le caractère délibératif des parlements.

Reposer sur l'expertise externe – groupes d'intérêts, *think tanks* et universitaires – permet de pallier ce dilemme. Par exemple, syndicalistes comme représentants du patronat disposent d'une information de première main sur les besoins du marché du travail. Écouter leurs recommandations permet de voter des lois mieux adaptées : il y a une délégation épistémique.

Cet appel massif[3] à l'expertise extérieure requiert des pratiques saines et des règles strictes pour les mettre en place. En effet, la sincérité d'un expert ne saurait être tenue pour acquise. Un représentant patronal qui promeut une baisse des charges le fait-il par intérêt corporatiste ? Ou parce qu'il pense sincèrement que cela permettra de fluidifier le marché du travail, et par suite sera bénéfique à la société ? Symétriquement, comment interpréter un syndicaliste qui promeut l'augmentation du salaire minimum ? En bref, les experts ont des préférences politiques bien identifiées et l'enjeu est de savoir si un conseil est sincère ou motivé par la poursuite d'un agenda personnel.

Aussi, des règles précises doivent être mises en place pour garantir la transparence, l'équilibre entre différentes opinions politiques et l'exigence de sincérité.

Prenons le cas de l'Assemblée nationale. Légalement, chaque expert doit se présenter à une convocation. La majorité comme l'opposition peuvent inviter librement des experts, tout le spectre politique est représenté : cela garantit qu'il y ait une « concurrence idéologique » entre experts. Par ailleurs, les préférences des experts sont facilement identifiées : le biais d'un syndicaliste est établi ; ceux d'un économiste marxiste ou de l'école de Chicago également.

En principe, un expert est tenu de dire la vérité – et peut être poursuivi sinon. Cependant, la tromperie intentionnelle est difficile à établir : un mauvais conseil est-il dû à l'incompétence ou à l'insincérité ? Même les meilleurs experts se trompent çà et là.

Afin de se protéger contre l'incompétence et l'insincérité, les experts peuvent être licenciés à tout moment s'ils ne montrent pas satisfaction. Malgré la précarité de leurs positions, ils parviennent généralement à se maintenir et conseillent les députés de manière répétée.

Il convient donc d'interroger la capacité des députés à domestiquer efficacement les experts – de sorte que les médiocres comme les menteurs soient facilement identifiés. Cette discipline se manifeste par deux canaux. Premièrement, les experts sont disciplinés par la peur d'être remplacés par un pair de même biais idéologique. Secondement, chaque expert est discipliné par son homologue actuel de biais opposé (concurrence idéologique).

[3]Par exemple, 3 000 auditions ont lieu par législature à l'Assemblée nationale tandis que le Parlement européen commande environ 1 000 rapports techniques par an.

L'objectif de l'article est d'étudier l'efficacité de la discipline par concurrence idéologique en fonction du degré de compétence des experts. Pour ce faire, je développe un modèle de communication facile répétée entre une décideuse (la députée représentative) notée D et des experts aux biais transparents et opposés notés E et E'. Je distingue simultanément selon que les bons experts sont parfaits – leurs signaux sont exacts – ou non ; selon qu'il y a uniquement de bons experts ou que la décideuse ignore si les experts sont compétents ou non.

Le modèle. Les experts se distinguent selon deux paramètres : la préférence idéologique (connaissance commune) et la compétence (information privée). Soit E un expert progressiste ; $E = A$ s'il est compétent/bon, $E = B$ s'il ne l'est pas. (Les mêmes notations avec des primes désignent les experts conservateurs.) Un bon expert est un véritable joueur stratégique, un mauvais expert est un pur bruit. Soit $\mu \in [1/2, 1]$ la probabilité *a priori* que l'expert soit bon (indépendante de l'idéologie).

À chaque période l'état du monde est $\omega \in \{0, 1\}$ (équiprobable). L'expert reçoit un signal $s \in \{0, 1\}$ conditionnellement à ω. Si $E = A, \mathbb{P}(s = 1|\omega = 1) = \mathbb{P}(s = 0|\omega = 0) = p > 1/2$. Si $E = B$, ces probabilités valent $1/2$ (le signal n'est pas informatif). (De même pour E'.) La structure d'information est connaissance commune et les signaux des experts E et E' sont indépendants sachant ω. Période par période, les distributions sont indépendantes.

Soit $r \in [0, 1]$ la probabilité qu'un expert E soit compétent sachant les stratégies jouées et l'histoire jusqu'à présent. C'est sa *réputation*, elle sera la variable d'état. Par définition, $r = \mu$ quand E commence à être auditionné.

Tous les joueurs utilisent des stratégies quasi-markoviennes (stationnaires). La stratégie de E est une fonction $\{0, 1\} \times [0, 1] \to [0, 1]$ qui à (s, r) associe la probabilité d'envoyer le message $m = 1$ (plutôt que $m = 0$). Chaque message est public. Par hypothèse, un mauvais expert est bruité et sa stratégie est fixée uniformément à $1/2$ (quel que soit (s, r)). La stratégie de D a deux aspects. Tout d'abord, le choix de l'action $a \in \{0, 1\}$ sachant les messages et les réputations. Après que l'action a été choisie, les paiements sont distribués et l'état du monde révélé. Ensuite,

la probabilité de licencier E (ou E') sachant l'état du monde, les messages et les réputations[4].

En particulier, la décision de licencier ou non un expert ne dépend pas de l'action a choisie par D. Si E est licencié, un nouvel expert progressiste (aussi noté E) entre et le jeu répété reprend.

Finalement, les paiements. À chaque période, seule l'action a influe sur les paiements. D reçoit 1 si l'action et l'état du monde coïncident, 0 sinon. E reçoit 1 si $a = 1, 0$ sinon. Il n'est donc pas intéressé directement par l'état du monde. (Le paiement de E' est opposé.) Après licenciement, le paiement d'un expert est nul.

Les joueurs partagent, sans perte de généralité, le même taux d'escompte exponentielle $\delta \in]0,\ 1[$. Par construction – les messages sont publics et *ex post* l'état du monde est révélé –, la réputation d'un expert est connaissance commune.

Le concept de solution est l'équilibre markovien stationnaire (quand $\mu < 1$) ou sous-jeux parfait (quand $\mu = 1$).

Résultats principaux. Je regarde uniquement les équilibres qui sont (i) symétriques et (ii) informatifs. La symétrie signifie que les stratégies dépendent uniquement des variables « réelles » et non des identités. Le caractère informatif signifique que les bons experts disent la vérité au moins aussi fréquemment que les mauvais experts (c'est-à-dire au moins une fois sur deux). Par conséquent, si $r < \mu$ il est optimal pour D de licencier E (quels que soient r' et les messages présents).

Interprétation heuristique des messages. Quand $\omega = 1$ les intérêt de A et D coïncident. Ainsi, si A reçoit $s = 1$ il pense (avec probabilité $p > 1/2$) que ses intérêts présents coïncident avec ceux de A. On peut donc attendre que A révèle sincèrement tous les $s = 1$. À l'inverse, c'est quand $s = 0$ que A peut être tenté de mentir et d'envoyer $m = 0$. Ainsi, il y a deux interprétations possibles à $m \neq \omega$. Si $m = 1$ mais $\omega = 0$, c'est une « erreur opportuniste » ; si $m = 0$ et $\omega = 1$, c'est une « erreur sincère » due à l'incompétence – notons toutefois qu'elle peut être le fait d'un expert compétent qui se trompe de temps à autres.

[4]La stratégie de D dépend donc des réputations *et* de l'adéquation entre l'état du monde et le message : elle n'est donc pas markovienne. Les stratégies des experts, par contre, sont markoviennes (les réputations étant la variable d'état).

Posons la restriction suivante aux stratégies de D : pourvu que $r \geq \mu$, D ne licencie pas si l'erreur a été sincère ou si le message et l'état du monde ont coïncidé. Alors :

Proposition A.6. *Supposons qu'il y ait uniquement des experts E.*

- *À l'équilibre, A ne peut pas préférer d'envoyer $m = 1$ quand $s = 0$ et $m = 0$ quand $s = 1$ (quelle que soit $r \geq \mu$) ;*
- *Supposons qu'il existe $r \geq \mu$ tel que A préfère envoyer $m = 0$ indépendemment de son signal. Alors il existe une amélioration parétienne de cet équilibre dans lequel A envoie $m = 1$ en r indépendemment de son signal.*

Cela signifie que, sous condition de Pareto-optimalité, un bon expert révèle toujours honnêtement un signal favorable. Au contraire, donc, une erreur sincère est authentiquement sincère : si $m = 0$, alors $s = 0$.

Enfin, il est systématiquement optimal, pour D de choisir $a = 1$ quand $\mathbb{P}(\omega = 1|m, m', r, r') > 1/2$ (avec indifférence en $1/2$). Cela signifie que D n'a jamais intérêt à faire un sacrifice de court terme (en choisissant une action probablement mauvaise) pour discipliner un expert dans le futur. Ainsi, on peut supposer que les recommandations des experts sont directement suivies par D et que la seule stratégie de cette dernière est le choix de licencier.

La compétition idéologique entre experts est-elle bénéfique à la décideuse ?

Les résultats dépendent de la compétence des bons experts, c'est-à-dire de la valeur de p, d'une part. D'autre part, ils dépendent de si la compétence est privée ($\mu < 1$) ou non. Appelons les bons experts *parfaits* quand $p = 1$.

Les résultats sont les suivants : quand les bons experts sont parfaits ou quand il y a uniquement de bons experts ($\mu = 1$, absence de compétence cachée), la concurrence idéologique est inutile. En effet, que les bons experts soient systématiquement sincères et qu'ils soient licenciés à la première erreur opportuniste est un équilibre sous-jeux parfait. Ainsi, la décideuse extrait toute l'information possible à l'équilibre. Par contre, s'il y a à la fois compétence imparfaite ($p < 1$) et compétence cachée ($\mu < 1$), alors la concurrence idéologique est bénéfique. Plus précisément, pour tout équilibre sous-jeux parfait avec uniquement des experts E, il existe un équilibre avec experts E et E' tel que le paiement (espéré *ex ante*) de la décideuse est strictement supérieur.

En formules :

Proposition A.7. *(Compétence publique : $\mu = 1$.) Supposons que $\delta \geq 2/(2p+1)$. Alors sincérité chez les experts et licenciement pour toute erreur opportuniste est un équilibre sous-jeux parfait qu'il y ait concurrence ou non. Le paiement de D est alors $V_D = p/(1-\delta)$. Le paiement de A est $V_A = \dfrac{1}{2-\delta(1-p)}$ (*idem *pour A').*

Ainsi, quand tous les experts sont compétents et qu'il n'y a pas de doute sur leurs types, D parvient à extraire toute l'information sans concurrence.

Proposition A.8. *(Compétence parfaite : $p = 1$.) Supposons que $\delta \geq 2/3$. Alors sincérité chez les experts compétents et licenciement uniquement quand $r < \mu$ est un équilibre sous-jeux parfait*[5]*.*

Cette capacité à extraire toute l'information ne tient pas quand il y a en même temps compétences privée et imparfaite :

Théorème A.1. *Supposons que $p, \mu < 1$ et que $p \times \mu \times \delta$ soit suffisamment grand. Alors pour tout équilibre avec uniquement des experts E, il existe un équilibre avec des experts E et E' tel que le paiement de D est strictement supérieur.*

Ainsi, il y a là bénéfice à avoir de la concurrence idéologique. La décideuse améliore son paiement *via* deux canaux distincts. Tout d'abord, le canal direct. On peut montrer que, dans cette situation, il n'est pas possible d'obtenir la sincérité totale d'un expert quand il est seul et que sa réputation est proche de 1. Quand deux experts concurrents ont une *même* réputation élevée, ils peuvent rester sincères à l'équilibre. Ensuite, le canal indirect. Sachant l'existence du canal direct, un expert de haute réputation qui est seul (au sens où son concurrent babille) doit être plus sincère que dans le cas sans concurrence. En effet, l'option de sortie de D (qui

[5]Quand il n'y a pas de concurrence, le paiement espéré de D est $V_D(\mu) = \dfrac{\dfrac{\mu}{1-\delta} + \dfrac{1-\mu}{2-\delta}}{1-(1-\mu)\dfrac{\delta}{2-\delta}}$. S'il y a concurrence, le paiement de D est légèrement supérieur bien que les stratégies des experts soient les mêmes. La raison est que la décideuse essaie les experts jusqu'à « tomber » sur un expert parfait. Cet essayage est plus rapide avec deux experts en même temps. La réputation d'un expert après t messages exacts est $r^t = \dfrac{2^t \times \mu}{1+(2^t-1)\mu}$, ce qui converge exponentiellement rapidement vers 1.

est endogène aux stratégies) est plus élevée avec concurrence car avec probabilité positive, deux nouveaux experts vont obtenir de hautes réputations et rester sincères.

Le chapitre discute ensuite d'une alternative avec paiement courant constant (indépendemment des réputations) pour D et d'extensions du modèle – où un mauvais expert est stratégique plutôt que bruité et où il y a un paiement intrinsèque à conseiller.

A.4 Expérimentation stratégique et protection

Ce chapitre est à l'état de projet. Le modèle est bien défini. La résolution du problème de décision – c'est-à-dire du problème du planificateur social – est en cours.

Savoir comment protéger au mieux une société contre des catastrophes naturelles est, c'est notoire, un problème épineux. Le réchauffement climatique, l'instabilité qu'il a provoquée ou les catastrophes naturelles récentes ont remis cette question sur le devant de la scène. Les catastrophes naturelles ont trois particularités : elles sont massives, contiennent une part d'irréversibilité et sont difficiles à prévoir.

Chaque désastre engendre un scandale public. Par exemple, les inondations de juillet 2021 en Europe de l'Ouest ont fait 269 victimes et ont mis en lumière l'impéritie de l'État (en l'occurence, de l'Allemagne, la Belgique et les Pays-Bas). Les interrogations se résument à deux ensembles : comment se fait-il que les autorités aient été incapables de prévoir l'ocurrence de la catastrophe ? Comment se fait-il qu'aucune installation à même de prévenir les dégâts n'ait été disponible ?

La première question interroge le problème du manque d'information. La seconde question adresse le problème du manque de protection, sous hypothèse que la catastrophe survient. Un principe de précaution devrait encourager à dresser de la protection contre la catastrophe. Cependant, la survenance est incertaine et la protection peut s'avérer être un gâchis *a posteriori*. Des contraintes temporelles ou budgétaires ainsi que la multitude des enjeux suggèrent qu'il n'est pas souhaitable de sur-investir dans des risques très hypothétiques.

Par conséquent, il y a trois attitudes possibles face à une catastrophe concevable mais hautement hypothétique pour un politicien : (i) construire de la protection malgré le risque que ce

soit un gâchis *ex post* ; (ii) essayer de lever l'incertitude par la recherche (mais la détection de signaux faibles est ardue) ; (iii) adopter une attitude attentiste (qui libère des ressources).

Ce cadre se généralise à plusieurs pouvoirs publics. Dans le cas de l'inondation en Europe de l'Ouest, trois pays sont concernés. Jusqu'à un certain point, la protection est privée : la construction par la Belgique d'un barrage ne bénéficie qu'à elle. À l'inverse, la recherche est publique : dès que le gouvernement belge sait que la vallée du Rhin est sujette aux inondations dans un futur proche, l'Allemagne et les Pays-Bas le savent également.

L'asymétrie entre une protection privée[6] et une recherche publique, ainsi que la possibilité pour une partie d'être un passager clandestin sur les dépenses en recherche des autres font de cette situation un jeu stratégique.

Ce chapitre cherche à formaliser ce dilemme entre ces trois actions. Il a pour objectif de déterminer le plan d'actions optimal d'un politicien comme fonction de sa croyance subjective que la catastrophe *peut* survenir un jour. Ensuite, je compte étendre le modèle à deux joueurs afin de mesurer l'effet du passager clandestin.

Pour ce faire, un modèle d'expérimentation stratégique à trois actions, c'est-à-dire trois bras, est tout indiqué.

Le modèle. Le temps $t \in \mathbb{R}_+$ est continu et escompté au taux $r > 0$ par deux joueurs neutres au risque. Un *événement* (la catastrophe, l'inondation) est de type $\theta \in \{0, 1\}$. Si $\theta = 0$, alors l'événement ne survient jamais. Si $\theta = 1$, alors l'événement est distribué exponentiellement de paramètre λ. Notons p_0 le prior.

À chaque période $[t, t + dt[$, les joueurs ont une unité de ressource à répartir entre se protéger, chercher et attendre. Soit[7] k_t l'investissement en protection, q_t celui en recherche tels que $k_t + q_t \leq 1$. Jouer (k_t, q_t) coûte $(\alpha_k \times k_t + \alpha_q \times q_t)dt$: attendre se fait sans coût.

La probabilité instantanée d'apprendre θ par recherche est $\theta \times \mu \times q_t \times dt$. Ainsi, l'apprentissage se fait exponentiellement (de paramètre μ) et ne peut advenir que si $\theta = 1$. Cela revient à

[6] Manifestement, la protection ne peut être purement privée. Ainsi, la construction d'un barrage devrait bénéficier à toutes les parties en aval mais serait inutile aux parties en amont. Un tel enrichissement de l'histoire (et donc du modèle) serait un objectif pour l'avenir.

[7] On devrait indexer les stratégies par l'identité de chaque joueur. Je ne m'en embarasse pas.

supposer que : on peut prouver que quelque chose existe, pas qu'il n'existe pas ; et que l'absence de nouvelles est une bonne nouvelle.

Si l'événement survient à la date $t > 0$, le joueur reçoit un paiement forfaitaire $\min\{m_t - l, 0\}$ où $m_t = \int_0^t k_s ds$ est le stock de protection accumulé. L'opérateur min matérialise le fait que trop de protection est un gâchis.

Soit Y_t l'événement (probabiliste) « ni catastrophe, ni apprentissage de θ avant t ». La probabilité subjective que $\theta = 1$ en t est donc $p_t = \mathbb{P}(\theta = 1|Y_t)$. Puisque les actions et les résultats sont publics, p_t est connaissance commune. Les joueurs utilisent des stratégies markoviennes stationnaires symétriques où la variable d'état est le couple (probabilité subjective, stock de protection). Ainsi, une *stratégie* (ou *politique*) est une application $\sigma : [0,1] \times \mathbb{R}_+ \to \{(k,q) : k + q \leq 1\}$ qui est mesurable par rapport à l'information disponible. Dans la mesure où il est inutile d'avoir un stock m excédant l, on peut restreindre σ à $[0,1] \times [0,l]$.

Premiers résultats. *La loi d'évolution des croyances.*

De manière relativement classique[8], la probabilité subjective obéit à l'équation différentielle ordinaire $p'(t) = -p(t)\big(1-p(t)\big)\big(\lambda+\mu q_t\big)$ dont l'unique solution est $p_t = \dfrac{p_0 e^{-\lambda t - \mu \int_0^t q_s ds}}{1 - p_0 + p_0 e^{-\lambda t - \mu \int_0^t q_s ds}}$.

La fonction de paiement.

Quand $p = 0$, il est optimal d'attendre (c'est-à-dire que $k = q = 0$). Pour que le problème ne soit pas trivial, il faut que quand $p = 1$ (ou suffisamment proche), il soit optimal de se protéger plutôt qu'attendre (pour n'importe quel $m < l$). Pour ce faire, il faut que le coût α_k soit petit devant le temps d'arrivée de l'événement. Autrement dit, faisons l'hypothèse paramétrique que $\alpha_k \leq \lambda/(\lambda + r)$.

Notons u la fonction à deux variables de paiement du joueur. On peut calculer indépendemment $u(1,m)$ grâce à l'hypothèse paramétrique. Ensuite, partant de $(p_0, 0)$, le paiement espéré à

[8]Il existe toutefois une nuance. Dans la littérature, lorsque le bras dit « risqué » n'est pas activé, la mise à jour des croyances est gelée. Ce n'est pas le cas ici car du seul fait que la catastrophe ne soit pas survenue jusqu'à présent, la croyance qu'elle puisse survenir dans le futur diminue. La nature de l'équation, elle, est classique.

jouer la stratégie $(k_t, q_t)_{t\geq 0}$ est

$$u(p_0,0) = \int_0^{+\infty} e^{-rt}\Big(-(\alpha_k k_t + \alpha_q q_t) + p_0 e^{-\lambda t - \mu\int_0^t q_s ds}\big(\lambda(m_t - l) + \mu q_t u(1, m_t)\big)\Big)dt$$

L'équation de Bellman.

Par des arguments standards, l'équation de Bellman est

$$\begin{aligned} u(p,m) = \frac{1}{r}\max_{k+q\leq 1}\Big\{&\Big(r\lambda p(m-l)) - \lambda p(1-p)\frac{\partial u}{\partial p}(p,m) - \lambda p u(p,m)\Big)\\ &+q\Big(\mu p \times u(1,m) - r\alpha_q - \mu p(1-p)\frac{\partial u}{\partial p}(p,m) - \mu p u(p,m)\Big)\\ &+k\Big(-r\alpha_k + \frac{\partial u}{\partial m}(p,m)\Big)\Big\}\end{aligned}$$

Cela suppose implicitement que u est dérivable à gauche en p et à droite en m. Cette équation est linéaire en le contrôle (k, q). Ainsi, à l'optimum (k, q) est une solution en coin. Cela donne trois équations différentielles linéaires (dont une aux dérivées partielles quand $k = 1, q = 0$).

Il est alors possible de calculer les trois fonctions d'utilité d'un joueur s'il *s'engage* sur une stratégie éternellement. Ainsi, l'attentisme assure le paiement $(m-l)\dfrac{\lambda r}{\lambda + r}p$. Bien entendu, le modèle autorise le joueur à changer de bras.

La question est actuellement de déterminer les domaines (qui sont des courbes dans le plan (p, m)) où le joueur est indifférent entre deux actions. Une propriété usuellement retrouvée dans la littérature (mais pas systématiquement) est le *smooth pasting*, c'est-à-dire que la fonction d'utilité d'un joueur est C^1 sur ces domaines d'indifférence. Elle semble ne pas se vérifier dans ce modèle.

Bibliography

Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, *84*(3), 488–500. https://doi.org/10.2307/1879431

Andjiga, N.-G., Chantreuil, F., & Lepelley, D. (2003). La mesure du pouvoir de vote. *Mathématiques et sciences humaines*, (163).

Arrow, K. J., & Fisher, A. C. (1974). Environmental Preservation, Uncertainty, and Irreversibility. *The Quaterly Journal of Economics*, *88*(2), 312–319. https://www.jstor.org/stable/1883074?seq=1

Au, P. H. (2015). Dynamic information disclosure. *The RAND Journal of Economics*, *46*(4), 791–823. https://www.jstor.org/stable/43895617

Aumann, R. J. (1976). Agreeing to Disagree. *The Annals of Statistics*, *4*(6), 1236–1239. https://www.jstor.org/stable/2958591

Austen-Smith, D. (1990). Credible debate equilibria. *Social Choice and Welfare*, *7*(1), 75–93. https://www.jstor.org/stable/41105938

Benabou, R., & Laroque, G. (1992). Using privileged information to manipulate markets: Insiders, gurus, and credibility. *The Quarterly Journal of Economics*, *107*(3), 921–958. https://doi.org/10.2307/2118369

Bhattacharya, S., & Mukherjee, A. (2013). Strategic information revelation when experts compete to influence. *The RAND Journal of Economics*, *44*(3), 522–544. https://doi.org/10.1111/1756-2171.12029

Board, S., & Lu, J. (2018). Competitive information disclosure in search markets. *Journal of Political Economy*, *126*(5), 1965–2010. https://doi.org/10.1086/699211

Bolton, P., & Harris, C. (1999). Strategic Experimentation. *Econometrica*, *67*(2), 349–374. https://www.jstor.org/stable/2999588

Buchanan, J. M. (2001). Game theory, mathematics, and economics. *Journal of Economic Methodology*, *8*(1), 27–32. https://doi.org/10.1080/13501780010022802

Chatterjee, K., Dutta, B., Ray, D., & Sengupta, K. (1993). A noncooperative theory of coalitional bargaining. *The Review of Economic Studies*, *60*(2), 463.

Crawford, V., & Sobel, J. (1982). Strategic Information Transmission. *Econometrica*, *50*(6), 1431–1451.

Davies, G. (2024). *HS2: Update following cancellation of Phase 2* (tech. rep. No. HC 128). National Audit Office. https://www.nao.org.uk/reports/hs2-update-following-cancellation-of-phase-2/#downloads

Doyle, J. R. (2013). Survey of time preference, delay discounting models. *Judgment and Decision Making*, *8*(2), 116–135. https://doi.org/10.1017/S1930297500005052

Elomäki, A., & Haapala, T. (2024). Politics of expertise in the European Parliament: Discursive constructions and contestations of expertise by party-political actors. *Journal of European Integration*, *46*(3), 321–339. https://doi.org/10.1080/07036337.2023.2283699

Epstein, L. G. (1980). Decision Making and the Temporal Resolution of Uncertainty. *International Economic Review*, *21*(2), 269. https://doi.org/10.2307/2526180

Farrell, J., & Gibbons, R. (1989). Cheap Talk with Two Audiences. *The American Economic Review*, *79*(5), 1214–1223. https://www.jstor.org/stable/1831447

Forges, F. (2017). Coopération en information incomplète : Quelques modèles stratégiques: *Revue d'économie politique*, *Vol. 127*(4), 467–493.

Forges, F., & Serrano, R. (2013). Cooperative games with incomplete information: Some open problems. *International Game Theory Review*, *15*(2), 1340009.

Gillies, D. B. (1959). Solutions to general non-zero-sum games. In *Contribution to the Theory of Games IV* (pp. 47–85, Vol. 40). Annals of Mathematic Studies.

Gilligan, T. W., & Krehbiel, K. (1989). Asymmetric information and legislative rules with a heterogeneous committee. *American Journal of Political Science*, *33*(2), 459–490. https://doi.org/10.2307/2111156

Gollier, C., & Pratt, J. W. (1996). Risk Vulnerability and the Tempering Effect of Background Risk. *Econometrica*, *64*(5), 1109–1123. https://doi.org/10.2307/2171958

Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of Law & Economics*, *24*(3), 461–483. https://www.jstor.org/stable/725273

Harsanyi, J. C. (1966). A General Theory of Rational Behavior in Game Situations. *Econometrica*, *34*(3), 613. https://doi.org/10.2307/1909772

Harsanyi, J. C. (1968a). Games with Incomplete Information Played by "Bayesian" Players, I-III. Part II. Bayesian Equilibrium Points. *Management Science*, *14*(5), 320–334. https://www.jstor.org/stable/2628673

Harsanyi, J. C. (1968b). Games with Incomplete Information Played by "Bayesian" Players, I-III. Part III. The Basic Probability Distribution of the Game. *Management Science*, *14*(7), 486–502. https://www.jstor.org/stable/2628894

Harsanyi, J. C. (1973). Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, *2*(1), 1–23. https://doi.org/10.1007/bf01737554

Harsanyi, J. C. (2004). Games with Incomplete Information Played by "Bayesian" Players, I-III: Part I. The Basic Model. *Management Science*, *50*(12), 1804–1817. https://www.jstor.org/stable/30046151

Hart, S., & Kurz, M. (1983). Endogenous Formation of Coalitions. *Econometrica*, *51*(4), 1047–1064. https://doi.org/10.2307/1912051

Hart, S., & Mas-Colell, A. (1996). Bargaining and Value. *Econometrica*, *64*(2), 357–380. https://doi.org/10.2307/2171787

Hayek, F. A. (1945). The Use of Knowledge in Society. *The American Economic Review*, *35*(4), 519–530. https://www.jstor.org/stable/1809376

Heidhues, P., Rady, S., & Strack, P. (2015). Strategic experimentation with private payoffs. *Journal of Economic Theory*, *159*, 531–551. https://doi.org/10.1016/j.jet.2015.07.017

Hörner, J., & Skrzypacz, A. (2017, November). Learning, Experimentation, and Information Design. In B. Honoré, A. Pakes, M. Piazzesi, & L. Samuelson (Eds.), *Advances in Economics and Econometrics* (1st ed., pp. 63–98). Cambridge University Press. https://doi.org/10.1017/9781108227162.002

Jonas, H. (2008). *Le principe responsabilité : Une éthique pour la civilisation technologique*. Flammarion.

Jullien, B., & Park, I.-U. (2014). New, like new, or very good? reputation and credibility. *The Review of Economic Studies*, *81*(4), 1543–1574. https://doi.org/10.1093/restud/rdu012

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, *47*(2), 263. https://doi.org/10.2307/1914185

Keller, G., & Rady, S. (2015). Breakdowns: Breakdowns. *Theoretical Economics*, *10*(1), 175–202. https://doi.org/10.3982/TE1670

Keller, G., Rady, S., & Cripps, M. (2005). Strategic Experimentation with Exponential Bandits. *Econometrica*, *73*(1), 39–68. https://doi.org/10.1111/j.1468-0262.2005.00564.x

Klein, N., & Rady, S. (2011). Negatively Correlated Bandits. *The Review of Economic Studies*, *78*(2), 693–732. https://doi.org/10.1093/restud/rdq025

Konishi, H., Le Breton, M., & Weber, S. (1999). On coalition-proof nash equilibria in common agency games. *Journal of Economic Theory*, *85*(1), 122–139.

Konishi, H., & Ray, D. (2003). Coalition formation as a dynamic process. *Journal of Economic Theory*, *110*(1), 1–41.

Kreps, D. M., & Wilson, R. (1982a). Reputation and imperfect information. *Journal of Economic Theory*, *27*(2), 253–279. https://doi.org/10.1016/0022-0531(82)90030-8

Kreps, D. M., & Wilson, R. (1982b). Sequential equilibria. *Econometrica*, *50*(4), 863. https://doi.org/10.2307/1912767

Kuhn, H. W. (1953). Extensive Games and the Problem of Information. In *Contributions to the Theory of Games (AM-28), Volume II* (pp. 193–216). Princeton University Press. https://doi.org/10.1515/9781400881970-012

Kultti, K. K., & Miettinen, P. A. (2007). Herding with Costly Observation. *The B.E. Journal of Theoretical Economics*, *7*(1). https://doi.org/10.2202/1935-1704.1320

Lewis, D. K. (1969). *Convention: A Philosophical Study*. Cambridge: Harvard University Press.

Maciejewski, M. (2018). *Role of the European Parliament in promoting the use of independent expertise in the legislative process* (tech. rep.). Policy Department for Economic, Scientific and Quality of Life Policies, European Parliament. Luxemburg. http://www.europarl.europa.eu/supporting-analyses

Mailath, G. J., & Samuelson, L. (2001). Who wants a good reputation? *The Review of Economic Studies*, *68*(2), 415–441. https://www.jstor.org/stable/2695935

Mailath, G. J., & Samuelson, L. (2006). *Repeated games and reputations: Long-run relationships*. Oxford Univ. Press.

Mandelbaum, A. (1987). Continuous Multi-Armed Bandits and Multiparameter Processes. *The Annals of Probability*, *15*(4), 1527–1556. https://www.jstor.org/stable/2244018

Marlats, C., & Ménager, L. (2021). Strategic observation with exponential bandits. *Journal of Economic Theory*, *193*, 105232. https://doi.org/10.1016/j.jet.2021.105232

Ménager, L. (2023). Common Knowledge in Game Theory. *Revue économique*, *74*(4), 569–600. https://www.jstor.org/stable/48766952

Mercereau. (2025). *La filière EPR : Une dynamique nouvelle, des risques persistants* (tech. rep.). Cour des Comptes. Paris. https://www.ccomptes.fr/fr/publications/la-filiere-epr-une-dynamique-nouvelle-des-risques-persistants

Milgrom. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, *12*(2), 380–391. https://doi.org/10.2307/3003562

Milgrom & Roberts. (1982). Predation, reputation, and entry deterrence. *Journal of Economic Theory*, *27*(2), 280–312. https://doi.org/10.1016/0022-0531(82)90031-X

Milgrom & Roberts. (1986). Price and advertising signals of product quality. *Journal of Political Economy*, *94*(4), 796–821. https://www.jstor.org/stable/1833203

Minozzi, W. (2011). A jamming theory of politics. *The Journal of Politics*, *73*(2), 301–315. https://doi.org/10.1017/S0022381611000296

Misyak, J. B., & Chater, N. (2014). Virtual bargaining: A theory of social decision-making. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1655), 20130487. https://doi.org/10.1098/rstb.2013.0487

Montero, M. (2023). Coalition Formation in Games with Externalities. *Dynamic Games and Applications*, *13*(2), 525–548. https://doi.org/10.1007/s13235-022-00460-0

Moretti, S., & Patrone, F. (2008). Transversality of the Shapley value. *TOP*, *16*(1), 1–41. https://doi.org/10.1007/s11750-008-0044-5

Myerson, R. B. (1999). Nash Equilibrium and the History of Economic Theory. *Journal of Economic Literature*, *37*(3), 1067–1082. https://www.jstor.org/stable/2564872

Nash, J. F. (1950a). The Bargaining Problem. *Econometrica*, *18*(2), 155. https://doi.org/10.2307/1907266

Nash, J. F. (1950b). Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences*, *36*(1), 48–49. https://doi.org/10.1073/pnas.36.1.48

Ottaviani, M., & Sørensen, P. N. (2006a). Professional advice. *Journal of Economic Theory*, *126*(1), 120–142. https://doi.org/10.1016/j.jet.2004.08.005

Ottaviani, M., & Sørensen, P. N. (2006b). Reputational cheap talk. *The RAND Journal of Economics*, *37*(1), 155–175. https://doi.org/10.1111/j.1756-2171.2006.tb00010.x

Park, I.-U. (2005). Cheap-talk referrals of differentiated experts in repeated relationships. *The RAND Journal of Economics*, *36*(2), 391–411. https://www.jstor.org/stable/4135248

Paternotte, C. (2017). The Fragility of Common Knowledge. *Erkenntnis*, *82*(3), 451–472. https://doi.org/10.1007/s10670-016-9828-4

Paternotte, C. (2011). Being realistic about common knowledge: A Lewisian approach. *Synthese*, *183*(2), 249–276. https://doi.org/10.1007/s11229-010-9770-y

Peleg, B., & Sudhölter, P. (2007). *Introduction to the theory of cooperative games* (2nd ed). Springer.

Pérez-Castrillo, D., & Wettstein, D. (2001). Bidding for the surplus: A non-cooperative approach to the shapley value. *Journal of Economic Theory*, *100*(2), 274–294.

Prelec, D. (1998). The Probability Weighting Function. *Econometrica*, *66*(3), 497–527. https://doi.org/10.2307/2998573

Presman, E. L. (1991). Poisson Version of the Two-Armed Bandit Problem with Discounting. *Theory of Probability & Its Applications*, *35*(2), 307–317. https://doi.org/10.1137/1135038

Quiggin, J. (2003). Background Risk in Generalized Expected Utility Theory. *Economic Theory*, *22*(3), 607–611. https://www.jstor.org/stable/25055704

Ray, D., & Vohra, R. (1999). A theory of endogenous coalition structures. *Games and Economic Behavior*, *26*(2), 286–336.

Rosenberg, D., Solan, E., & Vieille, N. (2007). Social Learning in One-Arm Bandit Problems. *Econometrica*, *75*(6), 1591–1611. https://doi.org/10.1111/j.1468-0262.2007.00807.x

Rosenthal, R. W. (1981). Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory*, *25*(1), 92–100. https://doi.org/10.1016/0022-0531(81)90018-1

Rothschild, M. (1974). A two-armed bandit theory of market pricing. *Journal of Economic Theory*, *9*(2), 185–202. https://doi.org/10.1016/0022-0531(74)90066-0

Salamanca, A. (2020). On the values of bayesian cooperative games with sidepayments. *Mathematical Social Sciences*, *108*, 38–49.

Schnakenberg, K. E. (2015). Expert advice to a voting body. *Journal of Economic Theory*, *160*, 102–113. https://doi.org/10.1016/j.jet.2015.08.005

Schnakenberg, K. E. (2017). The downsides of information transmission and voting. *Public Choice*, *173*(1), 43–59. https://doi.org/10.1007/s11127-017-0462-8

Schottmüller, C. (2019). Too good to be truthful: Why competent advisers are fired. *Journal of Economic Theory*, *181*, 333–360. https://doi.org/10.1016/j.jet.2019.03.006

Shapley, L. S. (1953). A value for n-person games. *Annals of Mathematics Studies*, *28*, 307–318.

Shapley, L. S., & Shubik, M. (1969). Pure Competition, Coalitional Power, and Fair Division. *International Economic Review*, *10*(3), 337. https://doi.org/10.2307/2525647

Sobel, J. (1985). A theory of credibility. *The Review of Economic Studies*, *52*(4), 557. https://doi.org/10.2307/2297732

Spence, M. (1973). Job Market Signaling. *The Quarterly Journal of Economics*, *87*(3), 355–374. https://doi.org/10.2307/1882010

Stigler, G. J. (1961). The Economics of Information. *Journal of Political Economy*, *69*(3), 213–225. Retrieved January 14, 2025, from https://www.jstor.org/stable/1829263

Stiglitz, J. E. (1975). The Theory of "Screening," Education, and the Distribution of Income. *The American Economic Review*, *65*(3), 283–300. https://www.jstor.org/stable/1804834

Strulovici, B. H. (2010). Learning While Voting: Determinants of Collective Experimentation. *Econometrica*, *78*(3), 933–971. https://doi.org/10.3982/ECTA8011

Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior* (First edition). Princeton University Press.

Yaari, M. E. (1987). The Dual Theory of Choice under Risk. *Econometrica*, *55*(1), 95–115. https://doi.org/10.2307/1911158

Young, H. P. (1985). Monotonic solutions of cooperative games. *International Journal of Game Theory*, *14*(2), 65–72.

Résumé : ***Politique, jeux et communication stratégique.*** Cette thèse étudie la transmission stratégique d'information dans des jeux de science politique. Elle consiste en trois parties indépendantes. La première partie est consacrée à l'étude de la coopération mal informée qui apparaît dans les projets publics à grande échelle. Je développe un modèle de jeux coalitionnels simples où les joueurs connaissent pour sûr uniquement les valeurs des coalitions auxquelles ils appartiennent. Chaque coalition peut dévoiler sa valeur à un planificateur social si ses membres le décident à l'unanimité. Le planificateur social met ensuite à jour ses croyances à propos du véritable jeu coalitionnel et alloue les paiements en fonction. Je caractérise sous quelles conditions il est optimal de certifier toutes et seulement les coalitions gagnantes. Dans la seconde partie, j'étudie un modèle de communication facile (*cheap talk*) répétée entre des experts et un parlementaire. Les experts ont chacun un biais idéologique et sont inégalement compétents. Ils peuvent être licenciés librement s'ils sont insatisfaisants. Le parlementaire souhaite identifier les experts compétents et qu'ils le conseillent sincèrement. Je montre que la compétition idéologique, c'est-à-dire avoir des experts avec des biais opposés, est bénéfique au Parlement seulement quand (i) la compétence des experts est une information privée et (ii) même les experts compétents peuvent se tromper. La troisième partie est préliminaire. Je définis un modèle d'expérimentation stratégique où des politiciens font face à la possibilité d'une catastrophe naturelle. Ils disposent de trois actions : construction de protection privée, conduite de recherche publique ou attentisme (sans coût). Je résous actuellement le problème de décision en stratégies markoviennes.

Descripteurs : Jeux répétés, Jeux coalitionnels, Communication stratégique, Expérimentation stratégique, Bayésianisme.

Abstract: ***Politics, games and strategic communication.*** This thesis studies the strategic transmission of information in political science games. It consists in three independent parts. The first part is devoted to the study of misinformed cooperation as found in public large-scale projects. I develop a model of simple coalition games where players know for sure only the worth of the coalitions they belong to. Each coalition may disclose its worth to a social planner should its members unanimously agree. The social planner then updates her belief about the coalition game at play and allocates payments accordingly. I characterise under which conditions it is optimal to certify all and only winning coalitions. In the second part, I study a model of repeated cheap talk between experts and a Member of Parliament. The experts have each an ideological bias and are unequally competent. They can be fired freely should they be unsatisfactory. The Member of Parliament wants to identify competent experts and to have them advise sincerely. I prove that ideological competition, namely having experts with opposed biases, is beneficiary to the Parliament only when (i) the competence of experts is private information and (ii) even competent experts may be wrong. The third part is at an early stage. I develop a model of strategic experimentation where politicians face the possibility of a natural disaster. They must choose between three actions: they may build private protection, conduct public research or wait (costless). I currently solve the decision problem in Markovian strategies.

Keywords: Repeated games, Coalitional games, Strategic communication, Strategic experimentation, Bayesianism.