



HAL
open science

Essays in Bayesian Games

Alfonso Montes Sanchez

► **To cite this version:**

Alfonso Montes Sanchez. Essays in Bayesian Games. Methods and statistics. Institut Polytechnique de Paris, 2020. English. NNT : 2020IPPAX096 . tel-03167045

HAL Id: tel-03167045

<https://theses.hal.science/tel-03167045>

Submitted on 11 Mar 2021

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

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



INSTITUT
POLYTECHNIQUE
DE PARIS

NNT : 2020IPPAX096

Thèse de doctorat



Essays in Bayesian Games

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École Polytechnique

École doctorale n°626 École Doctorale de l'Institut Polytechnique de Paris (ED IP
Paris)

Spécialité de doctorat: Sciences économiques

Thèse présentée et soutenue à Palaiseau, le 30/11/2020, par

ALFONSO MONTES

Composition du Jury :

Olivier Gossner Directeur de recherche, CREST-École Polytechnique et CNRS	Directeur de thèse
Ina Taneva Professeur, University of Edinburgh	Rapporteur
Frédéric Koessler Directeur de recherche, Paris School of Economics et CNRS	Rapporteur
Vasiliki Skreta Professeur, UCL and UT Austin	Examineur
Tristan Tomala Directeur de recherche, HEC Paris et CNRS	Président du jury
Yukio Koriyama Professeur, CREST-École Polytechnique	Examineur

Acknowledgements

The road I have rode during my PhD has been far from easy. At its beginnings, I was living in the UK, and my wife Mariana was seven months pregnant of our daughter Amparo. We moved to France with Amparo being one-month old, and we initiated then a long and tough path that started to show its payoffs not very long ago. This ride has been anything but a lonely one. Several people have been there for me, to support me in both personal and professional issues. These following lines constitute an attempt to acknowledge the effort of those who help me succeed in the gigantic task of pursuing a PhD.

First I want to thank Olivier Gossner. During this whole process, he has been there not only to teach me about microeconomics and game theory, but also about how to become a better professional and academic. I thank him all the time he spent reading my drafts, answering my questions and giving me feedback. I also thank his support in several other issues that go beyond producing research, and for being always understanding about the difficulties that I faced as a foreign student. Without his efforts, this dissertation could have not been completed.

Second, I want to thank my wife Mariana Huepe, who, despite the many difficulties we encountered during our first years living in France, she stood up and hanged in here with me to help me to complete this process. I also thank my adorable daughter Amparo, for being a constant source of love and happiness to us. I thank also my father Alejandro Montes, who has always trusted me and has always been there to support me and my family in all matters.

Third, I want to thank to the people at IPP that were always there to help, share and laugh. In particular, I thank Weronika Leduc and Lyza Racon for always being supportive and efficient when it has came to help me out. Weronika and Lyza made my life infinitely easier, not only administratively, but also personally. I also thank Tristan Tomala, Yannik Lucotte, Wael Bousselmi, Antoine Ferey, Matías Nuñez, Rafael Vieil, Francis Kramatz and Thibaud Vergé for all their support. I also thank École Polytechnique for embracing me and my family, and for providing support in all aspects of the student life.

Finally, I want to thank all the members of my jury, Ina Taneva, Frédéric Koessler, Vasiliki Skreta, Tristan Tomala and Yukio Koriyama. Despite all the imperfections of this PhD thesis, they dived into it with the clear purpose of helping me to grow as a researcher. I feel very proud of having this remarkable ensemble of people as the jury of my PhD.

Contents

Acknowledgements	3
1 Résumé en Français	7
1.1 Environnements Stratégiques	7
1.2 Acquisition d'Informations Coûteuse	8
1.3 Perception Erronée des Informations	9
1.4 Résumé du Chapitre 3	9
1.5 Résumé du Chapitre 4	11
1.6 Résumé du Chapitre 5	13
2 Introduction in English	15
2.1 Strategic Environments	15
2.2 Costly Information Acquisition	16
2.3 Misperception of Information	16
2.4 Summary of Chapter 3	17
2.5 Summary of Chapter 4	19
2.6 Summary of Chapter 5	20
3 Bayesian Persuasion with Rational Inattentive Receiver	23
3.1 Introduction	23
3.2 The Model	25
3.3 Solving the Model	27
3.3.1 The receiver's problem	27
3.3.1.1 The sender's problem	28
3.3.1.2 A Simple Example	29
3.3.2 Discussion	30
3.4 Comparative Statics	30
3.4.1 Preamble	31
3.4.2 Payoff Monotonicity	31
3.4.3 The Sender	32
3.4.4 The Receiver	32
3.4.5 A positive Result	32
3.5 An Application: Informational Lobby	33
3.5.1 Basic Setup	33
3.5.2 Solution and Comparative Statics	34
3.6 Conclusion	35
4 Canonical Equilibria in Games with Flexible Information Acquisition	37

4.1	Introduction	37
4.2	The Model	40
4.2.1	The Bayesian Game	40
4.2.2	The Attention-Move Game	40
4.2.2.1	Payoffs	40
4.2.2.2	The Cost of Information	41
4.2.2.3	Equilibrium	42
4.3	Existence of Canonical Nash Equilibrium	42
4.3.1	Main Results	42
4.3.2	Properties of Equilibrium	43
4.3.3	Proof of Theorem 4.1	44
4.3.4	Proof of Theorem 4.2	45
4.4	Rational Inattention in Strategic Environments	45
4.4.1	Rational Inattention and Behavioral Implications	45
4.4.2	Main Results	46
4.4.3	Concavification Technique to Find a Solution	47
4.5	Examples	48
4.5.1	Example 4.5.1: Matching Pennies	48
4.5.2	Example 4.5.2: Matching Pennies with Uncertain Outside Option	49
4.5.3	Example 4.5.3: Matching Pennies with Costly Number of Signals	50
4.5.4	Example 4.5.4: Battle of Sexes	51
4.5.5	Example 4.5.5: Chicken Game.	52
4.6	Conclusion	53
4.7	Appendix	53
4.7.1	Proof of Lemma 4.1	53
4.7.2	Proof of Lemma 4.2	54
4.7.3	Proof of Lemma 4.3	55
4.7.4	Proof of Lemma 4.4	56
5	Ideology, Media Outlets and Political Extremism	59
5.1	Introduction	59
5.2	The Model	62
5.2.1	Basics	62
5.2.2	Information	62
5.2.3	Ideological Neglect	62
5.2.4	Strategies and Equilibrium	63
5.3	Solving the Model	63
5.3.1	Preliminaries	63
5.3.2	Linear Equilibrium	64
5.3.3	Comparative Statics	65
5.4	Political Extremism	66
5.5	Welfare Analysis	67
5.6	Conclusion	68
5.7	Appendix	70
5.7.1	Proof of Lemma 5.1	70

5.7.2	Proof of Proposition 5.1	70
5.7.3	Proof of Proposition 5.2	71
5.7.4	Proof of Proposition 5.3	73
5.7.5	Proof of Proposition 5.4	73
5.7.6	Proof of Lemma 5.2	74
5.7.7	Proof of Lemma 5.2	76
5.7.8	Proof of Proposition 5.5	76
5.7.9	Proof of Proposition 5.6	77
5.7.10	Proof of Proposition 5.7	78
5.7.11	Proof of Lemma 5.3	79
5.7.12	Proof of Lemma 5.3	81
5.7.13	Proof of Proposition 5.8	81
5.7.14	Proof of Proposition 5.9	81
5.7.15	Proof of Proposition 5.10	82
5.7.16	Proof of Proposition 5.11	83

Chapter 1

Résumé en Français

Cette thèse de doctorat propose une série d'essais sur les jeux bayésiens. Il concentre son attention sur le rôle de l'information dans les résultats de ces jeux, depuis les stratégies et l'équilibre jusqu'aux implications pour le bien-être. De manière générale, les jeux bayésiens sont des situations stratégiques dans lesquelles il existe des informations incomplètes par rapport à un état pertinent lors de la détermination des paiements. Les joueurs ont une croyance a priori sur le paramètre inconnu, et ils reçoivent des informations sur l'état avant de choisir leurs actions. Les joueurs sont appelés *Bayesian* car ils mettent à jour leurs informations en utilisant la règle de Bayes. Traditionnellement, les informations reçues par les joueurs sont supposées exogènes et correctement comprises. Cette thèse de doctorat explore les implications et les conséquences de l'abandon de ces hypothèses.

Les trois principaux chapitres de cette thèse abordent deux questions différentes liées à l'information. D'une part, il n'est plus supposé que les joueurs reçoivent des informations données de manière exogène sur l'état inconnu, mais qu'ils doivent les collecter tout en encourageant un coût. Cette hypothèse conduit les contributions et les résultats du chapitre 3 et 4. Par contre, dans le chapitre 5 il n'est plus supposé que les informations distribuées entre les joueurs sont bien comprises. Autrement dit, je considère que les joueurs reçoivent des informations biaisée et corrélées, et qu'ils n'en sont pas pleinement conscients.

1.1 Environnements Stratégiques

Cette thèse étudie trois cadres stratégiques différents. Premièrement, il étudie la persuasion bayésienne. Dans ce cadre, un désigner d'informations (l'expéditeur) choisit les informations à divulguer afin de *persuader* un décideur (le destinataire) de choisir une action particulière. Il existe un état de nature inconnu qui, avec l'action choisie par le receveur, détermine le gain que les deux joueurs reçoivent. La divulgation d'informations est donc stratégique. Contrairement aux jeux bayésiens standard, le récepteur est le seul joueur qui choisit une action, et l'expéditeur a la possibilité de concevoir librement les informations à divulguer. De plus, les signaux (et les stratégies d'information de l'expéditeur) sont observés publiquement, tandis que dans les jeux bayésiens, les signaux sont observés en privé par les joueurs. Dans ce contexte, je suppose que le récepteur est capable de recueillir des informations supplémentaires après avoir observé la réalisation du signal du récepteur. L'acquisition d'informations est supposée coûteuse pour le récepteur et il augmente selon la quantité d'informations qu'elle fournit.

Deuxièmement, il étudie les jeux bayésiens généraux (finis). Un jeu bayésien est composé d'un *jeu de base* (joueurs, états de la nature, croyance a priori et gains) et d'une *structure d'information* (les

informations reçues par les joueurs).¹ Les informations reçues par chaque joueur sont définies par une *expérimentation statistique*, qui consiste en un espace de réalisation (les signaux) et une famille de distributions de probabilités qui spécifie la probabilité d'observer chaque signal dans chaque état de nature particulier. Après avoir observé en privé leurs signaux, les joueurs mettent à jour leurs croyances sur l'état (en utilisant la règle de Bayes), puis choisissent des actions simultanément. Dans ce cadre, je suppose que les joueurs choisissent leurs expérimentations statistiques, et paient un coût qui augmente avec la quantité d'informations fournies par leurs signaux.

Enfin, il étudie un jeu particulier d'informations incomplètes qui transmet quelques aperçus théoriques dans le domaine de l'économie politique. Dans ce cas, je suppose une structure de paiement particulière pour les joueurs. Pour chaque joueur, la perte est donnée par la distance quadratique entre son action et l'état de nature, et la distance quadratique entre son action et l'action moyenne dans la population. Le premier composant du gain représente l'intérêt que les agents ont pour *faire la bonne chose*, car ils veulent affirmer l'état de nature réalisé. Le deuxième terme représente une préférence pour l'accord entre les individus au sein de la société. Dans ce contexte, je suppose que les joueurs perçoivent mal les informations qu'ils reçoivent.

1.2 Acquisition d'Informations Coûteuse

Au cœur des résultats des chapitres 3 et 4 se trouve l'hypothèse que l'information est endogène et coûteuse. Les développements récents de la théorie et des applications économiques ont embrassé l'idée que les individus ne tirent pas leurs informations de nulle part, mais qu'ils doivent consacrer du temps et des efforts à les collecter afin de prendre de meilleures décisions.

Par exemple, les acheteurs sur Internet doivent faire des recherches en ligne lorsqu'ils ne connaissent pas certaines qualités d'un produit particulier. Les gestionnaires, les jurés et les décideurs politiques sont également confrontés à des informations incomplètes lorsqu'ils doivent choisir leurs actions. Ils doivent lire des journaux, des rapports, des sondages; ils doivent parler à des experts, étudier des sujets pertinents et des développements récents, et donner un sens à toutes les informations qu'ils ont obtenues. Par conséquent, obtenir des informations est coûteux. Et les mêmes prémisses s'appliquent aux situations stratégiques. On peut penser aux entreprises qui décident d'entrer dans un marché oligopolistique, aux enchérisseurs dans une vente aux enchères EBay et aux électeurs qui choisissent leurs candidats aux élections générales.

L'acquisition d'informations coûteuses par un seul agent a été mise en avant par la littérature sur l'inattention rationnelle.² Les agents sont *inattentifs rationnels* lorsque le coût d'acquisition de l'information est linéaire dans la *information mutuelle* entre les signaux et l'état.³ Lorsque la connaissance de l'environnement est coûteuse et que le coût de l'information satisfait certaines propriétés, les décideurs afficheront certains modèles de comportement. Une implication comportementale importante est que les informations vont être conçues dans un but de prise de décision, en ce sens que les informations qui ne sont pas pertinentes pour passer d'une action à une autre seront ignorées. Cela implique que chaque signal sera associé de manière unique à une action et que les stratégies d'équilibre prendront la forme de probabilités de choix conditionnel.

Cette idée a été étendue à des environnements stratégiques tels que les jeux globales (*global games*), les négociations (*bargaining*) et les jeux potentiels (*potential games*), et elle constitue un domaine de

¹ Voir Gossner (2002, [38]).

² Les articles phares de cette littérature sont Sims (2003, [77]), Matejka et McKay (2015, [62]), Stainer et al. (2018, [79]).

³ Voir chapitre 3.

recherche en croissance dont les limites sont encore inconnues.⁴ Comme il sera discuté, ceci la thèse reprend cette idée et l'utilise pour étudier les jeux bayésiens *generales* et une configuration de persuasion bayésienne avec un récepteur rationnel inattentif.

1.3 Perception Erronée des Informations

Au cœur du chapitre 5 de cette thèse se trouve l'idée que les agents peuvent ne pas percevoir correctement les informations qu'ils reçoivent. Une vaste littérature dans le domaine de l'économie et de la psychologie a exploré et documenté les fausses perceptions individuelles de la réalité, telles que l'excès de confiance, les vœux pieux, la dissonance cognitive et l'auto-tromperie (Benabou 2008, [14]). L'idée que les attitudes des agents reflètent souvent des croyances déformées est courante en science politique et en psychologie politique.⁵

L'idée de la négligence corrélationnelle comme justification d'une mauvaise perception de l'information a été particulièrement populaire.⁶ Les structures d'information génèrent des signaux corrélés plutôt qu'indépendants les uns des autres, et les agents peuvent ne parviennent pas à reconnaître une telle corrélation. La corrélation peut être entre les signaux de différents agents ou dans (la séquence de) signaux reçus par un seul agent. Sous ses nombreuses formes, la négligence corrélationnelle peut entraîner un excès de confiance, l'extrémisme et la polarisation, et ces liens ont été explorés théoriquement et empiriquement.⁷

Les efforts consacrés à cette ligne sont particulièrement précieux aujourd'hui. Notre monde évolue très rapidement et des motifs épineux émergent dans le monde entier. Bien que nous vivions à une époque où l'information est devenue plus accessible que jamais, nous observons les électeurs et les politiciens du monde entier devenir de plus en plus polarisés (McCarty et al. 2006, [65]). De plus, à l'ère moderne, la désinformation, les fausses nouvelles, l'ignorance et les idéologies se mélangent pour permettre à des politiciens radicaux (basés sur le discours de haine) d'accéder au pouvoir, comme Donald Trump aux États-Unis, Jair Bolsonaro au Brésil, Viktor Orban en Hongrie, et Andrzej Duda en Pologne. Angela Merkel, la chancelière allemande, considère cette évolution comme une menace pour les démocraties libérales en général et pour l'UE en particulier.⁸

L'humble contribution du chapitre 5 dans cette thèse contribue avec un grain de sable à l'effort de compréhension de ce phénomène global.

1.4 Résumé du Chapitre 3

Une caractéristique intéressante du modèle de persuasion bayésien standard est qu'il peut être appliqué dans de nombreuses situations économiques.⁹ On pourrait penser dans un groupe de lobbying (l'expéditeur) commander une recherche dans le but d'influencer un politicien (le récepteur). Ou une entreprise essayant de faire une vente en fournissant des informations sur les produits à ses clients potentiels. Néanmoins, pour étudier ces situations avec la configuration de base de la persuasion, nous

⁴ Voir Yang (2015, [83]), Ravid (2018, [74]), Denti (2015, [29]), Matejka (2016, [61]), Martin (2015, [58]).

⁵ Voir Lane (1959, [51]), Hochschild (1981 [42]; 1996 [43]), Kuran (1997, [50]), Jost et Major (2001, [44]).

⁶ Voir De Marzo et al. (2003, [28]), Ortoleva et Snowberg (2015, [69]), Glaeser et Sunstein (2009, [36]) et Levy et Razin (2015, [55]).

⁷ Voir Ortoleva et Snowberg (2015, [69]) et Levy et Razin (2015, [55]).

⁸ Voir l'interview <https://www.theguardian.com/world/2020/jun/26/for-europe-survivre-économie-besoins-survivre-angela-merkel-interview-in-full>.

⁹ Voir Kamenica et Gentzkow (2011, [31]).

devons supposer que le récepteur n'a accès à aucune information supplémentaire. En réalité, cependant, les politiciens peuvent mener leurs propres recherches et les clients peuvent trouver des informations en ligne sur les caractéristiques des produits. Mais les informations doivent être collectées et traitées, ce qui est sans aucun doute coûteux.

Ce chapitre étend la configuration de base de la persuasion bayésienne en permettant au destinataire, après avoir observé les informations divulguées par l'expéditeur, de recueillir des informations supplémentaires. Plus précisément, après avoir observé les informations fournies par l'expéditeur, le récepteur met à jour ses croyances et décide de recueillir des informations supplémentaires. L'acquisition d'informations est coûteuse pour le récepteur, et son coût est supposé linéaire dans *mutual information*, donc elle est *rationnelle inattentive*.¹⁰ Le coût est mis à l'échelle par un paramètre de coût exogène, que je considère comme la capacité du récepteur à collecter des informations.

J'étudie l'impact sur le bien-être des changements dans la capacité du récepteur à recueillir des informations. D'une part, je montre qu'une plus grande capacité de collecte d'informations n'est pas nécessairement meilleure pour le récepteur. En effet, une plus grande capacité à collecter des informations peut conduire l'expéditeur à fournir moins d'informations sur l'état, et moins d'informations globales divulguées à l'équilibre, ce qui ne peut être meilleur pour le décideur. Comme pour les informations elles-mêmes, la capacité de recueillir des informations n'est pas toujours souhaitable dans les environnements stratégiques. Dans un jeu, les gains d'équilibre peuvent être plus faibles lorsque plus d'informations sont fournies aux joueurs.¹¹ De même dans ce chapitre, même si une plus grande capacité de collecte d'informations est toujours souhaitable dans un environnement non stratégique, elle peut être préjudiciable au bien-être des joueurs lorsque les décisions sont prises de manière stratégique.

Je montre également que l'expéditeur ne peut pas être mieux loti si le destinataire devient meilleur dans la collecte d'informations. Autrement dit, une plus grande capacité de collecte d'informations réduit l'ensemble des stratégies réalisables de l'expéditeur, de sorte qu'il ne peut pas être meilleur qu'avant.

Pour résoudre le modèle, je profite des méthodes déjà disponibles dans la littérature. En effet, l'article examine la division optimale des postérieurs pour résoudre le problème de l'expéditeur ainsi que celui du récepteur.¹² Cette méthode, initialement proposée par Aumann et Maschler (1967, [4]; 1968, [5]; 1995, [9]), permet à l'analyste d'utiliser des outils graphiques pour résoudre toute application particulière du modèle.

Pour faciliter encore plus la recherche de l'équilibre, j'utilise le fait que pour chaque équilibre dans lequel le récepteur décide de recueillir des informations supplémentaires, il existe un autre équilibre dans lequel le récepteur ne le fait pas. Autrement dit, on peut toujours rechercher un équilibre dans lequel le récepteur n'apprend pas. La raison est simple: si le récepteur apprend à un certain équilibre, l'expéditeur pourrait donner ces informations supplémentaires en premier lieu sans changer son gain attendu (parce que le récepteur apprend de toute façon). Notez cependant que le gain du destinataire peut changer, car il n'encourt plus les frais de collecte d'informations. Cette observation repose sur l'hypothèse que l'acquisition d'informations est gratuite pour l'expéditeur.

Littérature connexe. Cet article s'inscrit dans la littérature de la persuasion bayésienne. Les extensions du modèle standard vont de plusieurs expéditeurs (Kamenica et Gentzkow 2017, [33]) à un expéditeur informé en privé (Hedlund 2017, [41]; Kolotilin et al. 2017, [49]), récepteur informé en privé (Kolotilin et

¹⁰Voir par exemple Sims (2003, [77]; 2006, [78]), Matejka et McKay (2014, [62]), Caplin et Dean (2013, [21]) et Steiner et al. (2018, [79]).

¹¹Voir par exemple Bassan (2003, [13]), Gossner (2000, [38]) et Neyman (1991, [68])

¹²Voir Kamenica et Gentzkow (2011, [31]) et Caplin et Dean (2013, [21]). Pour utiliser la méthode sur le problème du récepteur, il n'est pas nécessaire de supposer des fonctions de coût basées sur l'entropie. Il suffit de regarder les fonctions de coût séparables postérieures, pour lesquelles les fonctions de coût basées sur l'entropie sont un cas particulier.

al. 2017, [49]), capacité de communication limitée (Treust et Tomala 2019, [52]) et priors hétérogènes (Alonso et Camara 2016, [2]), pour nommer quelques. Kamenika et Gentzkow (2014, [32]) introduisent une acquisition d'information coûteuse dans le modèle en supposant que l'expéditeur doit assumer un coût de divulgation d'informations au destinataire. Les articles les plus proches de celui-ci sont Bloedel et Segal (2018, [19]) et Matysková (2018, [64]). Dans Bloedel et Segal (2018, [19]), le récepteur observe la stratégie d'information de l'expéditeur et choisit, avant d'observer la réalisation du signal, une stratégie d'information coûteuse pour connaître les signaux de l'expéditeur. Le cadre de Matysková (2018, [64]) est analogue à cet article. Contrairement à Matysková (2018, [64]), cet article se concentre sur la technique de la double concavification pour résoudre des instances du modèle et fournit une application simple à la littérature du lobby informationnel en économie politique.

Cet article est actuellement un travail conjoint avec Ludmila Matysková, et le travail conjoint est en cours de révision et de re-soumission au *Journal of Economic Theory*. Ce qui s'est passé avec cet article, c'est que Ludmila et moi travaillions en parallèle sur le même modèle, et nous avons tous les deux soumis nos articles à la revue. L'éditeur nous a proposé de travailler sur une version conjointe de l'article et de le soumettre à nouveau lorsqu'il sera prêt.

1.5 Résumé du Chapitre 4

De nombreuses situations économiques impliquent une prise de décision stratégique dans l'incertitude. Dans de tels cas, les agents décident non seulement de l'action à entreprendre, mais aussi de la quantité à apprendre sur l'environnement. Par exemple, les entreprises doivent se renseigner sur les conditions du marché avant de pénétrer un nouveau marché. Ils peuvent obtenir des informations sur la demande des consommateurs, leur (propre) coût de production et les technologies utilisées par les autres entreprises du marché. Les soumissionnaires sur EBay peuvent collecter des informations sur le produit avant de faire leur offre. Ils peuvent en apprendre davantage sur la durabilité du produit, sa maniabilité et les services après-vente. Enfin, les électeurs lors d'une élection peuvent vouloir se renseigner sur les politiques proposées par différents politiciens avant de se rendre aux urnes.

Ce chapitre étudie l'acquisition d'information coûteuse par plusieurs décideurs confrontés à un environnement stratégique. Les joueurs choisissent un *signal* arbitraire sur un état de nature inconnu, et ils encourrent un coût qui augmente dans la réduction de leur incertitude. Une fois les informations choisies, les agents jouent à une partie bayésienne.

L'étude de jeux d'informations incomplètes avec une acquisition flexible d'informations présente plusieurs difficultés techniques. Premièrement, l'espace de réalisation des signaux peut être n'importe quoi, d'un alphabet fini à un espace de dimension infinie non séparable. Cela seul pose un problème d'existence car l'espace des distributions de probabilités sur l'espace de réalisation est potentiellement non compact. La deuxième difficulté technique est liée au calcul de l'équilibre. En particulier, pour qu'un équilibre existe, il peut être nécessaire pour les joueurs de randomiser sur des actions (après l'observation d'un signal) ou sur des expériences. Par conséquent, le problème de trouver l'équilibre peut devenir une tâche ardue.

La contribution de l'article est triple. Premièrement, je montre que pour chaque équilibre dans (une classe assez grande de) jeux avec acquisition flexible et coûteuse d'informations, il existe un équilibre de stratégie pure *équivalent* dans lequel l'espace des signaux est égal à l'espace d'action.¹³ Cet équilibre de stratégie pure n'implique aucune randomisation sur les actions ni sur les expériences, et je l'appelle *équilibre canonique de Nash*. Deuxièmement, je montre qu'un équilibre canonique existe toujours. Enfin,

¹³Deux Les équilibres de Nash sont *équivalents* si la distribution conjointe entre les actions et les états est la même.

j'étend certains résultats standards de l'acquisition d'informations par un seul agent aux environnements stratégiques.

L'existence d'équilibres canoniques dans les jeux avec acquisition flexible d'informations est intéressante pour au moins trois raisons. Premièrement, il fournit des outils simplificateurs pour analyser des situations qui, en principe, pourraient être très difficiles à aborder. L'équilibre canonique est un équilibre où les signaux recommandent des actions *pure* et ces actions sont suivies par les joueurs. Par conséquent, il fournit un *principe de révélation* pour les jeux avec une acquisition d'informations flexible.

Deuxièmement, l'équilibre canonique fournit une configuration simple pour étendre les résultats standards de l'extraction d'informations par un seul agent aux environnements stratégiques. En utilisant *information mutuelle* comme fonction de coût, cet article fournit des choix d'équilibre (probabilistes) et des stratégies d'apprentissage. Il étend également la *approche postérieure* de Caplin et Dean (2013, [21]), qui fournit des outils géométriques pour résoudre des instances du modèle. Enfin, les équilibres canoniques abordent la critique de l'équilibre des stratégies mixtes.¹⁴ En gros, cette critique dit que l'équilibre de stratégie mixte ne joue aucun rôle en théorie comportementale (au-delà de la garantie de l'existence de l'équilibre), car les gens ne font pas de randomisation en pratique. L'équilibre canonique de Nash est un équilibre de stratégie pure, à la fois dans *stratégies d'information* et dans *stratégies comportementales*.

Le document se termine par quelques exemples. Il discute des pièces de monnaie correspondantes (un jeu classique avec un équilibre unique dans les stratégies mixtes), et une modification de celui-ci dans laquelle les gains dépendent d'un état de nature inconnu. Il présente aussi le jeu de la bataille des sexes et le jeu du poulet, pour discuter de l'équilibre corrélé.¹⁵

Littérature connexe. De manière générale, cet article concerne trois domaines de recherche en économie. Premièrement, il s'inscrit dans la littérature sur l'inattention rationnelle, en particulier dans l'étude des jeux avec des joueurs rationnels inattentifs. Contrairement à ce qui a été fait, dans cet article nous ne regardons pas un environnement stratégique *spécifique*, mais des jeux finis arbitraires. De plus, pour les deux premières contributions principales de l'article (Theorem 4.1 et 4.2), je n'assume pas une forme fonctionnelle particulière de la fonction de coût, mais je garde quatre propriétés clés qui sont également satisfaites par les fonctions de coût basées sur l'entropie. Deuxièmement, cet article s'inscrit dans la littérature sur le design de l'information. Comme noté dans Bergemann et Morris (2019, [17]), les problèmes de design de l'information peuvent être *literal* ou *metaphorical*. Dans le premier, un véritable designer d'informations fournit des informations aux joueurs afin de servir son intérêt personnel. Dans ce dernier, la formulation comme problème de conception de l'information est une métaphore que l'analyste utilise comme outil. Cet article concerne la seconde interprétation de la conception de l'information en étudiant des jeux dans lesquels la structure de l'information résulte d'une interaction indépendante entre les joueurs. J'aborde la question de savoir si nous pouvons restreindre l'attention, sans perte de généralité, aux structures d'information dans lesquelles le signal de chaque joueur *recommande* une action et cette action est suivie. Autrement dit, si nous avons un *principe de révélation* garantissant aucune perte de généralité en regardant les structures d'information *direct*. Enfin, cet article concerne la littérature sur la purification. Lancé par Harsanyi (1973, [40]), ce domaine étudie les conditions dans lesquelles un équilibre de stratégie pure peut être obtenu à partir d'un équilibre de stratégie mixte dans le contexte de jeux d'informations incomplètes. Ce type de résultats est appelé *théorèmes de purification*. Cette littérature se concentre sur les structures d'information exogènes. Puisque l'équilibre canonique

¹⁴Pour une discussion et des réponses à une telle critique, voir par exemple Harsanyi (1973, [40]), Aumann et al. (1983, [8]), Radner et Rosenthal (1982, [72]), Milgrom et Weber (1985, [66]).

¹⁵Voir Auman (1987, [7]).

est un équilibre de stratégie pure, cet article est le premier à fournir un résultat de purification pour les jeux avec des informations endogènes.

1.6 Résumé du Chapitre 5

Il est difficile d'expliquer la montée de l'extrémisme politique sans une notion de *idéologie*.¹⁶ Malgré son importance, ni les économistes ni les philosophes politiques ne sont parvenus à s'accorder avec une définition unifiée du concept. Sous ses nombreuses formes, le terme est utilisé pour désigner les préférences, les opinions, les croyances et certains types de biais psychologiques. Le philosophe politique Terry Eagleton dit que l'idéologie est comme l'halitosis, car vous pouvez facilement la reconnaître chez les autres, mais vous pouvez difficilement la reconnaître en vous-même.¹⁷

Cet article propose un cadre dans lequel l'idéologie est considérée comme une perception erronée de la réalité que les agents reconnaissent chez les autres mais ne parviennent pas à se reconnaître en eux-mêmes. Plus précisément, deux groupes idéologiques obtiennent leurs informations privées de sources d'information biaisées. Je suppose que les agents du même groupe consomment des informations provenant de médias similaires, de sorte que les *signaux* qu'ils reçoivent sont corrélés. Les agents souffrent de *négligence de biais*, dans laquelle ils ne parviennent pas à reconnaître le biais dans leurs sources d'information. De plus, ils souffrent de *négligence corrélacionnelle*, où ils ne parviennent pas à reconnaître la corrélation entre les signaux au sein du groupe idéologique. Néanmoins, les agents évaluent correctement le biais et la corrélation dans les sources d'information de l'autre groupe, et ils reconnaissent que l'autre groupe ne les reconnaît pas.¹⁸ Les agents reçoivent également des informations d'une source impartiale accessible au public.

Les individus doivent choisir leur position politique sur une question (par exemple, les taux d'imposition, les dépenses publiques, etc.), et leurs bénéfices dépendent de la réalisation d'un état inconnu (par exemple, les taux d'intérêt, le chômage, etc.). Leurs bénéfices dépendent également du niveau d'accord avec le reste de la population, il existe donc une complémentarité stratégique entre les positions politiques individuelles. Dans ce contexte, j'étudie le rôle du biais et de la négligence corrélacionnelle sur l'extrémisme politique et le bien-être social.

Je montre que la précision de l'information publique a un effet ambigu sur l'extrémisme politique. Un facteur clé de cette relation est le degré de négligence corrélacionnelle. Si la négligence corrélacionnelle est faible et si l'information publique est bruyante, l'augmentation de la variance de l'information publique peut réduire l'extrémisme politique. La raison en est que les agents accorderont plus de poids à leurs signaux privés, motivés non seulement par son accroissement relatif de précision, mais aussi par l'augmentation de son pouvoir de coordination. Ces deux effets suffisent à réduire l'extrémisme politique, malgré l'augmentation de la variance du signal public. Si la négligence corrélacionnelle est importante, l'extrémisme augmente toujours dans la variance de l'information publique.

Le bruit dans les signaux privés est composé d'un choc idiosyncratique de moyenne nulle (indépendant de tous les agents) et d'un choc commun à tous les membres d'un même groupe idéologique. Comme pour l'information publique, la variance du bruit idiosyncratique et de groupe a un effet ambigu sur l'extrémisme politique. Concernant le bien-être, je montre que le bien-être social attendu augmente dans la précision du signal public. Je montre en outre que le bien-être attendu diminue dans la taille du biais

¹⁶Voir Ortleva et Snowberg (2015, [69]), Martin et Yurukoglu (2017, [60]), Gentzkow et Shapiro (2006, [34]), Prior (2013, [71]).

¹⁷Terry Eagleton (1990, [81]).

¹⁸Ils croient également que les agents de l'autre groupe idéologique croient à tort qu'ils ne parviennent pas à les reconnaître. reconnaître leurs propres préjugés.

des sources d'information, et que diminue le biais et la négligence corrélacionnelle.

Littérature connexe. Cet article apporte une modeste contribution théorique à la littérature sur l'économie politique comportementale. D'une part, cet article contribue à l'étude de la négligence corrélacionnelle. Les articles importants de cette littérature récente sont Ortoleva et Snowberg (2015, [69]) et Levy et Razin (2015, [55]), qui étudient l'effet de la négligence corrélacionnelle dans le comportement politique. Dans le premier cas, il n'y a pas d'interaction stratégique et les agents ne parviennent pas à reconnaître que les signaux *qu'ils* observent sont corrélés les uns aux autres. Cela implique que les agents évaluent à tort le caractère informatif de leurs signaux. Ils montrent que la négligence corrélacionnelle génère un excès de confiance et documentent son impact sur les préférences politiques, la participation électorale et l'identification partisane. Levy et Razin (2015, [55]) examinent un environnement stratégique dans lequel les deux signaux *chaque* joueur observés peuvent être corrélés. Ils montrent que *les électeurs comportementaux*, ceux qui perçoivent mal la corrélation, peuvent finir par faire de meilleurs choix que ceux qui sont *rationnels*, car ils peuvent *correctement* changer son comportement après avoir observé deux signaux apparemment indépendants. Dans cet article, nous examinons également un modèle stratégique, mais j'introduis la négligence corrélacionnelle comme une incapacité à reconnaître la corrélation entre les signaux au sein d'un groupe idéologique. J'introduis également la négligence de partialité, car je considère que les sources d'information sont biaisées.

Chapter 2

Introduction in English

This PhD dissertation provides a series of essays in Bayesian games. It centers its attention on the role of information in the outcomes of such games, from strategies and equilibrium to welfare implications. Broadly speaking, Bayesian games are strategic situations in which there is incomplete information regarding a payoff relevant state. Players have a prior belief about the unknown parameter, and they receive information about the state before choosing their actions. Players are called *Bayesian* because they update their information using Bayes rule. Traditionally, the information players receive is assumed to be exogenous and correctly understood. This PhD dissertation explores some implications and consequences of giving up these assumptions.

The three core chapters of this thesis tackle two different issues related to information. On the one hand, it is no longer assumed that players receive exogenously given information about the unknown state, but that they have to gather it while incurring a cost. This assumption drives the contributions and results of chapter 3 and 4. On the other hand, in chapter 5 it is no longer assumed that information distributed among the players is well understood. That is, I consider that players receive bias and correlated information and that they are not fully aware of it.

2.1 Strategic Environments

This dissertation studies three different strategic frameworks. First, it studies the Bayesian Persuasion setup. In this framework, an information designer (the sender) chooses what information to disclose in order to *persuade* a decision maker (the receiver) to choose a particular action. There is an unknown state of nature which, along with the action chosen by the receiver, determines the payoff that both players receive. Information disclosure is then strategic. As opposed to standard Bayesian games, the receiver is the only player that chooses an action, and the sender has the ability to freely design the information to be disclosed. Furthermore, signals (and sender's information strategies) are publicly observed, while in Bayesian games signals are privately observed by the players. In this context, I assume that the receiver is able to gather additional information after observing the realization of receiver's signal. Information acquisition is assumed to be costly for the receiver, and increasing in the amount of information it obtains.

Second, it studies general (finite) Bayesian games. A Bayesian game is composed by a *basic game* (players, states of nature, payoffs and prior beliefs) and an *information structure* (the information received by the players).¹ Information received by each player is defined by a *statistical experiment*, which consists on a realization space (the signals) and a family of probability distributions that specifies the probability of

¹See Gossner (2002, [38]).

observing each signal in every particular state of nature. After privately observing their signals, players update their beliefs about the state (using Bayes rule) and then choose actions simultaneously. Within this framework, I assume that players choose their statistical experiments, and pay a cost that is increasing in the amount of information it provides.

Finally, it studies a particular game of incomplete information that conveys some theoretical insights in the field of Political Economy. In this case, I assume a particular payoff structure for the players. For every individual, the loss is given by the quadratic distance between his action and the state of nature, and the quadratic distance between his action and the average action in the population. The first component of the payoff represents the interest agents have in *doing the right thing*, as they want to assert the realized state of nature. The second term represents a preference for agreement among individuals. In this context, I assume that players misperceive the information they receive.

2.2 Costly Information Acquisition

At the core of the results in chapters 3 and 4 is the assumption that information is endogenous and costly. Recent developments in economic theory and applications have embraced the idea that individuals do not get their information out of nowhere, but that they have to spend time and effort collecting it in order to make better decisions.

For instance, buyers on the internet have to do online research when they do not know about certain qualities of a particular product. Managers, jurors and policy makers also face incomplete information when they have to choose their actions. They have to read newspapers, reports, surveys; they have to talk to experts, study relevant topics and recent developments, and to make sense of all the information they have obtained. Hence, obtaining information is costly. And the same premises apply for strategic situations. One can think of firms deciding to enter in an oligopolistic market, bidders in an EBay auction, and voters choosing their candidates in general elections.

Single agent costly information acquisition has been brought to the table by the literature on rational inattention. Landmark papers in this literature are Sims (2003, [77]), Matejka and McKay (2015, [62]), Stainer et al. (2018, [79]). Agents are *rational inattentive* when the cost of information acquisition is linear in the *mutual information* between the signals and the state.² When learning about the environment is costly and the cost of information satisfies certain properties, decision makers will display certain patterns of behavior. An important behavioral implication is that information is going to be designed with the purpose of decision making, in the sense that information that is not relevant for changing from one action to another is going to be ignored. This implies that each signal is going to be uniquely associated to an action, and that equilibrium strategies will take the form of conditional choice probabilities.

This idea has been extended to strategic environments such as global games, bargaining and potential games, and it constitutes a growing research area whose limits are yet to be known.³ As it will be discussed, this dissertation takes this idea and uses it to study *general* Bayesian games, and a Bayesian persuasion setup with a rational inattentive receiver.

2.3 Misperception of Information

At the core of chapter 5 in this dissertation is the idea that agents may fail to correctly perceive the information they receive. A vast literature in the field of economics and psychology has explored and

²See chapter 3.

³See Yang (2015, [83]), Ravid (2018, [74]), Denti (2015, [29]), Matejka (2016, [61]), Martin (2015, [58]).

documented individual misperceptions of reality, such as overconfidence, wishful thinking, cognitive dissonance and self-deception (Benabou 2008, [14]). The idea that agents' attitudes often reflect distorted beliefs is common in political science and political psychology.⁴

The idea of correlational neglect as a rationale for misperception in information has been a particularly popular one.⁵ Information structures generate correlated rather than mutually independent signals, and agents may fail to recognize such correlation. Correlation may be among signals of different agents, or within (the sequence of) signals received by a single agent. In its many forms, correlational neglect can bring about overconfidence, extremism and polarization, and these links have been explored theoretically and empirically.⁶

Efforts devoted in this line are particularly valuable today. Our world is changing very fast, and some prickly patterns are emerging worldwide. Despite living in an era in which information has become more accessible than ever, we observe voters and politicians all over the world to become more and more polarized (McCarty et al. 2006, [65]). Furthermore, in this modern era, misinformation, fake news, ignorance and ideologies are all mixing up to allow radical (hate-speech based) politicians into power, such as Donald Trump in the US, Jair Bolsonaro in Brazil, Viktor Orban in Hungary, and Andrzej Duda in Poland. Angela Merkel, the German Chancellor, sees such developments as a threat to liberal democracies in general and for the EU in particular.⁷

The humble contribution of Chapter 5 in this dissertation contributes with a grain of sand to the endeavour of understanding this global phenomenon.

2.4 Summary of Chapter 3

An nice feature of the standard Bayesian persuasion model is that it can be applied into many economic situations.⁸ One could think in a lobbying group (the sender) commissioning a research with the goal of influencing a politician (the receiver). One can also think of a firm trying to make a sell by providing product information to his potential customers. Nevertheless, to study these situations with the basic persuasion setup, we need to assume that receiver does not have access to any additional information. In reality, however, politicians may carry out their own research, and costumers can find online information about the characteristics of the products. But this information needs to be gathered and processed, and this is undoubtedly costly.

This Chapter extends the basic Bayesian Persuasion setup by allowing the receiver, after observing the information disclosed by the sender, to gather additional information. Specifically, after observing the information provided by the sender, the receiver updates his beliefs and decides whether to gather additional information. Information acquisition is costly for the receiver, and its cost is assumed to be linear in *mutual information*, so she is *rational inattentive*.⁹ The cost is scaled by an exogenous cost parameter, which I take as the receiver's ability of gathering information.

I study the impact on wellbeing of changes in receiver's ability of gathering information. On the one

⁴See Lane (1959, [51]), Hochschild (1981 [42]; 1996 [43]), Kuran (1997, [50]), Jost and Major (2001, [44]).

⁵See De Marzo et al. (2003, [28]), Ortleva and Snowberg (2015, [69]), Glaeser and Sunstein (2009, [36]) and Levy and Razin (2015, [55])

⁶See Ortleva and Snowberg (2015, [69]) and Levy and Razin (2015, [55]).

⁷See the interview <https://www.theguardian.com/world/2020/jun/26/for-europe-survive-economy-needs-survive-angela-merkel-interview-in-full>.

⁸See Kamenica and Gentzkow (2011, [31]).

⁹See for instance Sims (2003, [77]; 2006, [78]), Matejka and McKay (2014, [62]), Caplin and Dean (2013, [21]) and Steiner et al. (2018, [79]).

hand, I show that higher ability of gathering information is not necessarily better for the receiver. Indeed, higher ability to gather information may lead the sender to provide less information about the state, and less overall information disclosed at equilibrium, which cannot be better for the decision maker. As it happens with information itself, the ability of gathering information is not always desirable in strategic environments. In a game, equilibrium payoffs may be lower when more information is provided to the players.¹⁰ Similarly in this chapter, even though higher ability of gathering information is always desirable in a non-strategic environment, it may be detrimental for the wellbeing of players when decisions are taken strategically.

I also show that the sender cannot be better off if the receiver becomes better at gathering information. That is, higher ability of gathering information reduces the set of sender's feasible strategies, so he cannot be better than before.

To solve the model I take advantage of methods already available in the literature. Indeed, the paper looks at optimal splitting of posteriors to solve the sender's problem as well as receiver's problem.¹¹ This method, originally proposed by Aumann and Maschler (1967, [4]; 1968, [5]; 1995, [9]), allows the analyst to use graphical tools to solve any particular application of the model.

To facilitate the search of equilibrium even further, I use the fact that for every equilibrium in which the receiver decides to gather additional information, there exists another equilibrium in which the receiver does not. That is, one can always look for an equilibrium in which the receiver does not learn. The reason is simple: if the receiver *is* learning at some equilibrium, the sender could give that extra information in the first place without changing his expected payoff (because the receiver is learning anyway). Note however that receiver's payoff may change, as he is no longer incurring the cost of gathering information. This observation relies on the assumption that information acquisition is costless for the sender.

Related Literature. This paper fits into the literature of Bayesian persuasion. Extensions of the standard model range from multiple senders (Kamenica and Gentzkow 2017, [33]) to a privately informed sender (Hedlund 2017, [41]; Kolotilin et al. 2017, [49]), a privately informed receiver (Kolotilin et al. 2017, [49]), limited communication capacity (Treust and Tomala 2019, [52]) and heterogeneous priors (Alonso and Camara 2016, [2]), to name a few. Kamenica and Gentzkow (2014, [32]) introduce costly information acquisition into the model by assuming that the sender faces a cost of disclosing information to the receiver. The closest papers to this one are Bloedel and Segal (2018, [19]) and Matysková (2018, [64]). In Bloedel and Segal (2018, [19]) the receiver observes the sender's information strategy and chooses, before observing the signal realization, a costly information strategy to learn about the sender's signals. The framework in Matysková (2018, [64]) is analogous to this paper. As opposed to Matysková (2018, [64]), this paper focuses on the double concavification technique to solve instances of the model and provides a simple application to the literature of informational lobby in political economy.

This paper is currently a joint work with Ludmila Matysková, and the joint work is under revise and re-submit at the Journal of Economic Theory. What happened with this paper is that Ludmila and I were working in parallel in the same model, and we both submitted our papers to the journal. The editor offered us to work on a joint version of the paper, and to re-submit it when it is ready.

¹⁰See for instance Bassan (2003, [13]), Gossner (2000, [38]) and Neyman (1991, [68])

¹¹See Kamenica and Gentzkow (2011, [31]) and Caplin and Dean (2013, [21]). To use the method upon the receiver's problem, there is no need to assume entropy based cost functions. It is sufficient to look at posterior separable cost functions, for which entropy-based cost functions are a particular case.

2.5 Summary of Chapter 4

Many economic situations entail strategic decision making under uncertainty. In such cases, agents decide not only which action to take, but also how much to learn about the environment. For example, firms need to learn about market conditions before entering into a new market. They can obtain information about consumers' demand, their (own) cost of production, and the technologies used by the other firms in the market. Bidders on eBay can collect information about the product before making their bids. They may learn about the durability of the product, its handiness, and after-sell services. Finally, voters in an election may want to learn about the policies proposed by different politicians before going to the ballots.

This Chapter studies costly information acquisition by multiple decision makers facing a strategic environment. Players choose an arbitrary *signal* about an unknown state of nature, and they incur in a cost that is increasing in the reduction of their uncertainty. Once information has been chosen, agents play a Bayesian game.

Studying games of incomplete information with flexible information acquisition presents several technical difficulties. First, the realization space of the signals can be anything, from a finite alphabet to a non-separable infinite dimensional space. This alone brings about a problem of existence, because the space of probability distributions over the realization space is potentially non-compact. The second technical difficulty is related to the computation of equilibrium. In particular, for an equilibrium to exist, it may be necessary for the players to randomize over actions (after the observation of a signal) or over experiments. Hence, the problem of finding an equilibrium may become a daunting task.

The contribution of the paper is threefold. First, I show that for every equilibrium in (a fairly large class of) games with flexible costly information acquisition there exists an *equivalent* pure strategy equilibrium in which the space of signals is equal to the action space.¹² This pure strategy equilibrium entails no randomization over actions nor over experiments, and I call it *canonical Nash equilibrium*. Second, I show that a canonical equilibrium always exists. Finally, I extend some standard results from single-agent information acquisition to strategic environments.

The existence of canonical equilibria in games with flexible information acquisition is appealing for at least three reasons. Firstly, it provides simplifying tools to analyze situations that, in principle, could be very difficult to tackle. Canonical equilibria is an equilibrium where signals recommend *pure* actions and those actions are followed by the players. Hence, it provides a *revelation principle* for games with flexible information acquisition. Second, canonical equilibrium provides a simple setup to extend standard results from single-agent information extraction to strategic environments. Using *mutual information* as the cost function, this paper provides equilibrium (probabilistic) choices and learning strategies. It also extends the *posterior based approach* of Caplin and Dean (2013, [21]), which provides geometrical tools to solve instances of the model. Finally, canonical equilibria addresses the criticism of mixed strategy equilibrium.¹³ Broadly speaking, this criticism says that mixed strategy equilibrium plays no role in behavioral theory (beyond guaranteeing existence of equilibrium), because people does not randomize in practice. Canonical Nash equilibrium *is* a pure strategy equilibrium, both in *information strategies* and in *behavioral strategies*.

The paper concludes with some examples. It discusses matching pennies (a classic game with unique equilibrium in mixed strategies), and a modification of it in which payoffs depend on an unknown state of nature. It presents as well the game battle of sexes and the chicken game, to discuss correlated

¹²Two Nash equilibria are *equivalent* if the joint distribution between actions and states are the same.

¹³For a discussion and replies to such critique, see for instance Harsanyi (1973, [40]), Aumann et al. (1983, [8]), Radner and Rosenthal (1982, [72]), Milgrom and Weber (1985, [66]).

equilibrium.¹⁴

Related Literature. Broadly speaking, this paper relates to three research areas in economics. First, it fits into the rational inattention literature, particularly into the study of games with rational inattentive players. As opposed to what has been done, in this paper we do not look at a *specific* strategic environment, but arbitrary finite games. Furthermore, for the first two main contributions of the paper (Theorem 4.1 and 4.2), I do not assume a particular functional form of the cost function, but I keep four key properties that are also satisfied by entropy based cost functions. Second, this paper fits into the literature of information design. As noted in Bergemann and Morris (2019, [17]), information design problems can be either *literal* or *metaphorical*. In the former, a real information designer provides information to players in order to serve his personal interest. In the latter, the formulation as an information design problem is a metaphor that the analyst wields as a tool. This paper relates to the second interpretation of information design by studying games in which the information structure results from independent interaction among players. I address the question of whether we can restrict attention, without loss of generality, to information structures in which every player's signal *recommends* an action and that action is followed. That is, if we have a *revelation principle* guaranteeing no loss of generality on looking at *direct* information structures. Finally, this paper relates to the purification literature. Started by Harsanyi (1973, [40]), this area studies the conditions under which a pure strategy equilibrium can be obtained from a mixed strategy equilibrium in the context of games of incomplete information. This type of results are called *purification theorems*. This literature focuses on exogenous information structures. Since canonical equilibrium *is* a pure strategy equilibrium, this paper is the first one to provide a purification result for games with endogenous information.

2.6 Summary of Chapter 5

It is hard to explain the rise of political extremism without a notion of *ideology*.¹⁵ Despite its importance, nor economists nor political philosophers have come to agree with a unified definition of the concept. In its many forms, the term is used to label preferences, opinions, beliefs and some types of psychological biases. Political philosopher Terry Eagleton says that ideology is like halitosis, because you can easily recognize it in others but can hardly recognize it in yourself.¹⁶

This paper studies a framework in which ideology is considered as a misperception of reality that agents acknowledge in others but fail to recognize in themselves. Specifically, two ideological groups obtain their private information from biased information sources. I assume that agents from the same group consume news from similar media outlets, so the *signals* they receive are correlated. Agents suffer from *bias neglect*, under which they fail to recognize the bias in their information sources. Furthermore, they suffer from *correlational neglect*, where they fail to recognize the correlation among signals within ideological group. Nonetheless, agents correctly assess the bias and correlation in the information sources of the other group, and they acknowledge that the other group is failing to recognize them. Note that this notion of correlational neglect is different from the one studied in the literature. In Ortoleva and Snowberg (2015, [69]) and Levy and Razin (2015) an individual suffers from correlational neglect when he fails to recognize the correlation among *his* signals. See discussion in section 5.2.¹⁷ Agents also receive information from

¹⁴See Auman (1987, [7]).

¹⁵See Ortoleva and Snowberg (2015, [69]), Martin and Yurukoglu (2017, [60]), Gentzkow and Shapiro (2006, [34]), Prior (2013, [71]).

¹⁶Terry Eagleton (1990, [81]).

¹⁷They also believe that agents in the other ideological group mistakenly believe that they fail to recognize their own bias.

an unbiased publicly available source.

Individuals have to choose their political position on an issue (e.g. tax rates, government spending, etc.), and their payoffs depends on the realization of an unknown state (e.g. interest rates, unemployment, etc.). Their payoffs also depend on the level of agreement with the rest of the population, so there is strategic complementarity among individual political positions. In this context, I study the role of bias and correlational neglect on political extremism and social welfare.

I show that the precision of public information has an ambiguous effect on political extremism. A key factor in that relation is the degree of correlational neglect. If correlational neglect is low, and public information is noisy, increasing the variance of public information may reduce political extremism. The reason is that agents will put more weight in their private signals, motivated not only by its relative increase in precision but also by its increase in coordination power. These two effects are enough to reduce political extremism, despite the increase in the variance of the public signal. If the correlational neglect is large, extremism is always increasing in the variance of public information.

The noise in the private signals is composed by a zero-mean idiosyncratic shock (independent across all agents), and a shock that is common to all members in the same ideological group. As with public information, the variance of the idiosyncratic and group noise have ambiguous effect in political extremism. Regarding welfare, I show that expected social welfare is increasing in the precision of the public signal. I further show that expected welfare is decreasing in the size of the bias of the information sources, and that is decreasing in bias and correlational neglect.

Related Literature. This paper makes a humble theoretical contribution to the literature on behavioral political economy. On the one hand, this paper contributes to the study of correlational neglect. Important papers in this recent literature are Ortoleva and Snowberg (2015, [69]) and Levy and Razin (2015, [55]), who study the effect of correlational neglect on political behaviour. In the former, there is no strategic interaction, and agents fail to recognize that the signals *they* observe are correlated with each other. This implies that agents wrongly assess the informativeness of their signals. They show that correlational neglect generates overconfidence, and they document its impact on political preferences, voter turnout and partisan identification. Levy and Razin (2015, [55]) look at a strategic environment in which the two signals *each* player observes may be correlated. They show that *behavioural voters*, those who misperceive the correlation, may end up making better choices than those who are *rational*, because they may *correctly* change his behaviour after observing two apparently independent signals. In this paper we also look at a strategic model, but I introduce correlational neglect as a failure to recognize the correlation among signals within an ideological group. I also introduce bias neglect, as I take information sources to be biased.

Chapter 3

Bayesian Persuasion with Rational Inattentive Receiver

Abstract

This paper extends the basic Bayesian Persuasion setup by allowing the receiver, after observing the information disclosed by the sender, to gather additional information. Information gathering is costly for the receiver, but it is not for the sender. The cost of information extraction is assumed to be proportional to the reduction of his uncertainty (as measured by the entropy of his beliefs). The cost is parametrized by a scalar, which I take as receiver's ability to gather information. I analyze the impact of changes in receiver's ability of gathering information on the well being of both, sender and receiver. On the one hand, I show that the sender cannot be better off if the receiver's ability of gathering information increases. On the other hand, and in contrast with straightforward intuition, I show that the receiver may be worse off when he is better at gathering information. The reason is that more ability to gather information could lead the sender to provide less information in order to prevent the receiver from choosing a particular action. This, in turn, could imply less total information disclosed at equilibrium, making the receiver to be worse off. Finally, I provide an application on informational lobby.

Keywords: Bayesian persuasion, Rational inattention, Information design.

JEL Classification: D81, D82, D83.

Note: This paper is currently a joint project with Ludmila Matysková (University of Bonn). The joint project is in R&R at the Journal of Economic Theory.

3.1 Introduction

An appealing feature of the simple Bayesian persuasion model is that it has several economic applications.¹ We can consider a setting where a lobbying group (the sender) commissions a research with the goal of influencing a benevolent politician (the receiver). This is the case, for instance, of tobacco and pharmaceutical companies. We could also consider a setting where a firm tries to make a sell by providing product information to his potential costumers. This is the case of car dealer (that let the

¹See Kamenica and Gentzkow (2011, [31]).

customer try the car before buying it), or a software producer (that allows to download a trial version of its product).

To analyze these applications with the basic persuasion setup, it is necessary to assume that receiver does not have any further access to information. In reality, however, politicians may carry out their own research about tobacco and drugs, and customers can find online information about the characteristics of the products. This process of gathering information may be costly for the decision maker. Information needs to be gathered and processed, and decision makers need to allocate their limited attention (Matejka and McKay 2014, [62]).

This paper looks at a simple setup of Bayesian persuasion in which the receiver is allowed to gather additional information. As in the standard setup, the sender influences the decision of the receiver by changing his beliefs about an unknown state of nature. After observing the information provided by the sender, the receiver updates his beliefs and decides whether to gather additional information. Information acquisition is costly for the receiver, and its cost is assumed to be proportional to the reduction of his uncertainty as measured by the entropy of his beliefs. Hence, I consider that the receiver is *rational inattentive*.² The cost is scaled by an exogenous cost parameter, which I take as the ability of gathering information.

I analyze the impact on (sender and receiver's) well being of changes in receiver's ability of gathering information. On the one hand, and against simple appealing intuition, I show that higher ability is not necessary better for the receiver. Indeed, higher ability to gather information may lead the sender to provide less information about the state, and less overall information disclosed at equilibrium. This cannot be better for the decision maker. On the other hand, and in line with the intuition, I show that the sender cannot be better off if the receiver becomes better at gathering information. The reason is that the higher the ability, the more information the receiver is able to gather, which reduces the set of sender's available strategies.

To solve the model I take advantage of methods already available in the literature. I look at optimal splitting of posteriors to solve the sender's problem as well as receiver's problem.³ This method, originally proposed by Aumann and Maschler (1967, [4]; 1968, [5]; 1995, [9]), allows the analyst to use geometrical tools to solve any particular instance of the model.

To facilitate the search of an equilibrium even further, I use the fact that for every equilibrium in which the receiver decides to gather additional information, there exists an equivalent equilibrium in which the receiver does not. That is, one can always look for an equilibrium in which the receiver does not learn. The reason is simple: if the receiver *is* learning at some equilibrium, the sender could give that extra information in the first place without changing his expected payoff (because the receiver is learning anyway). Note however that receiver's payoff may change, as he is no longer incurring in a cost of gathering information. This observation relies on the assumption that information acquisition is costless for the sender.

My results show that, as it happens with information itself, the ability of gathering information is not always desirable in strategic environments. It is well known that the value of information may be negative when decisions are taken strategically. That is, in a game, equilibrium payoffs may be lower when more information is provided to the players.⁴ This paper shows that the same counter intuitive phenomenon

²See Sims (2003, [77]; 2006, [78]), Matejka and McKay (2014, [62]), Caplin and Dean (2013, [21]) and Steiner et al. (2018, [79]).

³See Kamenica and Gentzkow (2011, [31]) and Caplin and Dean (2013, [21]). To use the method upon the receiver's problem, there is no need to assume entropy based cost functions. It is sufficient to look at posterior separable cost functions, for which entropy-based cost functions are a particular case.

⁴See for instance Bassan (2003, [13]), Gossner (2000, [38]) and Neyman (1991, [68])

happens with the ability of gathering information in the context of the Bayesian persuasion model with rational inattentive receiver. That is, even though higher ability of gathering information is always desirable in a non strategic environment, it may be detrimental for the well being of players when decisions are taken strategically.

Related Literature This paper fits into the literature of Bayesian persuasion. Extensions of the standard model range from multiple senders (Kamenica and Gentzkow 2017, [33]) to a privately informed sender (Hedlund 2017, [41]; Kolotilin et al. 2017, [49]), a privately informed receiver (Kolotilin et al. 2017, [49]), limited communication capacity (Treust and Tomala 2019, [52]) and heterogeneous priors (Alonso and Camara 2016, [2]), to name a few. Kamenica and Gentzkow (2014, [32]) introduce costly information acquisition into the model by assuming that the sender faces a cost of disclosing information to the receiver. The closest papers to this one are Bloedel and Segal (2018, [19]) and Matysková (2018, [64]). In Bloedel and Segal (2018, [19]) the receiver observes sender's information strategy and chooses, before observing the signal realization, a costly information strategy to learn about the sender's signals. The framework in Matysková (2018, [64]) is analogous to this paper. As opposed to Matysková (2018, [64]), this paper focuses on the double concavification technique to solve instances of the model and provides a simple application to the literature of informational lobby in political economy.

The rest of the paper is organized as follows. Section (3.2) develops the model of persuasion with rational inattentive receiver. Section (3.3) characterizes the solution and section (3.4) does comparative statics. Section (3.5) shows an application and Section (3.6) concludes.

3.2 The Model

Let $I = \{S, R\}$ be the set of players, a sender and a receiver. There is a unknown state $\omega \in \Omega$, where Ω is finite. Players share a common prior $\mu \in \Delta(\Omega)$. The receiver has available a set of actions A , also assumed finite. Payoff functions are u and v for the receiver and the sender respectively, with $u, v : A \times \Omega \rightarrow \mathbb{R}$.

Players can choose an *experiment*. For the sender, an experiment is denoted by $\pi = (\mathcal{S}, (\pi_\omega)_\omega)$, which consists on a realization space \mathcal{S} and a family of probability distributions $(\pi_\omega)_\omega$ that specifies the probability of receiving signal $s \in \mathcal{S}$ for each state of nature ω . Analogously for the receiver, his experiment is denoted by $\phi = (\mathcal{R}, (\phi_\omega)_\omega)$. Without loss of generality, I assume that $\mathcal{R} = A$.⁵

The game is played as follows. At the first stage, the sender chooses and runs an experiment π . The design and realization of the experiment are publicly observed. In the second stage, the receiver chooses and runs an experiment ϕ , whose design and realization are also publicly observed. In the third stage, the receiver chooses an action $a \in A$.

After observing the realization $s \in \mathcal{S}$ of experiment π , the sender and the receiver update beliefs $q_s \in \Delta(\Omega)$ by $q_s(\omega) = \frac{\pi_\omega(s)\mu(\omega)}{\pi(s)}$ for every ω , where $\pi(\cdot)$ denote the unconditional distribution of signals in \mathcal{S} , so that $\pi(s) = \sum_\omega \pi_\omega(s)\mu(\omega)$. After observing the realization $a \in A$ of experiment ϕ , the sender and the receiver updates beliefs $q_a^s \equiv q_a(q_s) \in \Delta(\Omega)$, obtained $q_a^s(\omega) = \frac{\phi_\omega(a)q_s(\omega)}{\phi(a|s)}$, where $\phi(a|s) \equiv \sum_\omega \phi(a|\omega)q_s(\omega)$ denotes the unconditional distribution of a given s .

Bayesian updating implies that every experiment can be represented by a collection of posteriors that average up the prior.⁶ That is, if $\eta = (\mathcal{S}, (\eta_\omega)_\omega)$ is an experiment performed under prior p , then we can look at $\hat{\eta} = (\mathcal{S}, (q_s, \hat{\eta}_s)_s)$, where $q_s \in \Delta(\Omega)$ for all s , $(\hat{\eta}_s)_s \in \Delta(\mathcal{S})$, and $\int_s q_s d\hat{\eta}_s = p$. The latter representation

⁵See Matejka and McKay (2014, [62]).

⁶See Kamenica and Gentzkow (2011, [31]).

of experiments is preferable if we want to characterize the solution as a splitting of posteriors. In what follows I work with the second definition of experiments, and suppress the *hat* from the notation.

For the sender, information is not costly. He can choose any experiment he wants at zero cost. For the receiver, however, information acquisition is costly. The cost is assumed to be linear in the *mutual information* between states of nature and signals. Mutual information is a measure of the amount of information that the signals provide about the state. Formally, given a prior p , the cost of experiment ϕ is defined as

$$c(\phi; p, \lambda) = \frac{1}{\lambda} \left[H(p) - \sum_a H(q_a) \phi_a \right] \quad (3.1)$$

where $H(\cdot)$ is the Shannon entropy function,⁷ and $\frac{1}{\lambda} \in \mathbb{R}_+$ is the unit cost of information. I consider $\lambda \in \mathbb{R}_+$ as a measure of the ability of a receiver to gather information. The higher the value of λ , the lower is the cost of acquiring information.⁸

At the last stage, when signal s and a have been observed, the receiver chooses a to maximize $E_{q_a^s} [u(a, \theta)]$. Let $a^*(q_a^s)$ denote the solution correspondence of the problem. Since the cost function is separable in posteriors and entropy is a strictly convex function, it is true that if $q_a^s \neq q_{a'}^s$, then $a^*(q_a^s) \neq a^*(q_{a'}^s)$. Hence, we can pair up actions and posteriors and formulate the receiver's problem of choosing experiments as follows

$$\begin{aligned} \max_{\{\phi\}} \quad & \sum_a E_{q_a^s} [u(a, \omega)] \phi_a - c(\phi; q_s, \lambda) \\ \text{s.t.} \quad & \\ \phi \in F(q_s) \equiv & \left\{ (q_a^s, \phi_a)_{a \in A} : \sum_a q_a^s \phi_a = q_s \right\} \end{aligned} \quad (3.2)$$

where $c(\phi; q_s, \lambda)$ is defined in (3.1). The prior q_s is an exogenous parameter in program (3.2), so I let $\phi^*(q_s) \equiv (q_a^*(q_s), \phi_a^*(q_s))_{a \in A}$ denote the solution correspondence for prior q_s . Note that the objective function in (3.2) is continuous in $(q_a, \phi_a)_a$, and $F(q_s)$ is a compact and convex valued correspondence that is continuous on q_s . Therefore, by the Berge Maximum Theorem,⁹ the solution correspondence $\phi^*(q_s)$ is non empty and upper hemicontinuous on q_s .

Given receiver's optimal behavior in the second and third stage, the sender will design his experiment to maximize his expected payoff. If posterior q_s is generated by some signal realization of π , then the expected payoff for the sender is $\sum_a E_{q_a^*(q_s)} [v(a, \omega)] \phi_a^*(q_s)$, so he solves

⁷The Shannon entropy of a discrete random variable with realization space X and probability distribution p is given by $H(p) = -\sum_{x \in X} p(x) \log(p(x))$ and it is a measure of the uncertainty of x . I consider natural logarithm and use the convention that $0 \cdot \log(0) = 0$, which is justified by the continuity of $H(\cdot)$.

⁸The properties of the cost function that we use to solve the model will be discussed in Section 3.3.1. In general, all results and methods apply for any uniformly posterior separable cost function (see Caplin and Dean 2015[24]).

⁹See Leininger (1984, [54]).

$$\begin{aligned}
& \max_{\{\pi\}} \sum_{a,s} E_{q_a^*(q_s)} [v(a, \omega)] \phi_a^*(q_s) \pi_s \\
& \text{s.t.} \\
& \pi \in G(\mu) \equiv \left\{ (q_s, \pi_s)_{s \in \mathcal{S}} : \int_s q_s d\pi_s = \mu \right\}
\end{aligned} \tag{3.3}$$

An equilibrium of the game is a pair $(\pi, (\phi_s)_s)$ such that π solves (3.3), and for every $s \in \mathcal{S}$, ϕ_s solves (3.2).

3.3 Solving the Model

In what follows I proceed to find Nash equilibrium of the model by taking advantage of methods already available in the literature. I will solve the game using backward induction, so we will look at sub-game perfect Nash equilibrium.¹⁰ The method I use consists on looking at optimal splitting of posteriors to solve the sender's problem as well as the receiver's problem. The section concludes with a simple example that shows how to use the graphical tools provided by the method.

3.3.1 The receiver's problem

To solve receiver's problem (program 3.2), I follow the posterior based approach proposed by Caplin and Dean (2013, [21]). This approach consists on looking at the *net benefit* given by each action/posterior pair, which corresponds to the expected benefit of using an action at a certain posterior, minus the information cost associated with that posterior. This method is useful not only for equation 3.1, but for any uniform posterior separable cost function.¹¹ Hence, we rewrite Problem (3.2) as

$$\begin{aligned}
& \max_{\{\phi\}} \sum_a N^a(q_a) \phi(a) - \frac{1}{\lambda} H(q_s) \\
& \text{s.t.} \\
& \phi \in F(q_s) \equiv \left\{ (q_a^s, \phi_a)_{a \in A} : \sum_a q_a^s \phi_a = q_s \right\}
\end{aligned} \tag{3.4}$$

where the net benefit of action a at posterior q_a is $N^a(q_a) \equiv \sum_{\omega} u(a, \omega) q_a(\omega) + \frac{1}{\lambda} H(q_a)$. Note that there are $|A|$ net benefit functions. Then define the upper envelope of the net benefit functions as $N(q) \equiv \max_{\{a \in A\}} \{N^a(q)\}$, and $U^*(q_s) \equiv \max_{\{\phi \in F(q_s)\}} \sum_a N^a(q_a) \phi(a) - \frac{1}{\lambda} H(q_s)$. Finally, for any function $f(p)$ define its concave closure as $cav(f)(p) = \min_{\{g\}} \{g(p)\}$ where the minimum is taken over all concave functions g such that $g(p) \geq f(p)$ for all p . The concave closure of a function g is also referred to as the *concavification*¹² of g .

The main insight given by Caplin and Dean (2013, [21]) is twofold. First, they show that $U^*(q_s) = cav(N)(q_s) - \frac{1}{\lambda} H(q_s)$. Second, they show that, given a prior q_s , the *concavification* of $N(q)$ can be

¹⁰A discussion over non sub-game perfect Nash equilibrium can be find at the end of the section.

¹¹See Caplin and Dean (2013, [24]).

¹²See Aumann and Maschler (1995, [9]).

obtained finding the highest hyperplane that lies weakly above the net benefit functions and that contains the prior. Actions whose net benefit functions are supporting the hyperplane, are the actions chosen with positive unconditional probability. For the second result, it is necessary to identify the subsets of actions that, given a prior q , can be used to build the concave closure of $N(q)$. Formally, let

$$B(q) \equiv \left\{ C \in \mathcal{P}(A) : \exists \phi \in \text{int} \{ \Delta(C) \} \text{ and } (q_c)_{c \in C} \text{ s.t. } \text{cav}(N)(q) = \sum_c N^c(q_c) \phi_c \right\} \quad (3.5)$$

where $\mathcal{P}(A)$ denotes the power set of A , and a set $C \in B(q)$ is a subset of actions from which $\text{cav}(N)(q)$ can be constructed by some experiment. If there is more than one set in $B(q)$, the solution is not unique.

Theorem 3.1. *Caplin and Dean (2013, [21])*

For every $q \in \Delta(\Omega)$

1. The value of problem (3.4) is $U^*(q) = \text{cav}(N)(q) - \frac{1}{\lambda} H(q)$.
2. For every $C \in B(q)$, a solution $(q_c, \phi_c)_c$ to problem (3.4) satisfies $\sum_{c \in C} q_c(\omega) \phi_c = q(\omega) \quad \forall \omega \in \Omega$ and, and for every pair of actions $c, c' \in C$,

$$\frac{q_c(\omega)}{q_{c'}(\omega)} = \exp \left(\frac{1}{\lambda} [u(c, \omega) - u(c', \omega)] \right) \quad \forall \omega \in \Omega \quad (3.6)$$

Equations (4.9), that hold for actions chosen with positive probability under prior q , are called the Invariant Likelihood Ratios for Chosen Acts (ILR). An interesting particular case of the theorem above is when the cardinality of a set in $B(q)$ with the smallest number of elements is equal to the cardinality of Ω . In this case, the highest hyperplane is uniquely defined and optimal posteriors can be obtained independently of the Bayesian constraints. That is, equations (3.6) are enough to obtain the posteriors, and the unconditional probability of actions can be obtained by using Bayesian constraints afterwards.¹³

3.3.1.1 The sender's problem

The sender's problem is analogous to the Bayesian persuasion problem in Kamenica and Gentzkow (2011, [31]). Indeed, let $\hat{v}(q_s) = \sum_a E_{q_a^*(q_s)} [v(a, \omega)] \phi_a^*(q_s)$, and $V^* = \max_{\{\pi \in F\}} \sum_s \hat{v}(q_s) d\pi_s$, where $(q_a^*(q_s), \phi_a^*(q_s))_a$ denotes again the solution correspondence of receiver's problem for prior q_s . The following result follows.

Proposition 3.1. *Kamenica and Gentzkow (2011, [31])*

1. There exists a solution $\pi = (\mathcal{S}, (q_s, \pi_s)_s)$ to problem (3.3) if and only if there exists a solution $\hat{\pi} = (\hat{\mathcal{S}}, (\hat{q}_s, \hat{\pi}_s)_s)$ such that $\hat{\mathcal{S}}$ is finite.
2. The value of problem (3.3) for a given prior μ , is given by $\text{cav}(\hat{v})(\mu)$.

Note that since the receiver is now choosing experiments (and not only actions), proposition (3.1) does not guarantee that the set of signals can be reduced to the set of actions. It does, however, guarantee that there are finitely many signals in \mathcal{S} . Indeed, since Ω and A are finite, Caratheodory's Theorem implies that only finitely many signal realizations are needed to obtain $\sum_s \hat{v}(q_s) \pi_s = V^*(\mu)$.

¹³In other words, if $B(q) = |\Omega|$, we need to find if $B(q) \times (|\Omega| - 1)$ unknowns, and the ILR conditions give $|\Omega| \times (B(q) - 1)$ equations. However, if $B(q) < |\Omega|$, ILR conditions give $|\Omega| \times (B(q) - 1)$ equations to solve for $B(q) \times (|\Omega| - 1)$ unknowns. Hence, we are missing $|\Omega| - B(q)$ equations to solve the problem. If we consider also the Bayesian constraints, we add $|\Omega|$ equations and $B(q)$ unknowns, so now the problem can be solved. Indeed, in this case, although there exists infinitely many highest hyperplanes, there exists a unique highest lower-dimensional hyperplane that satisfies Bayesian constraints.

As in the standard Bayesian persuasion model, Proposition (3.1) does not guarantee the existence of a solution. In Kamenica and Gentzkow (2011, [31]), existence is guaranteed by the assumption that the receiver chooses, in case of indifference, the *sender's preferred* action. This assumption guarantees that \hat{v} is upper hemicontinuous, so that its graph is closed and the concave closure can always be constructed. In our model, \hat{v} is continuous on $\phi^*(q_s)$, which is non empty and upper hemicontinuous on q_s . Hence, \hat{v} is upper hemicontinuous and the solution exists. There is no need to focus on sender-preferred equilibrium.

The solution to the whole model is directly obtained by putting together the sender's and receiver's solution. For a given λ and payoff function of the receiver, plot the net benefit functions and identify the collection of optimal posteriors (as function of the prior) by building the concave closure. Then obtain receiver's optimal experiments (choice probabilities) and plot the sender's expected payoff and obtain the solution by building again the concave closure.

3.3.1.2 A Simple Example

To fix ideas, I finish this section with a simple example that shows how the solution to the Bayesian persuasion model with rational inattentive receiver can be derived graphically. Suppose that $A = \{a, b\}$, that $\Omega = \{\omega_1, \omega_2\}$ and that $\mu \equiv p(\omega_1) = \frac{2}{5}$. Suppose first that $\lambda \rightarrow \infty$, and that the payoffs are given in the following tables.¹⁴

$u(a, \omega)$	ω_1	ω_2
a	1	0
b	0	1

Table 3.1 – Receiver's Payoffs

$v(a, \omega)$	ω_1	ω_2
a	0	0
b	1	1

Table 3.2 – Sender's Payoffs

When $\lambda \rightarrow 0$, we are back in the classic Bayesian persuasion setup. Optimal posteriors induced by the sender are $q_0 = 0$ and $q_1 = \frac{1}{2}$. Figure (3.1) and (3.2) show expected payoffs for the receiver and the sender respectively. Figure (3.2) also shows the solution to sender's problem.

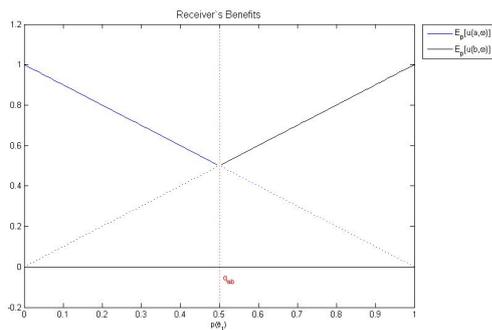


Figure 3.1 – Receiver's Expected Benefit

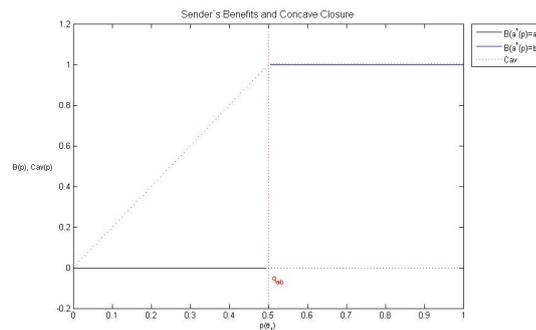


Figure 3.2 – Sender's Expected Benefits

Now assume that $\lambda = 1$. Then, I plot the net benefit functions to find receiver's solution. Figure (3.3) shows the net benefit functions and the its concave closure. Indeed, using Theorem (3.1), we obtain $q_a = 0.27$ and $q_b = 0.73$. Note that, $B(q) = \{a\}$ if $q < q_a$, $B(q) = \{a, b\}$ if $q \in [q_a, q_b]$, and $B(q) = \{b\}$ if $q_b < q$, so for $\mu \equiv \frac{2}{5}$, optimal posteriors are indeed q_a and q_b . Figure (3.4) shows the sender's solution to the problem, so the functions $\hat{v}(q)$ and $cav(\hat{v})(q)$ are depicted. Optimal posteriors induced by sender experiments are then going to be $q_0 = 0$ and $q_1 = q_b$.

¹⁴The payoff structure is equal to the one in Example 1 of Kamenica and Gentzkow (2011, [31]).

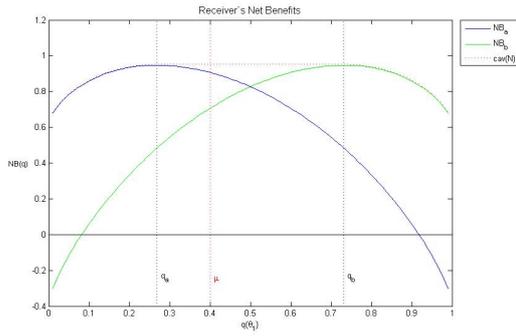


Figure 3.3 – Net Benefits and $cav(N)(p)$

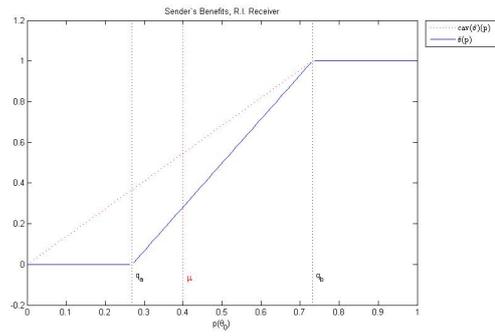


Figure 3.4 – Functions $\hat{v}(p)$ and $cav(\hat{v})(p)$

Two interesting things happen in this example. First, the lower is λ , the better is the sender (note that the limit when λ goes to zero gives the result of the standard Bayesian persuasion model). Second, the higher is λ , the better is the receiver. In both cases, the receiver does not gather additional information, but when $\lambda = 1$, the possibility of gathering information encourages the sender to disclose more information at equilibrium. In the next section I show that the first relation always hold, not only for this particular example. Then, I show that the second relation does not always holds. That is, it is not true that the higher is λ (i.e. the better is the receiver at gathering information), the better is the receiver.

3.3.2 Discussion

Up to this point we have solved the model by using backward induction, so that the equilibrium we find are all sub-game perfect Nash equilibrium. Nevertheless, there are many other Nash equilibrium that are not sub-game perfect. For instance, it can be observed that full information disclosure can always be implemented as a Nash equilibrium. Furthermore, any sender's information strategy that provides more information than the sub-game perfect Nash equilibrium can also be implemented as an equilibrium.

In particular, suppose that the sender is choosing a fully revealing experiment. Suppose also that the receiver is not gathering further information when the fully revealing experiment is performed, but that for any other experiment performed by the sender, the receiver commits to a fully revealing experiment. Then, it is certainly optimal for the sender to reveal the true state, as the state will be revealed anyway. On the other hand, the receiver will never see herself having to disclose full information (which is always not optimal), and so that strategy is also optimal. Hence, full information disclosure is a Nash equilibrium.

Given the discussion above, it seems that when the receiver has the power to commit to any off-equilibrium learning strategy, then all new Nash equilibria that appear would return a weakly higher equilibrium payoff for the receiver. This is certainly true for the binary states and actions setup. My conjecture is that the same holds true for arbitrary number of states and actions, although a detailed examination should be done.

3.4 Comparative Statics

In this section I present the main results of the paper. First, we will see that the sender is never better off when the ability of gathering information increases. Second, by means of an example, I show that the receiver is not necessarily better off if the ability of gathering information increases.

3.4.1 Preamble

To simplify notation, let $q_{as} \equiv q_a(q_s)$ and $\phi_{as} \equiv \phi_a(q_s)$, and recall that I denoted $B(q)$ the set of actions chosen with positive probability for prior q . Since sender's function \hat{v} is a polytope, the solution to sender's problem $cav(\hat{v})$ is built either on the kinks of \hat{v} or it overlaps with \hat{v} along one or more hyperplanes. Hence, if $s \in \mathcal{S}$ is observed at equilibrium and $q_{as} \neq q_{bs}$ (that is, the receiver learns), it has to be the case that the sender is indifferent between letting the receiver learn or providing the information from in the first place.

Hence, whenever the receiver does run an experiment after observing some s , it is possible to find another equilibrium in which the sender provides the information needed by the receiver, so the receiver never runs an experiment. Although this new equilibrium keeps the final (joint) distribution of actions and the states unchanged, it changes equilibrium payoffs. The reason is that, since the sender faces no cost of gathering information, he could provide (because he is indifferent) the additional information that the receiver is gathering. By doing this, the sender is saving the receiver the cost of gathering that information.

The fact that for every equilibrium in which the receiver learns there exists an equivalent equilibrium in which the receiver never runs an experiment depends crucially on the assumption that the sender has no cost of providing information. If information acquisition were to be costly for the sender as well, he would leave as much learning as possible in the hands of the receiver.¹⁵

Lemma 3.1. *For every equilibrium $\{(q_s, \pi_s)_s, (q_{as}, \phi_{as})_{a,s}\}$ such that for some $s \in \mathcal{S}$ and some $a, b \in B(q_s)$ $q_{as} \neq q_{bs}$, there exists another equilibrium $\{(\hat{q}_s, \hat{\pi}_s)_s, (\hat{q}_{as}, \hat{\phi}_{as})_{a,s}\}$ such that for every $s \in \hat{\mathcal{S}}$, $q_{as} = \hat{q}_s$ for all $a \in B(\hat{q}_s)$.*

Proof. Suppose there exists a triplet $a, b \in A$ and $s \in \mathcal{S}$ such that $q_{as} \neq q_{bs}$, let $\hat{\mathcal{S}} \equiv (\mathcal{S} \setminus s) \cup B(q_s)$, and for $a \in B(q_s)$ let $\hat{q}_a = q_{as}$ and $\hat{\pi}_a = \pi_s \phi_a$. For $s \notin B(q_s)$, let $\hat{q}_s = q_{as}$ and $\hat{\pi}_s = \pi_s$. It is direct to see that for all $s \in \hat{\mathcal{S}}$ we must have $q_a(\hat{q}_s) = \hat{q}_s$. Under experiment $\hat{\pi}$, the joint distribution of actions and states remain unchanged, so the expected payoff of the sender remains unchanged. The expected payoff of the receiver is increased by $c(\phi^*(q_s); \lambda, q_s)$. □

3.4.2 Payoff Monotonicity

By lemma (3.1), we can always look at equilibria where the receiver never runs an experiment. After observing signal s , no information is gathered by the receiver and some action a is chosen. As equilibrium multiplicity can easily arise in our model, we need an adequate definition of payoff monotonicity in order to do comparative statics. In particular, we look at the lowest and highest payoffs in the non-learning equilibrium set, and we see how these values change as we change the cost parameter λ . We say that a payoff is *locally increasing (decreasing)* if, for a small increase in λ , both the lowest and the highest payoff in the equilibrium set increase (decrease).¹⁶

Formally, let NE_λ be the set of non-learning Nash equilibria at cost parameter λ . Let $F_\lambda(\hat{\pi}, \hat{\phi})$ denote the payoff of a player at equilibrium $(\hat{\pi}, \hat{\phi}) \in NE_\lambda$, and let F_λ^M and F_λ^m denote the maximum and minimum payoffs at cost parameter λ respectively.¹⁷ We say that a payoff F_λ is *locally increasing*

¹⁵Kamenika and Gentzkow (2014, [32]) study the case in which the receiver does not have the ability to gather information, but the sender faces a cost of providing public information. This changes the shape of function \hat{v} , but the concavification method can still be applied.

¹⁶Milgrom and Roberts (1994) develop a framework to make robust comparisons of equilibria as exogenous parameters in the model change.

¹⁷That is, $F_\lambda^m \equiv \inf_{\{(\hat{\pi}, \hat{\phi}) \in NE_\lambda\}} F_\lambda(\hat{\pi}, \hat{\phi})$ and $F_\lambda^M \equiv \sup_{\{(\hat{\pi}, \hat{\phi}) \in NE_\lambda\}} F_\lambda(\hat{\pi}, \hat{\phi})$.

in λ if there exists a small neighborhood $B_\epsilon(\lambda)$ of radius ϵ around λ such that for all $\lambda', \lambda'' \in B_\epsilon(\lambda)$, $\lambda' > \lambda'' \Rightarrow$ (i) $F_{\lambda''}^m \leq F_{\lambda'}^m$ and (ii) $F_{\lambda''}^M \leq F_{\lambda'}^M$. Payoff F_λ is said to be *locally decreasing* if the inequalities in (i) and (ii) are reversed. Finally, when a payoff F_λ is *locally increasing (decreasing)* for every λ , then we say that it is *increasing (decreasing)*.

3.4.3 The Sender

First, we show that the sender cannot be better off if the receiver becomes better at gathering information. That is, a lower information cost parameter λ reduces the set of sender's feasible strategies, so he cannot be better than before.

Proposition 3.2. *The sender's expected equilibrium payoff is decreasing in λ .*

Proof. Fix prior μ_0 and λ , and let V_λ^m and V_λ^M be the maximum and the minimum sender's payoff in the equilibrium set NE_λ . Note that for every equilibrium in NE_λ , optimal strategies for the receiver entail no learning and directly take an action, so equilibrium strategies are fully characterized by sender's experiments. Since λ scales the whole function up or down, it is true that if the receiver is not learning at λ' , she would not be learning at $\lambda'' < \lambda'$. Hence, any distribution of posteriors available to the sender under λ' are also available for λ'' . Hence, $NE_{\lambda''} \subseteq NE_{\lambda'}$, so $V_{\lambda''}^m \leq V_{\lambda'}^m$ and $V_{\lambda''}^M \leq V_{\lambda'}^M$, and sender's expected equilibrium payoff is decreasing in λ . \square

3.4.4 The Receiver

The next proposition states that the receiver may be worse off when it is cheaper for him to gather additional information. A lower cost of gathering information may lead the sender to provide less information about the state, and less overall information disclosed at equilibrium. This cannot be better for the decision maker.

Proposition 3.3. *The receiver's expected payoff is not necessarily locally decreasing in λ .*

Proof. The proof goes by a counter example, which can be found in section 3.5. \square

Proposition 3.3 shows that, as it happens with information itself, the ability of gathering information is not always desirable in strategic environments. Here, the same counter intuitive phenomenon happens with the ability of gathering information when information acquisition is costly. That is, even though higher ability of gathering information is always desirable in a non strategic environment, it may be detrimental for the well being of players when decisions are taken strategically.

3.4.5 A positive Result

It is hard to provide general conditions that fully characterize situations in which the receiver is always (weakly) better when it is cheaper for him to gather information. In general, three actions and two states are enough to build counter examples, even when assuming state independent preferences for the sender. In what follows, we center our attention in the binary state case with binary actions. For this case, I show that if the sender strictly prefers to provide information at prior μ , then the payoff of the receiver is locally increasing in λ . Formally, let $U_\lambda(\hat{\pi}, \hat{\phi})$ denote the receiver's payoff at equilibrium $(\hat{\pi}, \hat{\phi}) \in NE_\lambda$, and let U_λ^M and U_λ^m denote the maximum and minimum payoffs at cost parameter λ respectively. Then, we have the following result.

Proposition 3.4. *Suppose that $V_\lambda(\hat{\pi}, \hat{\phi}) > \hat{v}(\mu_0) \forall (\hat{\pi}, \hat{\phi}) \in NE_\lambda$, and $|\Omega| = |A| = 2$. Then, U_λ is locally increasing in λ .*

Proof. Without loss of generality, suppose that $\Omega = \{\omega_0, \omega_1\}$, that $\mu \equiv \text{prob}(\omega_0)$, and that $a^*(\mu = 0) = a$ and $a^*(\mu = 1) = b$. It clear that there exists an interval, say $[p_\lambda^a, p_\lambda^b]$, in the interior of $[0, 1]$ in which the receiver wants to run an experiment. Hence, because of lemma 3.1, the support of each $\hat{\pi}$ has two elements, $\mu_a \in [0, p_\lambda^a]$ and $\mu_b \in [p_\lambda^b, 1]$. Since $V_\lambda(\hat{\pi}, \hat{\phi}) > \hat{v}(\mu_0)$, the sender could be inducing posteriors $(0, p_b)$, $(p_a, 1)$ and $(0, 1)$. If there are multiple equilibria, two situations could be possible. First, the sender is inducing any $\mu_a \in [0, p_\lambda^a]$ along with p_b or 1. Second, he is inducing any $\mu_b \in [p_\lambda^b, 1]$ along with 0 or p_a . Note that if the sender is indifferent between inducing any $\mu_b \in [p_\lambda^b, 1]$ and any $\mu_a \in [0, p_\lambda^a]$, then $V_\lambda(\hat{\pi}, \hat{\phi}) = \hat{v}(\mu_0)$, so we do not look at that case. If $V_\lambda(\hat{\pi}, \hat{\phi}) > \hat{v}(\mu_0)$, there exists a small neighborhood $B_\epsilon(\lambda)$ of radius ϵ around λ such that for all $\lambda' \in B_\epsilon(\lambda)$, $V_{\lambda'}(\hat{\pi}, \hat{\phi}) > \hat{v}(\mu_0)$. When the receiver is not learning at λ' he is not learning at λ'' , so for $\lambda', \lambda'' \in B_\epsilon(\lambda)$, $\lambda' > \lambda''$ implies that $p_{\lambda'}^a \leq p_{\lambda''}^a$ and $p_{\lambda'}^b \leq p_{\lambda''}^b$. It is direct to see that for every possible case, $U_{\lambda'}^M \geq U_{\lambda''}^M$ and $U_{\lambda'}^m \geq U_{\lambda''}^m$, so the payoff of the receiver is locally increasing in λ . \square

3.5 An Application: Informational Lobby

In the literature of political economy, lobbyists are assumed to influence politicians through the provision of either money or information (Cotton and Dellis 2016, [26]). Although influence through payments is certainly viewed as corruption, influence through information is often seen as desirable, because it leads to better informed policy decisions (Austen-Smith and Wright 1992, [11]; Cotton 2009, [25]). However, it has been shown that even when there are no payments and no information distortions or manipulations, informational lobby could still be detrimental for society.

For instance, Cotton and Dellis (2016, [26]) show that when the policy maker can obtain additional information by paying a cost, persuasion by the lobbying group could lead to worse policy decisions. In their model, the policy maker has to prioritize among issues, and the lobby could shift the attention of the policy maker towards less promising policy decisions that are relevant only for the lobbying group. Other models of informational lobby show that inefficiencies can be generated by deception (Rasmusen 1993, [73]), asymmetric information (Esteban and Ray 2006, [30]) and not benevolent politicians (Daley and Snowberg 2011, [27]).

This application looks at a slightly different question. In our setup, a lobbyist group persuades a policy maker by providing public information. The policy maker, however, can gather his own private information. The cost is parametrized by a scalar, which represents the ability of the politician to gather his information. Hence, we ask the following question. When there is informational lobby and the policy maker can gather his own private information, is she always better off as she becomes better at gathering information? Is the policy maker going to make better decisions when he is better at gathering information? Our answer to this questions is **not necessarily**.

3.5.1 Basic Setup

We consider a policy maker who has three policy options $A = \{p_1, p_2, p_3\}$, and it has to choose which of them to implement. He can implement only one, because he lacks time and resources to implement more.¹⁸ We also assume that he will always implement one of the policies, because doing nothing could be seen as incompetence (and detrimental for future elections). The state of the economy is the unknown

¹⁸See Cotton and Dellis (2016, [26])

state of nature, and it could be either bad (ω_b) with probability $\mu = \frac{1}{2}$ or good (ω_g) with probability $1 - \mu$. From the point of view of the policy maker, p_1 is better to be implemented in the good state, p_3 in the bad state, and the p_2 corresponds to a safe policy, which is preferable when there is too much uncertainty about the state.

The lobbying group only cares about policy p_2 . This policy is the one related to the guild who the lobbyist is representing. Specifically, we assume that the lobbyist receives his status quo payoffs whenever policy p_1 and p_3 are implemented. Payoffs are represented in the following table.

$u(p, \omega)$	ω_g	ω_b
p_1	1	0
p_2	$\frac{3}{4}$	$\frac{3}{4}$
p_3	0	1

Table 3.3 – Policy Maker Payoffs

$v(p, \omega)$	ω_g	ω_b
p_1	a	b
p_2	αa	βb
p_3	a	b

Table 3.4 – Lobbyist Payoffs

We assume that $a > b$, so what the lobbying group obtains in the good state is higher than what it gets on the bad state. Parameters α and β represent how does the public policy affects its payoff in each of the states.

If $\alpha, \beta > 1$, then the lobbying group always prefers policy p_2 over the status quo. If $\alpha > 1 > \beta$, then the policy is increasing his payoff in the good state and reducing in in the bad state. Hence, for such parameters, policy p_2 could be though as a pro-cyclical policy. In turn, if $\beta > 1 > \alpha$, the policy delivers a lower payoff in the good state and higher payoff in the bad state. In this case, it could be interpreted as a counter-cyclical policy. If $1 > \alpha, \beta$, the lobbying strictly prefers the status quo over policy p_2 .

3.5.2 Solution and Comparative Statics

The solution to the receiver problem is straight forward. At μ he will not learn anything and he will choose policy p_2 . Nevertheless, to find the solution for the whole game we need to identify the critical posteriors under which the receiver changes his behavior.

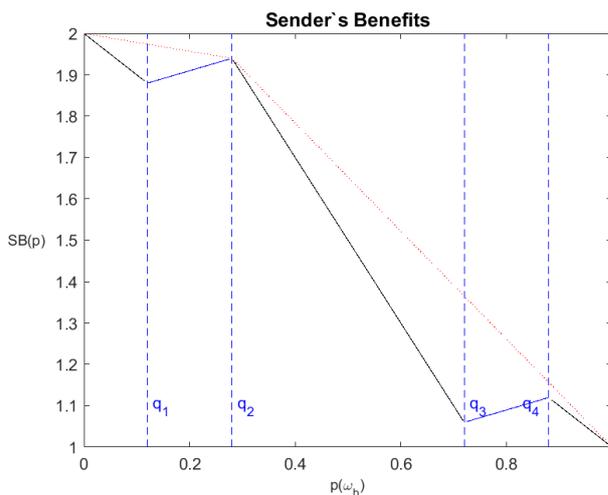


Figure 3.5 – Sender's Solution for λ

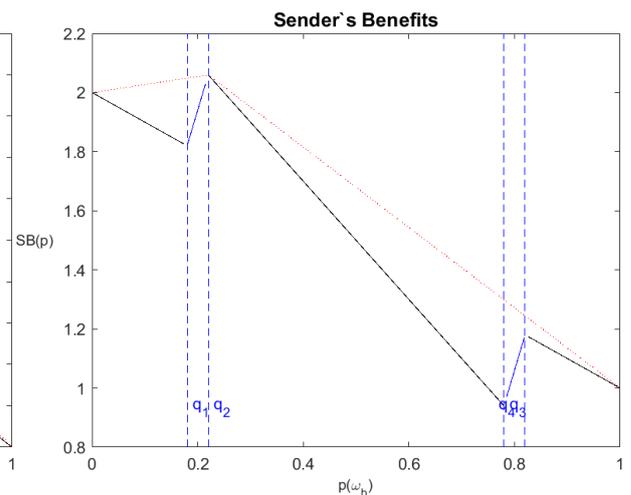


Figure 3.6 – Sender's Solution for $\lambda' < \lambda$

Given the payoffs of the policy maker, it is clear that there exists an ordered quadruple (q_1, q_2, q_3, q_4) such that, $B(q) = \{p_1\}$ if $q < q_1$, $B(q) = \{p_1, p_2\}$ if $q \in [q_1, q_2]$, $B(q) = \{p_2\}$ if $q \in [q_2, q_3]$, $B(q) = \{p_2, p_3\}$ if $q \in [q_3, q_4]$ and $B(q) = \{p_3\}$ if $q_4 < q$.

Suppose now that $\alpha > 1 > \beta$. In this case, the posteriors induced by the sender's experiment are $(q_2, 1)$. Now, note that the interval $[q_1, q_2]$ are the posteriors under which the receiver decides to gather further information. Consequently, since λ scales the whole cost function up, a lower $\lambda' < \lambda$ would generate an interval $[q'_1, q'_2] \subset [q_1, q_2]$, which implies that $q'_2 < q_2$. Consequently, the policy maker is getting more information when his ability is lower.

Intuitively, the lobbyist would like the policy maker to choose p_2 whenever the good state is more likely, and policy p_3 otherwise.¹⁹ However, the lobbyist knows that if he provides too much information about the good state, the policy maker would start learning and choosing with positive probability policy p_1 , which is not good for the sender. Hence, in order to keep the receiver away from choosing policy p_1 , the sender would provide *just enough* information to the receiver.

In turn, if the policy maker becomes better at gathering information, then the lobbyist would have to provide even less information in order to keep him away from choosing policy p_1 . Hence, the policy maker would get less information and be worse off.

The same argument applies for the case in which $\alpha < 1 < \beta$. At $\mu = \frac{1}{2}$, the lobbying group would induce $(0, q_3)$, and he will try to keep the policy maker away from choosing policy p_3 .

3.6 Conclusion

This paper considers a basic setup of Persuasion in which the receiver, after observing the sender's experiment, can gather his own information. Information gathering is costly for the receiver, but it is not for the sender. I assume that the receiver behaves like a rational inattentive decision maker, with an exogenous cost parameter. We take this cost parameter as the ability to gather information. On the one hand, I show that the sender cannot be worse off if the receiver's ability of gathering information is reduced. On the other hand, I show that the receiver is not necessarily better off if his ability of gathering information increases. That is, my results show that, as it happens with information itself, the ability of gathering information is not always desirable in strategic environments.

To provide the results, I look for an optimal splitting of posteriors to solve the sender's problem as well as receiver's problem. This method, was originally proposed by Aumann and Maschler (1967 [4], 1968 [5] and 1995 [9]), and it allows the analyst to use graphical tools to solve any instance of the model.

¹⁹The sender is actually indifferent between policy p_1 and p_3 , but he knows that when the bad state is more likely, the policy maker would never choose p_1

Chapter 4

Canonical Equilibria in Games with Flexible Information Acquisition

Abstract

I study costly information acquisition by multiple decision makers facing a strategic environment. First, each agent chooses an arbitrary *signal* about an unknown state of nature, and she incurs in a cost that is increasing in the amount of information contained in the signal. Then, once information has been acquired, agents choose their actions in a Bayesian game. I show that for every equilibrium there exists an *equivalent* pure strategy Bayes Nash equilibrium in which the space of signals is equal to the action space. The existence of such equilibrium, which I call *canonical equilibrium*, provides a *revelation principle* for games with flexible information acquisition. I also prove existence of equilibrium, and I extend some results from single-agent costly information acquisition to strategic environments.

Keywords: Canonical equilibrium, revelation principle, purification, rational inattention, games with endogenous information.

JEL Classification: D81, D82, D83.

4.1 Introduction

Many economic situations entail strategic decision making under uncertainty. In such cases, agents decide not only which action to take, but also how much to learn about the environment. For instance, firms need to learn about market conditions before entering into a new market. They can obtain information about consumers' demand, their cost of production and the technologies used by the other firms in the market. Bidders on EBay can collect information about the product before making their bids. They may learn about the durability of the product, its handiness and after-sell services. Finally, voters in an election may want to learn about the policies proposed by different politicians.

In this paper I study costly information acquisition by multiple agents that face a strategic environment. Players share a common prior belief about an unknown state of nature, and they freely choose a costly *statistical experiment* that provides information about the state.¹ The cost the players face is increasing

¹Experiments consist on a realization space (the signals) and a map that goes from states of nature to probability distributions over signals.

in the reduction of uncertainty the experiment provides. After each player designs his information, they choose strategies in the Bayesian game.

Studying games of incomplete information with flexible information acquisition presents several technical difficulties. First, the realization space of the signals can be anything, from a finite alphabet to a non-separable infinite dimensional space. This alone brings about a problem of existence, because the space of probability distributions over the realization space is potentially not compact. The second technical difficulty is related to the computation of equilibrium. In particular, for an equilibrium to exist, it may be necessary for the players to randomize over actions (after the observation of a signal) or over experiments. Hence, the problem of finding an equilibrium may become a daunting task.

The contribution of the paper is threefold. First, I show that for every equilibrium in a game with flexible costly information acquisition there exists an *equivalent* pure strategy equilibrium in which the space of signals is equal to the action space.² This pure strategy equilibrium entails no randomization over actions nor over experiments, and I call it *canonical Nash equilibrium*. Second, I show that a canonical equilibrium always exists. Finally, I extend some standard results from single-agent information acquisition to strategic environments.

The existence of canonical equilibria in games with flexible information acquisition is appealing for at least three reasons. First, it provides theoretical foundations to simplify the analysis of situations that, in principle, could be very difficult to tackle. Canonical equilibria is an equilibrium where signals recommend *pure* actions and those actions are followed by the players. Hence, it provides a *revelation principle* for games with flexible information acquisition. Second, existence of canonical equilibrium provides a simple setup to extend standard results on optimal single-agent information extraction to strategic environments. By focusing on *mutual information* as the functional form of the cost of information, I am able to provide equilibrium (probabilistic) choices and learning strategies. I also extend the *posterior based approach* of Caplin and Dean (2013, [21]), which provides geometrical tools to solve instances of the model.³ Finally, existence of canonical equilibria provides an additional response to the criticism of mixed strategy equilibrium.⁴ In particular, it is argued that mixed strategy equilibrium plays no role in behavioral theory beyond guaranteeing existence of equilibrium, because one does not observe people adopting mixed strategies in practice. Canonical Nash equilibrium is a pure strategy equilibrium, both in *experiments* and in *behavioral strategies*.

Related Literature. This paper relates to three research areas in economics. First, it fits into the rational inattention literature, particularly into the study of games with rational inattentive players.⁵ At the heart of the rational inattention literature is the assumption that agents face a cost of information acquisition that is proportional to the expected reduction of entropy of the beliefs induced by the experiment. For instance, Yang (2015, [83]) studies global games and provides a characterization of optimal experiments. He shows that multiplicity of equilibrium re-appear in coordination games when the the cost is low. Fixing as well the functional form of the cost function, Ravid (2017, [74]) studies a two players bargaining game, Martin (2017, [59]) studies strategic pricing in a game with a seller and a rational inattentive buyer, and

²Two Nash equilibria are *equivalent* if the joint distribution between actions and states are the same.

³This approach, developed for single agent problems, proposes to look at the *net benefit* generated by each action/posterior pair, which corresponds to the expected utility of using that action at that posterior minus the information cost associated with that posterior. The most general insight given by the *posterior based approach* is that the maximum attainable benefit for a player is given by the concave closure of the upper envelope of the net benefit functions. This method is analogous to the *concavification* method of Aumann (1995, [9]).

⁴For a discussion and replies to such critique, see for instance Harsanyi (1973, [40]), Aumann et al. (1983, [8]), Radner and Rosenthal (1982, [72]), Milgrom and Weber (1985, [66]).

⁵Landmark papers in rational inattention with a single decision maker are Sims (2003, [77]), Matejka and McKay (2014,[62]), Caplin and Dean (2013, [21]) and Steiner et al. (2018, [79]).

Denti (2015, [29]) studies potential games. In this paper we do not look at a *specific* strategic environment, but arbitrary finite games. Furthermore, for the first two main contributions of the paper (Theorem 4.1 and 4.2), I do not assume a particular functional form of the cost function, but I keep four key properties that are also satisfied by entropy based cost functions.

Second, this paper fits into the literature of information design. As noted in Bergemann and Morris (2019, [17]), information design problems can be either *literal* or *metaphorical*. In the former, a real information designer provides information to players in order to serve his personal interest.⁶ In the latter, the formulation as an information design problem is a metaphor that the analyst wields as a tool.⁷ This paper relates with the second interpretation of information design by studying games in which the information structure results from independent interaction among players. I address the question of whether we can restrict attention, without loss of generality, to information structures in which every player's signal *recommends* an action and that action is followed. That is, if we have a *revelation principle* guaranteeing no loss of generality on looking at *direct* information structures.

Counting on a revelation principle is a first step towards simplifying situations that could be very difficult to study. For instance, Taneva (2019, [80]) studies a self interested information designer who chooses the information structure that governs players' private information in a strategic situation. She shows that the information designer can rely on direct information structures, and uses this simplification to characterize the designer's optimal behavior. Nevertheless, a revelation principle is not always available. Lipnowski and Mathevet (2018, [56]) studies an information designer who designs information for an agent whose payoff depends not only on his action and the state but also on his beliefs.⁸ They show that the revelation principle may not hold when agents do not value information only for instrumental purposes (i.e. only for decision making). In our setup, player's net (of information cost) payoffs depends not only on actions and the state, but also on posterior beliefs. Furthermore, since players interact in a non cooperative game, randomization -either over experiment or actions- may be necessary to sustain equilibrium. Hence, it is not clear whether a revelation principle is always available. This paper provides such result by showing existence of canonical Nash equilibrium.

Finally, this paper relates to the purification literature. Started by Harsanyi (1973, [40]), this area studies the conditions under which a pure strategy equilibrium can be obtained from a mixed strategy equilibrium in the context of games of incomplete information.⁹ When information in a game satisfy certain technical conditions, it is possible to incorporate the randomization of the mixed strategy into the information itself, in a way in which players choose pure strategies after the observation of each signal.¹⁰ This type of results are called *purification theorems*. This literature focuses on exogenous information structures. Since canonical equilibrium *is* a pure strategy equilibrium, this paper is the first one to provide a purification result for games with endogenous information.

The rest of the paper is organized as follows. Section 4.2 presents the model. Section 4.3 studies equilibria and states the main results. Section 4.4 extends results from single-agent costly information extraction problems to strategic environments. Section 4.5 presents some examples and Section 4.6

⁶See for instance Brocas and Carrillo (2007, [20]), Rayo and Segal (2010, [75]), Kamenika and Gentzkow (2011, [31]), (2014, [32]) and (2017, [33]), Alonso and Camara (2015, [1]) and (2016, [2]), Treust and Tomala (2019, [52]), Koessler et al. (2018, [48]), Taneva (2019, [80]) Bergemann and Morris (2016a, [16]).

⁷See for instance Bergemann et al. (2015, [18]), Bergemann and Morris (2016b, [15]), Mathevet et al. (2020, [63]).

⁸This types of games are often called *Psychological games*

⁹Seminal papers on purification theorems and the interpretation of mixed strategies are Harsanyi (1973, [40]), Aumann et al. (1983, [8]), Radner and Rosenthal (1982, [72]), Milgrom and Weber (1985, [66]), Rubinstein (1991, [76]), Balder (2008, [12]) and Kahn and Rath (2009, [46]).

¹⁰See for instance Khan et.al (1999, [47]), Loeb and Sun (2006, [57]), Podczeck (2009, [70]) and Wang and Zhang (2012, [82]). See also Greinecker and Podczeck (2015, [39]) for review and discussion.

concludes.

4.2 The Model

4.2.1 The Bayesian Game

I consider a finite set of players $I = \{1, 2, \dots, N\}$ and a finite set of possible states K . Players have a common prior $\mu \in \Delta(K)$. Each player i has a finite action space A_i , and the payoff function is $g : K \times A \rightarrow \mathbb{R}^N$, where $A = \prod_i A_i$. A *basic game* is a five-tuple $G = (I, K, A, g, \mu)$. An *information structure* $\pi = (\pi_i)_i$ consists on a statistical experiment $\pi_i = (S_i, (\pi_k^i)_k)$ for each player i . Specifically, a statistical experiment consists on a realization space S_i along with a family of probability distributions $(\pi_k^i)_k$ that specifies the probability of receiving signals $s_i \in S_i$ for each state of nature k . I assume that statistical experiments only provide information about k , so signals are independent given the state.

The basic game G along with the information structure π define a standard incomplete information game $\Gamma(G, \pi)$ as follows.¹¹ First, the state k is drawn according to μ and $s_i \in S_i$ is drawn according to $\pi_k^i(\cdot)$. Then each player is informed about $s_i \in S_i$ and beliefs $q_{s_i} \in \Delta(K)$ are updated using Bayes' rule, so that $q_{s_i}(k) = \frac{\pi_k^i(s_i)\mu(k)}{\pi_i(s_i)}$ for every k , where $\pi_i(s_i) \equiv \sum_k \pi_k^i(s_i)\mu(k)$ denotes the unconditional probability of signal realization s_i . Finally, every player i chooses a mixed strategy $\sigma_i \in \Sigma_i$, where $\Sigma_i = \Delta(A_i)$.

A behavioral strategy for player i in $\Gamma(G, \pi)$ is a map $\beta_i : S_i \rightarrow \Sigma_i$. Let B_S denote the set of all behavioral strategies with domain S . Given a profile $\beta = (\beta_i)_i$ and an information structure π , the expected payoff of each player is written in the usual way, and a Nash equilibrium is a profile β such that no player can increase his expected payoff by changing his own behavioral strategy alone.

4.2.2 The Attention-Move Game

I consider now a game in which each player i chooses his experiments π_i , and then play the game $\Gamma(G, \pi)$. Furthermore, I assume that players can randomize over different experiments. That is, they can choose a collection $(\pi_{h_i})_{h_i \in H_i}$ of experiments, and a probability distribution $\sigma_i \in \Delta(H_i)$. This *mixture of experiments* selects, independently of the state, an experiment h_i from H_i with probability $\sigma_i(h_i)$ and then performs the experiment as usual.¹² For every possible experiment chosen by σ_i , player i has associated a behavioral strategy. With a slight abuse of notation, I let $\pi_i \equiv (\pi_{h_i})_{h_i}$ and $\beta_i \equiv (\beta_{h_i})_{h_i}$, so that player i chooses $(\sigma_i, \pi_i, \beta_i)$.

4.2.2.1 Payoffs

A profile of mixture of experiments along with a profile of behavioral strategies induces a joint distribution over actions as states. Let

$$\varphi_{h_j}(a_j|k) \equiv \int_{s \in S_{h_j}} \beta_s^{h_j}[a_j] d\pi_k^{h_j}(s) \quad (4.1)$$

be the total probability of action a_j in state k when experiment h_i is selected by σ_i . In 4.1, $\beta_s^{h_j}[a_j]$ denotes the probability that mixed strategy $\beta_{h_j}(s)$ attaches to action a_j . Consequently, the total probability of action a_j in state k is given by

$$\varphi_j(a_j|k) \equiv \sum_{h_j} \sigma_i(h_j) \varphi_{h_j}(a_j|k) \quad (4.2)$$

¹¹See Gossner (2000, [38])

¹²See Torgersten (1995).

Since signals are independent given the state, the probability of profile a in state k and the joint probability of profile a and state k are respectively $\varphi(a|k) \equiv \prod_j \varphi_j(a_j|k)$ and $\varphi(a, k) \equiv \varphi(a|k)\mu(k)$. I also let $\varphi_{-i}(a_{-i}|k) \equiv \prod_{j \neq i} \varphi_j(a_j|k)$ and $\varphi_{-i}(a_{-i}, k) = \varphi_{-i}(a_{-i}|k)\mu(k)$. Although φ depends on σ, π and β , to keep notation simple I will refer to $\varphi(\sigma, \pi, \beta)$ only when necessary, and omit the extra notation when no confusion arises.

The joint distribution between actions and states is a key ingredient in all results in the next section. Indeed, if two different profiles of behavioral strategies and mixture experiments share the same distribution φ , then, both are statistically indistinguishable: an outside observer would not be able to distinguish between the two by systematically observing ex post action profiles. Hence, we say that two profiles (σ, π, β) and (σ', π', β') are *equivalent* if $\varphi(\sigma, \pi, \beta) = \varphi(\sigma', \pi', \beta')$.

When experiment h_i is selected by σ_i , the expected payoff of player i is given by

$$E^{h_i}[g_i|\beta, \pi] \equiv \sum_{k,a} g_i(a, k) \varphi_{h_i}(a_i|k) \varphi_{-i}(a_{-i}, k) \quad (4.3)$$

I assume that players choose simultaneously their mixture of experiments, and that they do not observe the mixture of experiments of the other players when choosing their behavioral strategies.¹³

4.2.2.2 The Cost of Information

I also assume that experiments are costly. That is, the cost for player i of performing experiment π_i is a function $c_i : \Delta(S_i)^{|K|} \times \Delta(K) \rightarrow \mathbb{R}$. I assume that the information cost functions satisfy four properties. First, I assume that no information has zero cost. This is a standard normalization. Second, I assume that, for a fixed realization space S , c is continuous in $\pi \in \Delta(S)^{|K|}$. Continuity in experiments is necessary to have continuous objective functions for every player, which is necessary to show existence.

Third, I assume that the cost of information is quasi convex in experiments. Formally, fix a realization space S , and let $\lambda \in [0, 1]$. Then, for all $\pi_1, \pi_2 \in \Delta(S)^{|K|}$ and for all $\lambda \in [0, 1]$, $c(\lambda\pi_1 + (1 - \lambda)\pi_2, \mu) \leq \max\{c(\pi_1, \mu), c(\pi_2, \mu)\}$. Quasi-convexity guarantees that when two experiments π_1, π_2 are optimal, the convex combination of the two experiments is optimal as well. This assumption is not bold. It has been shown in Pomatto et al. (2019) that cost functions satisfying quite general properties (such as continuity and Blackwell informativeness) are convex functions, which are also quasi-convex.¹⁴

Finally, I assume that, at a given prior, more information (in Blackwell sense) is more expensive. Formally, let $q \in \Delta(K)$ be the uniform prior assigning equal probability to each element of K , and let π_1 and π_2 be two experiments inducing distributions (over posteriors) ϕ_1 and ϕ_2 under the uniform prior. Then, π_1 is more informative than π_2 in the Blackwell sense, denoted by $\pi_1 \succeq \pi_2$, if

$$\int_{\Delta(K)} f(p) \phi_1(p) \geq \int_{\Delta(K)} f(p) \phi_2(p) \quad (4.4)$$

for every convex function $f : \Delta(K) \rightarrow \mathbb{R}$. Furthermore, an experiment π_1 is said to be strictly more informative than π_2 (denoted $\pi_1 \triangleright \pi_2$) if $\pi_1 \succeq \pi_2$ and $\pi_2 \not\succeq \pi_1$.¹⁵ I assume that if $\pi_1 \succeq \pi_2$, then $c(\pi_1, \mu) \geq c(\pi_2, \mu)$, and that if $\pi_1 \triangleright \pi_2$, then $c(\pi_1, \mu) > c(\pi_2, \mu)$.

¹³Note that if I were to assume otherwise, an equilibrium would still be a profile (σ, π, β) , and the results of next section would all hold. However, in such a case, we would need to check whether an equilibrium is sequential, and if a sequential equilibrium always exists.

¹⁴The other two properties that cost functions have to satisfy are (i) additivity with respect to independent experiments and (ii) linearity with respect to a probability of success of the experiment. See Pomatto et al. (2019) for details.

¹⁵See Lehmann (1988, [53]).

4.2.2.3 Equilibrium

The game in which players choose mixture of (costly) experiments and then play the subsequent Bayesian game is called the attention-move game, which I denote by $\Lambda(G, c)$, where $c = (c_i)_i$. In the attention-move game $\Lambda(G, c)$, given a profile of strategies $(\sigma_{-i}, \pi_{-i}, \beta_{-i})$ each player i solves

$$\begin{aligned} & \max_{\{\sigma_i, \pi_i, \beta_i\}} \sum_{h_i} \sigma_i(h_i) (E^{h_i}[g_i | \beta_{h_i}, \pi] - c(\pi_{h_i}, \mu)) \\ & \text{s.t.} \\ & \sigma_i \in \Delta(H_i), \quad \beta_{h_i} \in B_{S_{h_i}} \quad \forall h_i \\ & \pi_{h_i} \in F_i \equiv \left\{ (S_{h_i}, (\pi_k^{h_i})_k) : \pi_k^{h_i} \in \Delta(S_{h_i}) \quad \forall k \right\} \quad \forall h_i \end{aligned} \tag{4.5}$$

An equilibrium of the game $\Lambda(G, c)$ is a profile (σ, π, β) such that for every player i , $(\sigma_i, \pi_i, \beta_i)$ solves (4.5). An equilibrium in $\Lambda(G, c)$ is said to be *pure in experiments* if σ_i is degenerate. An equilibrium (σ, π, β) is said to be *pure in behavioral strategies* if $\beta_i(s_i)$ is a unit vector for all i and s_i . An equilibrium in $\Lambda(G, c)$ is said to be *pure* if it is pure in experiments and pure in behavioral strategies. That is, each player does not randomize between experiments, and plays a pure strategy after the realization of each signal.

An equilibrium profile (σ, π, β) is a *canonical Nash equilibrium* of $\Lambda(G, c)$ if it is pure and the set of signals is equal to the set of actions. That is, $S_i = A_i$ and $\beta_i = \text{id}_{A_i}$, so the map β_i is the identity function over the set of actions.

It is important to note that the set of Nash equilibrium of the attention move game is not the whole set of correlated equilibrium as defined by Bergemann and Morris (2019). Since in the attention move game signals are independent given the state, not all correlated equilibria are available. Although the model supports equilibria that display correlation beyond the payoff relevant states of nature (see example 3 in section 4.5), high coordination can only be reached when the cost is low enough.¹⁶ Furthermore, even if certain correlation can be produced using payoff irrelevant states, candidate experiments must be part of the equilibrium as well. This requirement rules out some available correlated equilibrium (see example 4 in section 4.5).¹⁷ To introduce correlation of signals given the state of nature one could use a public randomization device freely available to every player.

4.3 Existence of Canonical Nash Equilibrium

4.3.1 Main Results

In this section I provide the main results of the paper. Theorem 4.1 asserts that for every equilibrium there exists a canonical Nash equilibrium that induces the same joint distribution between actions and states.

Theorem 4.1. *For every Nash equilibrium profile $(\hat{\sigma}, \hat{\pi}, \hat{\beta})$ in the game $\Lambda(G, c)$ there exists a canonical Nash equilibrium $(\tilde{\sigma}, \tilde{\pi}, \tilde{\beta})$ such that $\varphi(\hat{\sigma}, \hat{\pi}, \hat{\beta}) = \varphi(\tilde{\sigma}, \tilde{\pi}, \tilde{\beta})$.*

This result identifies equivalence classes of Nash equilibria. Furthermore, all Nash equilibria with the same canonical representation are indistinguishable from the point of view of an outsider observing states

¹⁶It would be possible to explore correlation further by considering that players have available a public costless correlation device. The public correlation device along with a payoff relevant state could be a starting point to study correlation explicitly.

¹⁷In example 4 I show that experiments producing payoff outside of the convex hull of the Nash equilibrium of the chicken game are not part of any equilibrium profile.

and actions. This equivalence allows me to focus on canonical Nash equilibrium to show existence of Nash equilibrium, which is the second main result of this paper.

Theorem 4.2. *Every game $\Lambda(G, c)$ has a Nash Equilibrium.*

In what follows, I show the main properties of equilibrium and provide the proof of Theorem 4.1. Then, using Theorem 4.1, I provide a simple proof for Theorem 4.2.

4.3.2 Properties of Equilibrium

To proof Theorem 4.1 I need four lemmas. Lemma 4.1 states that for every equilibrium of the game $\Lambda(G, c)$ there exists an equivalent equilibrium that is pure in experiments. Lemmas 4.2 and 4.3 together establish that we can focus on equilibrium information structures in which each signal realization gives a recommendation of what to do. Since recommendations can be mixed strategies, Lemma 4.4 provides a purification of such mixed strategies recommendations.

Lemma 4.1. *For every Nash equilibrium $(\hat{\sigma}, \hat{\pi}, \hat{\beta})$ there exists an equivalent equilibrium $(\hat{\sigma}, \hat{\pi}, \hat{\beta})$ that is pure in experiments.*

Proof. See Appendix. □

When a player is randomizing his choice of an experiment, he could also be choosing a single experiment that contains the same information and it is weakly cheaper than the expected cost of such randomization. Since every randomization among experiments can be replaced by a single composed experiment, in what follows we consider only equilibria that are pure in experiments and omit the randomization σ from the notation.

Lemma 4.2 shows that for signals giving different information, different mixed strategies must be followed by the players. Moreover, these mixed strategies must not only be different, but also have disjoint support. Formally, I say that strategy σ is played with positive unconditional probability if there exists a pair (s_i, k) such that $\beta(s_i) = \sigma$ and $\pi_k^i(s_i), \mu(k) > 0$. For a given σ , I identify the set of signals that recommend mixed strategies whose support is not disjoint from that of σ . That is, let $x_i(\sigma) \equiv \{s_i \in S_i : \text{supp}\{\beta(s_i)\} \cap \text{supp}\{\sigma\} \neq \emptyset\}$. Then, the result goes as follows.

Lemma 4.2. *Let $(\hat{\pi}, \hat{\beta})$ be a Nash equilibrium of $\Lambda(G, c)$ in which, for some player i , strategy σ is played with positive unconditional probability. Then, there exists a $\tilde{q} \in \Delta(K)$ such that $q_{s_i} = \tilde{q}$ for all $s_i \in x_i(\sigma)$.*

Proof. See Appendix. □

In words, Lemma 4.2 says that all signals that recommend mixed strategies that share elements in their support should all generate the same posterior. The proof relies on the Blackwell-informative consistency of the cost function and in the best respond condition of Nash equilibrium. In particular, suppose that two signals recommend mixed strategies that share at least one element in their support. Then, the best response condition allows us to replace such strategies by a pair that puts all the probability in one of the shared element of the support without changing the payoff of the player. Then, we obtain two signals recommending the same action. Finally, when no different actions are induced by two signal realizations, it must be true that the associated posteriors are the same. When the same action is induced, the same expected benefit is obtained by replacing such split of posteriors by a single average posterior, and the new experiment would yield a strictly lower cost.

Lemma 4.2 does not rule out the possibility of having different signals that recommend the same mixed strategy as long as they provide the same information. However, Lemma 4.3 (below) states that

whenever two signals (that provide the same information) recommend the same strategy, there exists another equilibrium in which (i) there is only one signal for each strategy, and (ii) the joint probability of action and states remains unchanged. The idea of the proof for this lemma is similar to that of proposition 1 in Kamenica and Gentzkow (2011, [31]). Let $y(\sigma) \equiv \{s_i \in \hat{S}_i : \hat{\beta}(s_i) = \sigma\}$ be the set of signals that recommend mixed strategy σ . Then, the result goes as follows.

Lemma 4.3. *Let $(\hat{\pi}, \hat{\beta})$ be a Nash equilibrium of $\Lambda(G, c)$ in which, for some player i , mixed strategy σ is played with positive unconditional probability, and that $|y(\sigma)| > 1$. Then, there exists another Nash Equilibrium $(\tilde{\pi}, \tilde{\beta})$ such that $|y(\tilde{\beta}(s_i))| = 1$ for all $s_i \in \hat{S}_i$, and $\varphi(\hat{\pi}, \hat{\beta}) = \varphi(\tilde{\pi}, \tilde{\beta})$.*

Proof. See appendix. □

To complete the proof of Theorem 4.1, I need to take care of the fact that when a signal realization is uniquely paired with a mixed strategy (as it happens by Lemma 4.2 and 4.3), it is, by construction, recommending several pure actions. Lemma 4.4 shows that whenever a signal recommends a mixed strategy, the mixed strategy can be incorporated into the information structure so that each signal realization recommends a pure strategy. This can be done without changing the joint probability of action and states. That is, we can incorporate the randomization of the mixed strategy into the experiment itself.

Lemma 4.4. *Let $(\hat{\pi}, \hat{\beta})$ be a Nash equilibrium of $\Lambda(G, c)$ in which, for some player i , strategy σ is played with positive unconditional probability, and that $|\text{supp}\{\sigma\}| > 1$. Then, there exists a pure strategy equilibrium $(\tilde{\pi}, \tilde{\beta})$ such that $\varphi(\hat{\pi}, \hat{\beta}) = \varphi(\tilde{\pi}, \tilde{\beta})$.*

Proof. See Appendix. □

The proof is by construction and it relies on the fact a when a mixed strategy σ is induced by a signal realization s_i , it is possible to construct an equally informative experiment $\tilde{\pi}_i$ (hence, with the same cost), in which s_i is replaced by a set of new signals whose distribution, conditional on k , mimics the distribution of σ . A key ingredient of the proof is the independence of signals conditional on the state. Otherwise, changing signals may change equilibrium information strategies of the other players, changing the whole distributions of actions and states. This procedure is illustrated in example 2. The spirit of the argument is analogous to that of Greinecker and Podczeck (2015, [39]), in the sense that players can use randomization devices that do not provide any further information. This is also in line with Aumann (1964, [6]; 1974, [10]), where mixed strategies are implemented by using randomization devices.

4.3.3 Proof of Theorem 4.1

In what follows, I show how Lemmas 4.1-4.4 together imply that any Nash equilibrium can be transformed into a canonical Nash equilibrium such that the joint distribution of actions and states are the same.

Proof. Suppose $(\hat{\pi}, \hat{\beta})$ is a Nash Equilibrium in $\Lambda(G, c)$ that is pure in experiments (Lemma 4.1). By Lemma 4.2, for $(\hat{\pi}, \hat{\beta})$ to be a Nash equilibrium it has to be the case that, for every player i , each $s_i \in \hat{S}_i$ is associated with one strategy $\hat{\beta}(s_i)$ that shares no support with any other induced strategy. That is, if an action a_i is played with positive probability under $\hat{\beta}(s_i)$, for all $s'_i \neq s_i$, $\hat{\beta}_{s'_i}[a_i] = 0$. Hence, signals in \hat{S}_i generate a partition of A_i .

By Lemma 4.3, we can convert Nash equilibrium $(\hat{\pi}, \hat{\beta})$ into a new Nash equilibrium $(\tilde{\pi}, \tilde{\beta})$ in which redundant signals are packed into one single signal and $\varphi(\hat{\pi}, \hat{\beta}) = \varphi(\tilde{\pi}, \tilde{\beta})$. Hence, in this new equilibrium, there is one signal associated to each element of the partition of A_i .

Then, by Lemma 4.4, each $\check{\beta}(s_i)$ can be incorporated into an experiment in which information-irrelevant signals are incorporated to mimic the distribution of $\check{\beta}$ over each element in the partition of A_i . Hence, there exists a pure strategy equilibrium $(\check{\pi}, \check{\beta})$ such that $\varphi(\check{\pi}, \check{\beta}) = \varphi(\tilde{\pi}, \tilde{\beta})$.

Finally, since Lemma 4.4 is applied to each player i and each element of the partition of A_i generated by $(\hat{\pi}, \hat{\beta})$, each $s_i \in \tilde{S}$ is uniquely associated to an action, and \tilde{S}_i can be taken to be A_i . Hence, the new Nash equilibrium $(\tilde{\pi}, \tilde{\beta})$ is canonical, and $\varphi(\hat{\pi}, \hat{\beta}) = \varphi(\tilde{\pi}, \tilde{\beta})$, which concludes the proof. \square

4.3.4 Proof of Theorem 4.2

When restricting attention to canonical experiments I can use simple tools to show existence of equilibrium. First, I use the Berge Maximum Theorem¹⁸ to show upper hemi-continuity of the solution correspondence, and then, Kakutani Fixed Point Theorem¹⁹ to show existence.

Proof. For a canonical Nash equilibrium, each player chooses a family of probability distributions over the finite set A_i , so $\pi_i \in \hat{F}_i \equiv \{((\pi_k^i)_k, A_i) : \pi_k^i \in \Delta(A_i) \forall k\}$. Additionally, behavioral strategies are equal to the identity. Let $V_i(\pi) \equiv E[g_i|\pi] - c(\pi_i, \mu)$, and note that the set \hat{F}_i corresponds to $(\Delta^{|A_i|-1})^{|K|}$, where $\Delta^{|A_i|-1}$ denotes the simplex of dimension $|A_i| - 1$. Hence, \hat{F}_i is non empty, compact and convex, and it does not depend on π_{-i} .

The objective function $V_i : \hat{F}_i \times \prod_{j \neq i} \hat{F}_j \rightarrow \mathbb{R}$ is continuous in π_k^j for every $j \in I$ and $k \in K$. Furthermore, $V_i(\pi_i, \pi_{-i})$ is the sum of a linear ($E[g_i|\pi]$) and concave ($-c(\pi_i, \mu)$) function of π , so it is concave function of π . Then, the Berge's maximum theorem implies that the correspondence defined by $\psi_i(\pi_{-i}) := \operatorname{argmax}_{\pi_i \in \hat{F}_i} V_i(\pi_i, \pi_{-i})$ is a non empty convex valued and upper hemi-continuous correspondence.

Finally, $\psi : \prod_j \hat{F}_j \rightarrow \prod_j \hat{F}_j$, defined as $\psi = (\psi_i)_i$, is a non-empty, compact and convex valued upper hemi-continuous correspondence, so by Kakutani Fixed Point Theorem, there exists a π such that $\pi \in \psi(\pi)$, and the equilibrium exists. \square

4.4 Rational Inattention in Strategic Environments

4.4.1 Rational Inattention and Behavioral Implications

Costly information acquisition has been widely studied by the rational inattention literature, mainly in the context of a single decision maker. By assuming a particular functional form of the cost function, this literature has been able to provide not only a solution to the problem of the *rational inattentive* decision maker but also testable behavioral implications of such models. In this section I provide some extensions of what has been established so far.

In particular, I provide the analogous optimal *conditional choice probabilities* for games of incomplete information and the associated learning strategies. Both are obtained, for each agent, given a profile of conditional choice probabilities of the other players. Finally, I recover some graphical tools to help the analyst to find optimal learning strategies. In particular, I extend the posterior based approach of Caplin and Dean (2013, [21]) by looking at the problem of each player as that of an optimal splitting of posteriors. Hence, we give the analyst the techniques introduced by Auman and Maschler (1967, [4]; 1968, [5]; 1995, [9]) and Kamenica and Gentzkow (2011, [31]) to solve instances of the model.

¹⁸See Leininger (1984,[54]).

¹⁹See Kakutani (1941,[45]).

4.4.2 Main Results

The rational inattention literature assumes that the cost of information acquisition is proportional to the *mutual information* between the signals and the state. Mutual information is a measure of the amount of information about the state of nature contained in the signals. Hence, the cost for player i is defined as

$$c(\pi_i; \mu, \lambda) = \lambda_i \left[H(\mu) - \sum_{s_i} H(q_{s_i}) \pi(s_i) \right] \quad (4.6)$$

where $H(\cdot)$ is the Shannon entropy function,²⁰ and λ_i is the unit cost of information for player i that is assumed to be common knowledge. The cost is linear in the expected reduction of entropy of the state of nature given by the experiment.

Note that the standard rational inattentive cost function satisfies the three basic properties we imposed in section 4.2.2. That is, experiments that convey no information have zero cost, the cost of information is convex in experiments, and more information (in Blackwell sense) is more expensive. Additionally, this cost function is differentiable and strictly convex on experiments, and it has unbounded derivatives at the extreme points of its domain, so an interior solution is always guaranteed.²¹

Theorem 4.3. *Fix a prior belief $\mu \in \Delta(K)$, a profile of experiments π_{-i} for players other than i , and a cost parameter $\lambda_i > 0$. Then, optimal experiments of player i satisfy*

$$\pi_i(a|k) = \frac{\pi_i(a) \exp\left(\frac{1}{\lambda_i} E[g_i(a, a_{-i}, k)|k]\right)}{\sum_{a'} \pi_i(a') \exp\left(\frac{1}{\lambda_i} E[g_i(a', a_{-i}, k)|k]\right)} \quad (4.7)$$

If action a is chosen with positive probability under rule π_i , then optimal posterior beliefs satisfy

$$q_a(k) = \frac{\exp\left(\frac{1}{\lambda_i} E[g_i(a, a_{-i}, k)|k]\right) \mu(k)}{\sum_{a'} \pi_i(a') \exp\left(\frac{1}{\lambda_i} E[g_i(a', a_{-i}, k)|k]\right)} \quad (4.8)$$

where the expectation is taken under $\varphi(a_{-i}|k)$.

Proof. The proof equation 4.7 follows directly from Theorem 1 in Matejka and McKay (2014,[62]) by replacing the payoff of action a in state of nature k by the expected payoff conditional on k given by a profile of experiments of the other players. This can be done because signals are assumed to be independent across players for a given state of nature. The proof of equation 4.8 follow directly from Theorem 1 in Caplin and Dean (2013, [21]) and Caplin et al. (2017, [24]) by (again) replacing the payoff of action a in state of nature k by the expected payoff conditional on k given by a profile of experiments of the other players. \square

A corollary of Theorem 4.3 is that, whenever two actions are chosen with positive probability at equilibrium, the ratio between the log difference in posteriors beliefs and the difference in utilities across

²⁰The Shannon entropy of a discrete random variable with realization space X and probability distribution p is given by $H(p) = -\sum_{x \in X} p(x) \log(p(x))$, and it is a measure of the uncertainty of x . I consider the natural logarithm, so entropy is expressed in nats, and we use the convention that $0 \cdot \log(0) = 0$, which is justified by the continuity of $H(\cdot)$.

²¹To establish the second part of the Theorem below and the concavification approach developed on section 4.4.3 we only need the cost functions satisfy the three basic properties discussed and to be posterior separable. There is no need to use mutual information as the functional form of the cost. A cost function W is said to be posterior separable if it can be written as

$$W(\pi, \mu) = \left[F(\mu) - \int_{s_i} F(q_{s_i}) d\pi(s_i) \right]$$

where $F : \Delta(K) \rightarrow \mathbb{R}$ is a strictly convex function. Mutual information is a particular instance of such cost functions.

chosen actions is constant across states. That is, by taking the ratio between the posteriors of two actions chosen with positive probabilities, one obtain

$$\frac{q_{a_i}(k)}{q_{a'_i}(k)} = \exp\left(\frac{1}{\lambda} E[g(a_i, a_{-i}, k) - g(a'_i, a_{-i}, k) | k]\right) \quad \forall k \in K \quad (4.9)$$

This property, called invariant likelihood ratio, pins down the rate at which choice mistakes respond to changes in the cost of those mistakes.²²

4.4.3 Concavification Technique to Find a Solution

To see how we can use the concavification techniques developed in the literature on information design, we need to take a deeper look at the posterior based approach proposed by Caplin and Dean (2013, [21]) for a single decision maker. This approach looks at the problem of a single rational inattentive decision maker as that of a splitting of posteriors. Indeed, recall that Bayesian updating implies that every experiment can be represented by a collection of posteriors that average up the prior.²³ That is, if $\hat{\eta} = (S, (\hat{\eta}_k)_k)$ is an experiment performed under prior p , then we can look at

$$\eta \in \mathcal{F}(p) \equiv \left\{ (S, (q_s, \eta_s)_s) : q_s \in \Delta(K) \forall s, (\eta_s)_s \in \Delta(S) \text{ and } \sum_s q_s \eta_s = p \right\} \quad (4.10)$$

This approach relies fundamentally on the fact that optimal *rational inattentive* behavior is inconsistent with the choice of the same action at two different posteriors. The key to extend the approach to games is to note that the approach is consistent with the choice of two different actions at the same posterior. Hence, we can use Theorem 4.1 and write down the *net benefit* generated by pairs actions/posterior, which corresponds to the expected benefit of using that action at that posterior minus the information cost associated with that posterior. Recall that net benefit functions for player i depend on the experiments of the other players, but I will omit such dependency to keep notation simple. Then, the *net benefit* of posterior q_{a_i} under action a_i is defined as

$$N_{a_i}(q_{a_i}) \equiv \sum_{k, a_{-i}} g(a_i, a_{-i}, k) \varphi(a_{-i} | k) q_{a_i}(k) + \lambda H(q_{a_i}) \quad (4.11)$$

Note that here are $|A|$ net benefit functions, which are all bounded and strictly concave. Then, rewrite problem (4.5) as

$$\max_{\{\pi_i \in \mathcal{F}(\mu)\}} \sum_{a_i} N_{a_i}(q_{a_i}) \pi_i(a_i) - \lambda H(\mu) \quad (4.12)$$

where $\lambda H(\mu)$ is a constant that can be ignored for optimization purposes. Hence, it can be seen from program (4.12) that the maximum is attained by some convex combination of certain values of the net benefit functions. Hence, define the upper envelope of the net benefit functions as $N_i(q) \equiv \max_{\{a_i\}} \{N_{a_i}(q)\}$, and the *concave closure* of $N_i(q)$ by $cav(N_i(q)) \equiv \min f(q)$, where the minimum is take over all *concave* functions f such that $f(q) \geq N_i(q)$ for all $q \in \Delta(K)$. The concave closure of a function g is also referred to as the *concavification*²⁴ of g . Then, the following result states that the the value of program (4.12) for player i given prior μ and profile π_{-i} is given by $cav(N_i(\mu)) - \lambda H(\mu)$.

²²Behavioral implications of rational inattention have been studied empirically for a single decision maker in Caplin and Dean (2013, [21]), Caplin and Martin (2014, [23]) and Caplin and Dean (2015, [22]). Additionally, an attempt to test these behavioral implications in strategic environments has been undertaken by Martin (2016, [58]).

²³See Kamenica and Gentzkow (2011, [31]).

²⁴See Aumann and Maschler (1995, [9]).

Theorem 4.4. Fix a prior belief $\mu \in \Delta(K)$, a profile of experiments π_{-i} for players other than i , and a cost parameter $\lambda_i > 0$. The value of problem (4.12) is $U^*(\mu) = \text{cav}(N_i)(\mu) - \lambda H(\mu)$.

Proof. The proof follows directly from lemmas 1, 2 and 3 in Caplin and Dean (2013, [21]), by replacing the payoff of action a in state of nature k by the expected payoff conditional on k given by a profile of experiments of the other players. Once again, this can be done because signals in my model are assumed to be independent across players for a given state of nature. \square

Note that this theorem implies that the *concavification* of $N_i(q_i)$ can be obtained by finding the lowest hyperplane that lies weakly above the net benefit functions and that contains the prior. Optimal posteriors are then given by those who support such hyperplane. Strategies whose net benefit functions are supporting the hyperplane, are the actions that are chosen with positive unconditional probability. An important behavioral implication of this approach in single agent decision problems is that information strategies are locally constant with respect to the prior. That is, changing the prior does not change optimal posteriors. However, it does change conditional choice probabilities. Hence, since the net benefit functions depend on the conditional choice probabilities of the other players, changes in the prior change the shape of the net benefits functions and hence optimal posteriors. That is, optimal learning strategies are no longer constant with respect to the prior.

4.5 Examples

In this section I present five examples that serve both to illustrate the main results of the paper and to highlight the scope and limitations of the model we have studied so far. The first example shows how to obtain a pure Nash equilibrium in a two-person two-action game in which the only equilibrium is in mixed strategies. In this example we consider a payoff irrelevant state of nature that allows to build uninformative experiments that incorporate the randomization of the mixed strategy. The second example shows how to obtain a pure Nash equilibrium in a two-person three-action game of incomplete information in which there is an equilibrium where recommendations are mixed strategies. The construction of such an equilibrium illustrates how Lemma 4.4 works.

The third example studies what happens when the cost function violates Blackwell informativeness. In particular, it is assumed that there is a small fixed cost for the use of every signal. It shows that in such a case, canonical equilibrium cannot be obtained. The fourth example shows how a payoff irrelevant state can be used as a correlation device to coordinate the actions of the players. Hence, this example shows how my model supports the existence of endogenously generated *correlated equilibrium*, as in Aumann (1987, [7]). That is, it shows an equilibrium in which signals are correlated beyond the payoff relevant parameter, and that information is gathered only by coordination purposes. In my model, learning about the correlation device is costly, and it can therefore be interpreted as the cost of coordination.

The last example shows that not all correlated equilibrium can be implemented as an equilibrium in my model. Indeed, as opposed to Aumann (1987, [7]), the use of the correlation device is endogenous, and it has to be part of the equilibrium. Throughout these examples I use mutual information as the cost function.

4.5.1 Example 4.5.1: Matching Pennies

Consider the game of matching pennies, and consider an unknown parameter $k \in K \equiv \{k_0, k_1\}$. Suppose that the prior on K is $\mu = (.1, .9)$, and that λ is 1. Payoffs are presented in table 4.1.

P1/P2	<i>H</i>	<i>T</i>
<i>H</i>	1, -1	-1, 1
<i>T</i>	-1, 1	1, -1

Table 4.1 – Payoffs Example 4.5.1

The unique equilibrium of this game is a profile $(\sigma_1, \sigma_2) = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$. I now build an information structure that implements this equilibrium as a canonical equilibrium of a Bayesian game. Let $S_i = (H, T)$ and $(\pi_{k_0}^i(H), \pi_{k_1}^i(H)) = (\frac{1}{2}, \frac{1}{2})$ for $i = 1, 2$. The cost of such experiment is zero, because it does not provide any information about k . Additionally, both players are happy to play H upon observing H , because they know the other player is randomizing, with half probability in both states of nature. Hence, behavioral strategies are the identity, and this equilibrium is canonical.

This example has been revised in Radner and Rosenthal (1982, [72]). They show that when the (exogenous) information structure is atomless and not independent given the state, a pure strategy equilibrium cannot be constructed. In my setup, the assumption that players can only learn about the state guarantees the existence of a pure strategy Nash equilibrium. Since the design of information structures is endogenous, it is possible to convert this pure strategy equilibrium into a canonical one. Indeed, an equilibrium of the form $S_i = [0, 1]$ and $(\pi_{k_0}^i(\cdot), \pi_{k_1}^i(\cdot))$ uniformly distributed for $i = 1, 2$ with $\beta_{s_i}(H) = 1$ if $s_i \leq \frac{1}{2}$ and $\beta_{s_i}(H) = 0$ otherwise is a pure strategy equilibrium that is not canonical.

4.5.2 Example 4.5.2: Matching Pennies with Uncertain Outside Option

Consider again the game of matching pennies and suppose that the players have an additional option that depends on the realization of an unknown parameter $k \in K \equiv \{-1, 1\}$. Suppose that the prior on K is $\mu = (.5, .5)$, and that λ is 1. Payoffs are presented in table 4.2. To find an equilibrium in this example

P1/P2	<i>H</i>	<i>T</i>	<i>O</i>
<i>H</i>	1, -1	-1, 1	0, k
<i>T</i>	-1, 1	1, -1	0, k
<i>O</i>	$k, 0$	$k, 0$	k, k

Table 4.2 – Payoffs Example 4.5.2

it is enough to assume that player j is playing a strategy that puts equal probability to actions H and T in every state of nature. In this case, player i is always indifferent between H and T , obtaining a payoff of zero in both cases independently of the state of nature. Hence, since playing O reports 1 or -1 depending on the state, it is profitable to learn about the state so that actions H or T are played with higher probability when the state is -1 , and to play action O otherwise.

An equilibrium is then one where $S_1 = S_2 = (s_{HT}, s_O)$, and to play mixed strategy $\sigma = (\frac{1}{2}, \frac{1}{2})$ after observing signal s_{HT} , and to play action O after observing signal s_O . For this example, the probability of state $k = -1$ after observing signal s_{HT} is $q_{HT} = 0.73$. Analogously, $q_O = 0.27$ and $\pi(s_{HT}) = \pi(s_O) = \frac{1}{2}$. This in turn implies that $\pi(s_{HT}|k = -1) = 0.73$ and $\pi(s_{HT}|k = 1) = 0.23$. The solution is shown in Figure 4.1, where the net benefit functions and the lowest hyperplane (weakly above them) are depicted.

Note that in this example, the recommendation of signal s_{HT} is indeed a mixed strategy. Hence, let us now convert this equilibrium in a pure strategy equilibrium. Define $S'_1 = S'_2 = (s_H, s_T, s_O)$, and $\pi(s_H|k = -1) = \sigma(H)\pi(s_{HT}|k = -1)$ and $\pi(s_H|k = 1) = \sigma(H)\pi(s_{HT}|k = 1)$. Analogously with $\pi(s_T|k)$.

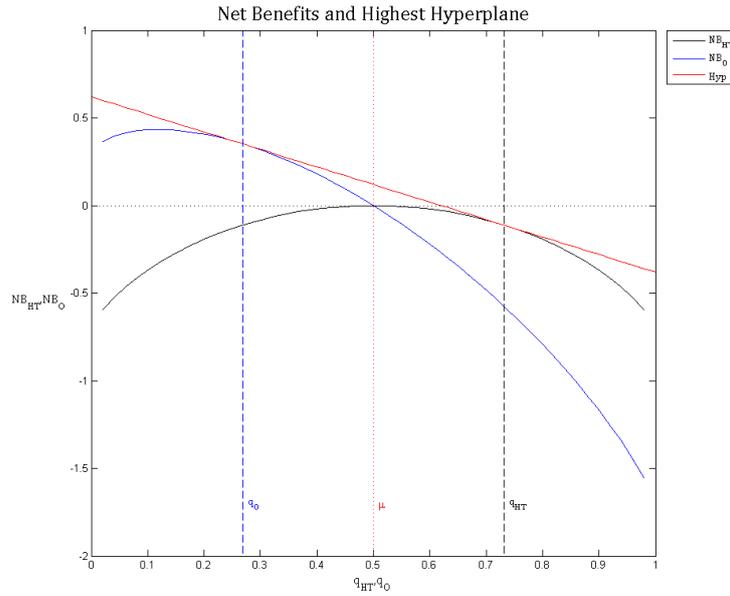


Figure 4.1 – Solution Example 4.5.2

Then, we obtain $\pi(s_H|k = -1) = 0,365$, $\pi(s_H|k = 1) = 0,115$.

Note that

$$\begin{aligned}
 q_H &= \frac{\pi(s_H|k = -1)\mu(k = -1)}{\pi(s_H|k = -1)\mu(k = -1) + \pi(s_H|k = 1)\mu(k = 1)} \\
 &= \frac{\sigma(H)\pi(s_{HT}|k = -1)\mu(k = -1)}{\sigma(H)\pi(s_{HT}|k = -1)\mu(k = -1) + \sigma(H)\pi(s_{HT}|k = 1)\mu(k = 1)} \\
 &= \frac{\pi(s_{HT}|k = -1)\mu(k = -1)}{\pi(s_{HT}|k = -1)\mu(k = -1) + \pi(s_{HT}|k = 1)\mu(k = 1)} \\
 &= q_{HT}
 \end{aligned}$$

It is easy to check that $q_T = q_{HT}$ as well. Hence, at this new equilibrium, players obtain exactly the same information about the state when observing signal s_H or s_T . This new equilibrium is such that only pure strategies are played after the observation of each signal.

4.5.3 Example 4.5.3: Matching Pennies with Costly Number of Signals

In this subsection we revise examples 4.5.1 and 4.5.2 considering a cost function that does not satisfies Blackwell informativeness. In particular, suppose that the cost depends only on the number of signals. That is, no matter how much information you gather with n signals, the cost is $c \cdot n$, with c small.

For example 4.5.1, it should be clear that no pure strategy Nash equilibrium is ever obtained. Since the state of nature is payoff irrelevant, and it is costly to use a randomization devise, players just play mixed strategies.

More interestingly, the unique equilibrium in example 4.5.2 is full information revelation with only two signals, provided that c is small. Specifically, the equilibrium may entail signals s_{HT} and s_O , and behavioral strategies $\beta_{s_{HT}}[H] = \beta_{s_{HT}}[T] = \frac{1}{2}$ and $\beta_{s_O}[O] = 1$. With a fully informative experiment, the expected payoff of each player is 0.5, and any less informative experiment (with these two signals) would yield a lower expected payoff. Another possible deviation would be to not pay for any signal (no player would ever pay for only one signal), and to not gather any information at all. For this case, the expected payoff is only zero, so it will not be a profitable deviation.

Note that this unique equilibrium is not canonical, as it is not pure in behavioral strategies. Essentially, canonical equilibrium can be obtained because adding non-informative randomization into information acquisition is for free. When the cost depends on the number of signals, such extra noise is costly. Similar results can be obtained if we suppose that there is a cap in the number of signals the players can use.

4.5.4 Example 4.5.4: Battle of Sexes

Consider the game “the battle of sexes”, where payoffs are independent of a state $k \in K \equiv \{k_0, k_1\}$. Assume that the prior is $\mu = (.5, .5)$ and that λ is $\frac{13}{10}$. Payoffs are presented in table 4.3.

P1/P2	L	R
U	4, 3	0, 0
D	0, 0	3, 4

Table 4.3 – Payoffs Example 4.5.4

If information is for free, then we obtain perfect coordination through the use of unknown state k , as the toss of a coin in Aumann (1987, [7]). If information is costly, but the cost is low enough and has bounded derivatives in the boundary of its domain, we could have such a correlated equilibrium. If the cost is high enough or has unbounded derivatives in the boundaries of the domain, as it is the case with entropy based cost functions, we will have certain degree of correlation, but not perfect correlation. Information is acquired only for coordination purposes, as it is payoff irrelevant. An equilibrium of this game is $\pi(U|s_U) = .91$, $\pi(D|s_D) = .84$ for the row player, and $\pi(R|s_R) = .84$, $\pi(L|s_L) = .91$ for the column player.

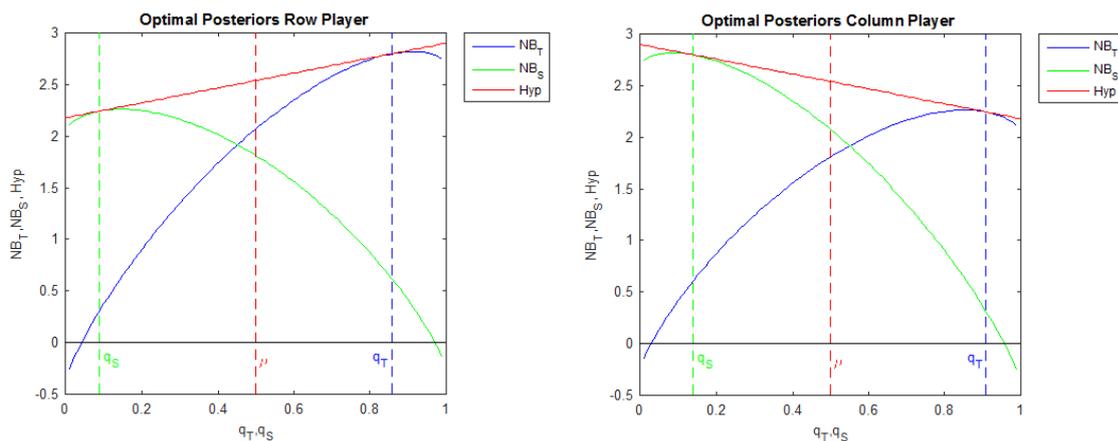


Figure 4.2 – Solution Example 4.5.4

We can see that in this example there is information extraction only for coordination purposes. As opposed to Aumann (1987, [7]), learning about the state that it is used to coordinate is costly. Indeed, if information is for free, then the expected payoff for each player would be 3, 5 for each player. In our case, the cost of information leads to imperfect coordination, which, net of cost of information, gives a payoff of 2.532 for each player. Hence, the difference of approximately one unit of payoff could be taken to be the cost of coordination.

4.5.5 Example 4.5.5: Chicken Game.

Consider a game (“chicken”) where payoffs are independent of the state k . Assume that the state space is $K \equiv \{k_0, k_1\}$ and that the prior is $\mu = (.5, .5)$. Payoffs are presented in table 4.4. The original example of Aumann (1987, [7]) showed how the joint distribution of signals given in table 4.5 can be implemented as a correlated equilibrium. That is, upon observing s_U , player one assigns half probability to player two observing s_L and s_R , so it is convenient for him to play action U .

P1/P2	L	R
U	6, 6	2, 7
D	7, 2	0, 0

Table 4.4 – Payoffs Example 4.5.5

$\varphi(s_1, s_2)$	s_L	s_R
s_U	$\frac{1}{3}$	$\frac{1}{3}$
s_D	$\frac{1}{3}$	0

Table 4.5 – Joint Distribution of Signals

The distribution showed in table 4.5 is feasible in this model. Indeed, experiments $(\pi_{k_0}^1(U), \pi_{k_1}^1(U)) = (1, \frac{1}{3})$ and $(\pi_{k_0}^2(L), \pi_{k_1}^2(L)) = (\frac{1}{3}, 1)$ generate the desired distribution. However, it remains to check whether such profile of experiments is an equilibrium of the game. Indeed, for cheap enough information acquisition, it is not possible to obtain such distribution, even when the marginal cost of information acquisition is bounded at the boundaries of its domain.²⁵ The reason is that, since player one is playing U in state k_0 , it is not optimal for player two to put any probability in action L . That is, if possible, he would set $\pi_{k_0}^2(L) = 0$. This would imply that equilibrium distribution of signals puts positive probability to the two Nash equilibria of the game.

$\varphi(s_1, s_2)$	s_L	s_R
s_U	0.18	0.39
s_D	0.39	0.04

Table 4.6 – Joint Distribution of Signals

However, it may be possible that a high enough information cost would prevent player two of setting $\pi_{k_0}^2(L) = 0$. This would also be the case if the marginal cost of information acquisition is not bounded at the boundaries of its domain. In such a case, it would be optimal for player two to assign some positive probability to action L in state k_0 . The relative size of probability mass assigned to the non stable action profiles will depend on the particular payoffs. For instance, for entropy based cost functions with $\lambda = \frac{1}{2}$, in this example we obtain $(\pi_{k_0}^1(U), \pi_{k_1}^1(U)) = (0.96, 0.19)$ and $(\pi_{k_0}^2(L), \pi_{k_1}^2(L)) = (0.19, 0.96)$. The joint distribution of signals induced by such experiment profile is showed in table 4.6.

²⁵Otherwise it would be trivial to show that such distribution is not an equilibrium because it would not satisfy individual optimality.

4.6 Conclusion

This paper studies finite games of incomplete information in which players chose *what* and *how much* information to acquire. For each player information acquisition is costly, and the cost is assumed to be proportional to the reduction of her uncertainty. I show that for every equilibrium there exists an *equivalent* pure strategy Bayes Nash equilibrium in which the space of signals is equal to the action space. I call such equilibrium a *canonical equilibrium*. Then, I prove existence of equilibrium. Finally, I show how standard results in single-agent information acquisition can be extended to strategic environments.

The existence of canonical equilibria in games with endogenous information provides theoretical foundations to simplify the analysis of situations that, in principle, could be very difficult to tackle. Canonical equilibria allows the analyst to find an equilibrium where signals recommend pure actions and those actions are followed by the players. Existence of canonical equilibrium provides a *revelation principle* for games of incomplete information with costly information acquisition.

4.7 Appendix

4.7.1 Proof of Lemma 4.1

For every Nash equilibrium $(\hat{\sigma}, \hat{\pi}, \hat{\beta})$ there exists an equivalent equilibrium $(\hat{\sigma}, \hat{\pi}, \hat{\beta})$ that is pure in experiments.

Proof. Let $(\hat{\sigma}, \hat{\pi}, \hat{\beta})$ be an equilibrium under which player i is randomizing among experiments π_1 and π_2 with probabilities $\sigma(1)$ and $\sigma(2)$ respectively, and let $\hat{\varphi} \equiv \varphi(\hat{\sigma}, \hat{\pi}, \hat{\beta})$. Let π_1^e (π_2^e) denote the *expansion* of experiment π_1 (π_2) to the realization space $S_1 \cup S_2$. That is, $\pi_1^e = (S_1 \cup S_2, (\pi_{1,k}^e)_k)$, where $\pi_{1,k}^e(s) = \pi_1^e(s)$ if $s \in S_1$ and 0 otherwise. Now, define $\bar{\pi}_i = (\bar{S}_i, (\bar{\pi}_{k,\sigma}^i)_k)$, where $\bar{S}_i \equiv \bigcup_{h \in \{1,2\}} (\{h\} \times S_h)$, and $\bar{\pi}_k(s) = \sum_{h \in \{1,2\}} \sigma(h) \pi_{k,h}^e(s)$. Note that experiments π_1^e , π_2^e and $\bar{\pi}_i$ have all the same realization space. Furthermore, let $\bar{\beta}$ be the map that goes from \bar{S}_i to $\Sigma(A)$ defined by $\bar{\beta}(s) = \beta_h(s)$ if $s \in S_h$. Finally, let $\bar{\sigma}$ put probability one in experiment $\bar{\pi}_i$. Now I show that the joint distribution of actions and states is the same under the mixture experiment σ_i over π_1 and π_2 and under experiment $\bar{\pi}_i$. Let $\hat{\varphi} \equiv \varphi(\hat{\sigma}, \hat{\pi}, \hat{\beta})$, and $\bar{\varphi} \equiv \varphi(\bar{\sigma}, \bar{\pi}, \bar{\beta})$. Then, I will show that $\bar{\varphi} = \hat{\varphi}$. Indeed,

$$\begin{aligned}
 \bar{\varphi}(a_j|k) &\equiv \int_{s \in \bar{S}_i} \bar{\beta}_s[a_j] d\bar{\pi}_k^i(s) \\
 &= \int_{s \in \bar{S}_i} \bar{\beta}_s[a_j] d \sum_{h \in \{1,2\}} \sigma(h) \pi_{k,h}^e(s) \\
 &= \sigma(1) \int_{s \in S_1} \hat{\beta}_s^1[a_j] d\hat{\pi}_k^i(s) + \sigma(2) \int_{s \in S_2} \hat{\beta}_s^2[a_j] d\hat{\pi}_k^i(s) \\
 &= \hat{\varphi}(a_j|k)
 \end{aligned}$$

Suppose without loss of generality that $c(\hat{\pi}_1^e, \mu) \geq c(\hat{\pi}_2^e, \mu)$. Since $c(\pi, \mu)$ is a quasi-convex function, we

have that

$$\begin{aligned}
V(\bar{\pi}_i, \hat{\pi}_{-i}) &\equiv \sum_{k,a} g_i(a, k) \bar{\varphi}(a_i|k) \hat{\varphi}(a_{-i}|k) \mu(k) - c(\bar{\pi}_i, \mu) \\
&= \sum_{k,a} g_i(a, k) \hat{\varphi}(a, k) - c(\sigma(1)\pi_1^3 + (1 - \sigma(1))\pi_2^e, \mu) \\
&\geq \sum_{k,a} g_i(a, k) \hat{\varphi}(a, k) - \max\{c(\pi_1^e, \mu), c(\hat{\pi}_2^e, \mu)\} \\
&= V(\hat{\pi}_1, \hat{\pi}_{-i})
\end{aligned}$$

Since $\hat{\pi}_1$ was a best response for player i , $\hat{\pi}_i$ is also optimal for player i . Since $\varphi(\hat{\pi}) = \varphi(\bar{\pi})$, the payoff of players $j \neq i$ remain unchanged, so $\hat{\pi}$ is still a Nash equilibrium. The extension for non binary randomization among experiments is trivial, provided that $c(\sum_h \sigma(h)\pi_h) \leq \max_h \{c(\pi_1, \dots, \pi_H)\}$. This later result can be found in Dragomir and Pearce (2012). \square

4.7.2 Proof of Lemma 4.2

Let $(\hat{\pi}, \hat{\beta})$ be a Nash equilibrium of $\Lambda(G, \mu)$ in which, for some player i , mixed strategy σ is induced with positive probability. Let $x(\sigma) := \{s_i \in \hat{S}_i : \text{supp}\{\beta_i(s_i)\} \cap \text{supp}\{\sigma\} \neq \emptyset\}$ be the set of signal realizations that induce mixed strategies whose support is not disjoint with the support of σ . Then, there exists a $\tilde{q} \in \Delta(K)$ such that $q_{s_i} = \tilde{q}$ for all $s_i \in x(\sigma)$.

Proof. By contradiction. Take a mixed strategy σ played with positive unconditional probability in Nash equilibrium $(\hat{\pi}, \hat{\beta})$. Assume that $q_{s_i} \neq q_{s'_i}$ for a pair $s_i, s'_i \in x(\sigma)$. We construct a new feasible strategy $(\tilde{\pi}_i, \tilde{\beta}_i)$ that generates the same expected payoff but at a lower cost. Let $\tilde{S}_i = \hat{S}_i \setminus \{s_i, s'_i\} \cup s_i^*$ be the new signal realization space, and let

$$\tilde{\pi}_k^i = \begin{cases} \hat{\pi}_k^i(s_i) + \hat{\pi}_k^i(s'_i) & \text{if } s_i = s_i^* \\ \hat{\pi}_k^i(s_i) & \text{if } s \neq s_i^* \end{cases} \quad (4.13)$$

be the new family of probability distributions of the new experiment $\tilde{\pi}_i$. It is clear that

$$q_{s_i^*}^i(k) = \frac{[\pi_k^i(s_i) + \pi_k^i(s'_i)] \mu(k)}{\sum_k [\pi_k^i(s_i) + \pi_k^i(s'_i)] \mu(k)} = \lambda_{s_i} q_{s_i}^i + \lambda_{s'_i} q_{s'_i}^i \quad (4.14)$$

Where $\lambda_s = \frac{\pi_i(s)}{\pi_i(s_i^*)}$ for $s \in \{s_i, s'_i\}$. Let $\tilde{\beta}_i(s_i) = \hat{\beta}_i(s_i)$ if $s_i \neq s_i^*$, and let $\tilde{\beta}_i(s_i) = \tilde{\sigma}$ otherwise, where $\tilde{\sigma}$ is any mixed strategy with support in $\bigcap_{\beta \in x(\sigma)} \text{supp}\{\beta\}$. First, I show that $\tilde{\sigma}$ is optimal when observing s_i^* (that is, when the posterior belief is $q_{s_i^*}^i$). Indeed,

$$\begin{aligned}
E[g_i|s_i^*, \tilde{\beta}, \tilde{\pi}] &= \sum_{k,a} g_i(a, k) \tilde{\varphi}(a, k) \\
&= \sum_{k, a_i, a_{-i}} g_i(a, k) \tilde{\varphi}(a_{-i}|k) \tilde{\beta}_{s_i^*}[a_i] q_{s_i^*}^i(k) \\
&= \sum_{s \in \{\hat{s}_i, \hat{s}'_i\}} \lambda_s \left(\sum_{k,a} g_i(a, k) \tilde{\varphi}(a_{-i}|k) \tilde{\sigma}[a_i] q_{s_i}(k) \right)
\end{aligned}$$

Since the support of $\tilde{\sigma}$ is contained in the support of $\hat{\beta}(\hat{s}_i)$ and $\hat{\beta}(\hat{s}'_i)$, the best response condition implies that

$$\begin{aligned}
\sum_{k,a} g_i(a, k) \tilde{\varphi}(a_{-i}|k) \tilde{\sigma}[a_i] q_{\hat{s}_i}(k) &= \sum_{k,a} g_i(a, k) \tilde{\varphi}(a_{-i}|k) \hat{\beta}_{\hat{s}_i}[a_i] q_{\hat{s}_i}(k) \\
\sum_{k,a} g_i(a, k) \tilde{\varphi}(a_{-i}|k) \tilde{\sigma}[a_i] q_{\hat{s}'_i}(k) &= \sum_{k,a} g_i(a, k) \tilde{\varphi}(a_{-i}|k) \hat{\beta}_{\hat{s}'_i}[a_i] q_{\hat{s}'_i}(k)
\end{aligned}$$

Finally, since $\hat{\beta}_{s_i}$ and $\hat{\beta}_{s'_i}$ are optimal upon observing \hat{s}_i and \hat{s}'_i , we conclude that for all $\beta' \in \Sigma$

$$\begin{aligned} E[g_i | s_i^*, \tilde{\beta}, \tilde{\pi}] &= \sum_{s \in \{\hat{s}_i, \hat{s}'_i\}} \lambda_s \left(\sum_{k,a} g_i(a, k) \tilde{\varphi}(a_{-i}|k) \tilde{\sigma}[a_i] q_{s_i}(k) \right) \\ &\geq \sum_{s \in \{\hat{s}_i, \hat{s}'_i\}} \lambda_s \left(\sum_{k,a} g_i(a, k) \tilde{\varphi}(a_{-i}|k) \beta'[a_i] q_{s_i}(k) \right) \\ &= E[g_i | s_i^*, \beta', \tilde{\pi}] \end{aligned}$$

Hence, we conclude that $\tilde{\sigma}$ is optimal upon observing s_i^* . Then, I show that the $\tilde{\pi}_i$ does not change the benefit and it reduces the cost. We can see again that

$$\begin{aligned} \tilde{\varphi}(a_i|k) &= \int_{s_i \in \tilde{S}_i} \tilde{\beta}_{s_i}[a_i] d\tilde{\pi}_k^i(s_i) \\ &= \int_{s_i \in \tilde{S}_i \setminus s_i^*} \tilde{\beta}_{s_i}[a_i] d\tilde{\pi}_k^i(s_i) + \tilde{\beta}_{s_i^*}[a_i] d\tilde{\pi}_k^i(s_i^*) \\ &= \int_{s_i \in \tilde{S}_i \setminus s_i^*} \tilde{\beta}_{s_i}[a_i] d\tilde{\pi}_k^i(s_i) + \sigma[a_i] \int_{s_i \in y(\sigma)} \hat{\pi}_k^i(s_i) \\ &= \int_{s_i \in \hat{S}_i \setminus y(\sigma)} \hat{\beta}_{s_i}[a_i] d\hat{\pi}_k^i(s_i) + \int_{s_i \in y(\sigma)} \hat{\beta}_{s_i}[a_i] \hat{\pi}_k^i(s_i) \\ &= \hat{\varphi}(a_i|k) \end{aligned}$$

Hence, since $\tilde{\varphi}(a_i|k) = \hat{\varphi}(a_i|k)$, expected payoffs are the same under $(\hat{\pi}, \hat{\beta})$ and $(\tilde{\pi}, \tilde{\beta})$. That is,

$$\begin{aligned} E[g_i | s_i^*, \tilde{\beta}, \tilde{\pi}] &= \sum_{k,a} g_i(a, k) \tilde{\varphi}(a, k) \\ &= \sum_{k,a} g_i(a, k) \hat{\varphi}(a, k) \\ &= E[g_i | s_i^*, \hat{\beta}, \hat{\pi}] \end{aligned}$$

Now it remains to show that $\hat{\pi} \triangleright \tilde{\pi}$, so $c(\hat{\pi}, \mu; \lambda) < c(\tilde{\pi}, \mu; \lambda)$. Let $\hat{\phi}$ and $\tilde{\phi}$ denote the probability distribution over posteriors induced by experiments $\hat{\pi}$ and $\tilde{\pi}$ (respectively) under the uniform prior. By (12), it is clear that $\tilde{\phi}(\tilde{q}_{s^*}) = \hat{\phi}(q_s) + \hat{\phi}(q_{s'})$, where \tilde{q}_{s^*} is defined by (13) when taking the uniform prior instead of μ . Hence,

$$\begin{aligned} \int_{\tilde{S}} f(q_s) \hat{\phi}(q_s) &= \int_{\hat{S} \setminus \{s, s'\}} f(q_s) \hat{\phi}(q_s) + f(q_s) \hat{\phi}(q_s) + f(q_{s'}) \hat{\phi}(q_{s'}) \\ &\geq \int_{\tilde{S} \setminus \{s, s'\}} f(q_s) \hat{\phi}(q_s) + f(q_{s^*}) (\hat{\phi}(q_s) + \hat{\phi}(q_{s'})) \\ &= \int_{\tilde{S}} f(q_s) \tilde{\phi}(q_s) \end{aligned}$$

for every convex function $f : \Delta(K) \rightarrow \mathbb{R}$, so we conclude that $\hat{\pi} \triangleright \tilde{\pi}$. Now, we show that there exists a function $g : \Delta(K) \rightarrow \mathbb{R}$ such that

$$\int_{\tilde{S}} f(q_s) \hat{\phi}(q_s) > \int_{\tilde{S}} f(q_s) \tilde{\phi}(q_s)$$

so that $\tilde{\pi} \not\triangleright \hat{\pi}$. Taking g to be any strictly convex function shows the desired result. Hence, we conclude that $c(\hat{\pi}, \mu; \lambda) < c(\tilde{\pi}, \mu; \lambda)$, which proves that $\hat{\phi}$ could not be an equilibrium strategy. \square

4.7.3 Proof of Lemma 4.3

Let $(\hat{\pi}, \hat{\beta})$ be a Nash equilibrium of $\Lambda(G, c)$ in which, for some player i , mixed strategy σ is played with positive unconditional probability, and that $|y(\sigma)| > 1$. Then, there exists another Nash Equilibrium $(\tilde{\pi}, \tilde{\beta})$ such that $|y(\tilde{\beta}(s_i))| = 1$ for all $s_i \in \tilde{S}_i$, and $\varphi(\hat{\pi}, \hat{\beta}) = \varphi(\tilde{\pi}, \tilde{\beta})$.

Proof. The proof goes by construction. Let s_i^* be the new signal associated with strategy σ , and let $\tilde{S}_i = \hat{S}_i \setminus y(\sigma) \cup s_i^*$ to be the new realization space. Let the new experiment be $\tilde{\pi}_k^i(s_i^*) = \int_{s_i \in y(\sigma)} d\hat{\pi}_k^i(s_i)$ and $\tilde{\pi}_k^i(s_i) = \hat{\pi}_k^i(s_i)$ for $s_i \neq s_i^*$. The new behavioral strategy is $\tilde{\beta}(s_i^*) = \sigma$ and $\tilde{\beta}(s_i) = \hat{\beta}(s_i)$ for $s_i \neq s_i^*$. Let $\tilde{\varphi} \equiv \varphi(\tilde{\pi}, \tilde{\beta})$, and $\tilde{\varphi} \equiv \varphi(\tilde{\pi}, \tilde{\beta})$. Then,

$$\begin{aligned} \tilde{\varphi}(a_i|k) &= \int_{s_i \in \tilde{S}_i} \tilde{\beta}_{s_i}[a_i] d\tilde{\pi}_k^i(s_i) \\ &= \int_{s_i \in \tilde{S}_i \setminus s_i^*} \tilde{\beta}_{s_i}[a_i] d\tilde{\pi}_k^i(s_i) + \tilde{\beta}_{s_i^*}[a_i] \tilde{\pi}_k^i(s_i^*) \\ &= \int_{s_i \in \tilde{S}_i \setminus \tilde{s}_i} \tilde{\beta}_{s_i}[a_i] d\tilde{\pi}_k^i(s_i) + \sigma[a_i] \int_{s_i \in y(\sigma)} d\hat{\pi}_k^i(s_i) \\ &= \int_{s_i \in \tilde{S}_i \setminus y(\sigma)} \hat{\beta}_{s_i}[a_i] d\hat{\pi}_k^i(s_i) + \int_{s_i \in y(\sigma)} \hat{\beta}_{s_i}[a_i] d\hat{\pi}_k^i(s_i) \\ &= \hat{\varphi}(a_i|k) \end{aligned}$$

Hence, since $\tilde{\varphi}(a_i|k) = \hat{\varphi}(a_i|k)$, expected payoffs are the same under $(\hat{\pi}, \hat{\beta})$ and $(\tilde{\pi}, \tilde{\beta})$. Note also that $y(\sigma) \subseteq x(\sigma)$, so by Lemma 4.2, the posterior is constant for all elements in $y(\sigma)$. This, together with the fact that $\tilde{\beta}(s_i^*) = \sigma = \hat{\beta}(s_i)$ for all $s_i \in y(\sigma)$, implies that replacing $\hat{\beta}(s_i)$ by $\tilde{\beta}(s_i^*)$ is still optimal for player i . Finally, since $\hat{\pi}$ and $\tilde{\pi}$ provide exactly the same information, the cost of both experiments is the same. Thus, $(\tilde{\pi}, \tilde{\beta})$ is also a Nash equilibrium. This procedure can be applied simultaneously to all strategies σ' such that $y(\sigma') > 1$, which gives desired result. \square

4.7.4 Proof of Lemma 4.4

Let $(\hat{\pi}, \hat{\beta})$ be a Nash equilibrium of $\Lambda(G, \mu)$. Suppose that for some mixed strategy $\tilde{\beta}$ played with positive probability by player i , $|\text{supp}\{\tilde{\beta}\}| > 1$. Then, there exists another Nash Equilibrium $(\tilde{\pi}, \tilde{\beta})$ such that $|\text{supp}\{\tilde{\beta}_i(s_i)\}| = 1$ for all $s_i \in \tilde{S}_i$, $i \in I$ and $\varphi(\hat{\pi}, \hat{\beta}) = \varphi(\tilde{\pi}, \tilde{\beta})$.

Proof. Let $(\hat{\pi}, \hat{\beta})$ be a Nash equilibrium of $\Lambda(G, \mu)$, and suppose that there exists a $\beta_i(s_i^*)$ such that $|\text{supp}\{\beta_i(s_i^*)\}| = M > 1$. Making a slight abuse of notation, let M denote the support of $\beta_i(s_i^*)$, and denote $m_k \in M$ an arbitrary element of M . We will construct a new equilibrium strategy $(\tilde{\pi}_i, \tilde{\beta}_i)$ such that $|\text{supp}\{\tilde{\beta}_i(s_i)\}| = 1$ for every $s_i \in \tilde{S}_i$ and $\varphi(\hat{\pi}, \hat{\beta}) = \varphi(\tilde{\pi}, \tilde{\beta})$.

Let \tilde{S}_i be a new signal realization space that replaces $s_i^* \in \hat{S}_i$ with an alphabet $\Delta \equiv \{m_k\}_{k=1}^M$. That is, $\tilde{S}_i \equiv \hat{S}_i \setminus s_i^* \cup \Delta$. Let $\tilde{\pi}_i(\Delta|k) = \hat{\pi}_k^i(s_i^*)$, and let $\tilde{\pi}_i(m_k|\Delta) = \hat{\beta}_i(s_i^*)[m_k]$. Now, let

$$\tilde{\pi}_i(s_i|k) = \begin{cases} \tilde{\pi}_i(s_i|\Delta)\tilde{\pi}_i(\Delta|k) & \text{if } s_i \in \Delta \\ \hat{\pi}_k^i(s_i) & \text{if } s_i \notin \Delta \end{cases}$$

It is easy to see that for all s'_i and s''_i in Δ and for all $k \in K$,

$$\tilde{q}_{s'_i}(k) = \frac{\tilde{\pi}(s'_i|k)\mu(k)}{\sum_{k'} \tilde{\pi}(s'_i|k')\mu(k')} = \frac{\tilde{\pi}_i(s'_i|\Delta)\tilde{\pi}_k^i(\Delta|k)\mu(k)}{\sum_{k'} \tilde{\pi}_i(s'_i|\Delta)\tilde{\pi}_k^i(\Delta|k')\mu(k')} = \tilde{q}_{s''_i}(k)$$

Consequently, since the distribution of posteriors remains unchanged, the cost of $\tilde{\pi}$ is equal to the cost of $\hat{\pi}$. Next, let $\tilde{\beta}(s_i) = \hat{\beta}(s_i)$ if $s_i \notin \Delta$, and $\tilde{\beta}(s_i) = \mathbf{1}_{s_i}$ otherwise. We want to show that $\varphi(\hat{\pi}, \hat{\beta}) = \varphi(\tilde{\pi}, \tilde{\beta})$.

Indeed,

$$\begin{aligned}
\tilde{\varphi}(a_i|k) &= \int_{s_i \in \tilde{S}_i} \tilde{\beta}_{s_i}[a_i] d\tilde{\pi}_k^i(s_i) \\
&= \int_{s_i \in \tilde{S}_i \setminus \Delta} \tilde{\beta}_{s_i}[a_i] d\tilde{\pi}_k^i(s_i) + \int_{s_i \in \Delta} \tilde{\beta}_{s_i}[a_i] d\tilde{\pi}_k^i(s_i) \\
&= \int_{s_i \in \tilde{S}_i \setminus \Delta} \hat{\beta}_{s_i}[a_i] d\hat{\pi}_k^i(s_i) + \int_{s_i \in \Delta} \tilde{\beta}_{s_i}[a_i] \tilde{\pi}_i(s_i|\Delta) \tilde{\pi}_i(\Delta|k) \\
&= \int_{s_i \in \tilde{S}_i \setminus \Delta} \hat{\beta}_{s_i}[a_i] d\hat{\pi}_k^i(s_i) + \sum_{m_k \in \Delta} \tilde{\pi}_i(m_k|\Delta) \tilde{\pi}_i(\Delta|k) \\
&= \int_{s_i \in \tilde{S}_i \setminus \Delta} \hat{\beta}_{s_i}[a_i] d\hat{\pi}_k^i(s_i) + \tilde{\pi}_i(\Delta|k) \\
&= \int_{s_i \in \tilde{S}_i \setminus \Delta} \hat{\beta}_{s_i}[a_i] d\hat{\pi}_k^i(s_i) + \hat{\pi}_k^i(s_i^*) \\
&= \int_{s_i \in \hat{S}_i} \hat{\beta}_{s_i}[a_i] d\hat{\pi}_k^i(s_i) \\
&= \hat{\varphi}(a_i|k)
\end{aligned}$$

Since $\tilde{\varphi}(a_i|k) = \hat{\varphi}(a_i|k)$, and since the cost of $\hat{\pi}$ and $\tilde{\pi}$ are the same, $(\tilde{\pi}, \tilde{\beta})$ is an equilibrium. Note that the equilibrium $(\tilde{\pi}, \tilde{\beta})$ does not have mixed strategy $\beta_i(s_i^*)$. By proceeding inductively on strategies and players, it is possible to obtain an equilibrium such that, $|\text{supp}\{\tilde{\beta}_i(s_i)\}| = 1$ for all $s_i \in \tilde{S}_i$, $i \in I$, and $\varphi(\hat{\pi}, \hat{\beta}) = \varphi(\tilde{\pi}, \tilde{\beta})$. \square

Chapter 5

Ideology, Media Outlets and Political Extremism

Abstract

This paper studies the role of *ideology* in political extremism and social welfare. I model ideology as a misperception of reality that agents recognize in others but fail to acknowledge in themselves. Specifically, two ideological groups receive information that is biased and correlated, and they suffer from *bias neglect* and *correlational neglect*. Agents correctly assess the bias and correlation in the information sources of the other group, and they acknowledge that the other group is failing to recognize them. Agents also receive information from an unbiased publicly available source. Individuals have to choose their political position on an issue, and their payoffs depend not only on the realization of an unknown state but also on the level of agreement with the rest of the population. In this context, I show that the precision of the public and private information have an ambiguous effect in political extremism. I further show that welfare is increasing in the precision of the public signal, and decreasing in bias and correlational neglect.

Keywords: Ideology, political extremism, bias neglect, correlational neglect.

JEL Classification: D71, D72, D83.

5.1 Introduction

It is hard to explain the rise of political extremism without a notion of *ideology*.¹ Despite its importance, nor economists nor political philosophers have come to agree with a unified definition of the concept. In its many forms, the term is used to label preferences, opinions, beliefs and some types of psychological biases. Political philosopher Terry Eagleton says that ideology is like halitosis, because you can easily recognize it in others but can hardly recognize it in yourself.²

This paper proposes a framework in which ideology is considered as a misperception of reality that agents acknowledge in others but fail to recognize in themselves. Specifically, two ideological groups obtain their private information from biased information sources. I assume that agents from the same

¹ See Ortoleva and Snowberg (2015, [69]), Martin and Yurukoglu (2017, [60]), Gentzkow and Shapiro (2006, [34]), Prior (2013, [71]).

² Terry Eagleton (1990, [81]).

group consume news from similar media outlets, so the *signals* they receive are correlated. Agents suffer from *bias neglect*, under which they fail to recognize the bias in their information sources. Furthermore, they suffer from *correlational neglect*, where they fail to recognize the correlation among signals within ideological group.³ Nonetheless, agents correctly assess the bias and correlation in the information sources of the other group, and they acknowledge that the other group is failing to recognize them.⁴ Agents also receive information from an unbiased publicly available source.

Individuals have to choose their political position on an issue (e.g. tax rates, government spending, etc.), and their payoffs depends on the realization of an unknown state (e.g. interest rates, unemployment, etc.). Their payoffs also depend on the the level of agreement with the rest of the population, so there is strategic complementarity among individual political positions. In this context, I study the role of bias and correlational neglect on *political extremism* and *social welfare*. More precisely, I define political extremism as the variance of the average action in the population, and social welfare as the expected sum of individual payoffs conditional on the state.⁵ Although the notion of welfare is fairly natural, the definition of political extremism is somewhat arbitrary. Furthermore, since political extremism is defined as a *variance* and the bias in information sources is a constant, bias neglect has no role political extremism. Hence, it remains to explore other definitions.

Information in the model is assumed to be normal and the payoffs to be quadratic, so equilibrium strategies are linear in the signals that the players observe. Despite linearity of equilibrium, tractability is at times elusive. Hence, the main results of this paper are provided for the symmetric scenario, that is, when the variance of private information and correlational neglect are the same for both groups and the bias neglect in one group is the negative of the other.⁶ For this case, I show that the precision of public information has an ambiguous effect on political extremism. A key factor in that relation is the degree of correlational neglect. If correlational neglect is low, and public information is noisy, increasing the variance of public information may reduce political extremism. The reason is that agents will put more weight in their private signals, motivated not only by its relative increase in precision but also by its increase in coordination power. These two effects are enough for reducing political extremism, despite the increase in the variance of the public signal. If correlational neglect is large, extremism is always increasing in the variance of public information.

The noise in the private signals is composed by a zero-mean idiosyncratic shock (independent across all agents), and a shock that is common to all members in the same ideological group.⁷ As with public information, the variance of the idiosyncratic and group noise have ambiguous effect in political extremism. If correlational neglect is low, political extremism is always increasing in the idiosyncratic variance. If correlational neglect is high, political extremism can be increasing or decreasing in in the same variable. In particular, if the variance of the group noise is large (but not acknowledged), agents put too much weight in their private information. Hence, an increase in the idiosyncratic variance leads players to put less weight on their private signals, reducing political extremism.

³Note that this notion of correlational neglect is different from the one studied in the literature. In Ortoleva and Snowberg (2015, [69]) and Levy and Razin (2015) an individual suffers from correlational neglect when he fails to recognize the correlation among *his* signals. Hence, correlational neglect generates overconfidence. See discussion in section 5.2.

⁴They also believe that agents in the other ideological group mistakenly believe that they fail to recognize their own bias.

⁵Ortoleva and Snowberg (2015, [69]) call the variance of the average action in the population *ideological extremeness*. They also define ideology as the average political position chosen by the agents.

⁶Hence, robustness of the results for the general case is yet to be proven.

⁷A central problem of my specification is that correlational neglect is also an under estimation of the variance of the common shock of the members in the same ideological group. Hence, correlational neglect accounts for an over estimation of the precision of your signals (and those of the member of your group) *and* an under estimation of the correlation among signals. This issue could be solved by normalizing the variance of the private signal. See discussion in Section 5.2.

The variance of the common shock has two effects on the private signal. First, it adds noise, making the information of the private signal less precise. Second, it acts as a coordinating tool, as it allows players in the same group to better coordinate their actions. For this reason, if correlational neglect is low, the coordinating power of signals is fully acknowledged and incorporated into optimal actions. Hence, increasing the variance in the group noise would always increase political extremism.

In turn, if the variance of the common shock is large but not acknowledged, agents put too much weight in their private information, and this mistake is acknowledged by agents in the other group. Then, a reduction in the weight that agents in one group put in their private signals leads individuals in the other group to do the same. Hence, an increase in the variance of the common shock leads agents to put significantly less weight in their private signals, and the total effect leads to a reduction in political extremism. Regarding welfare, I show that expected social welfare is increasing in the precision of the public signal. I further show that expected welfare is decreasing in the size of the bias of the information sources, and that is decreasing in bias and correlational neglect.

Related Literature. This paper makes a humble theoretical contribution to the literature on behavioral political economy. On the one hand, this paper contributes to the study of correlational neglect. Important papers in this recent literature are Ortoleva and Snowberg (2015, [69]) and Levy and Razin (2015, [55]), who study the effect of correlational neglect in political behavior. In the former, there is no strategic interaction, and agents fail to recognize that the signals *they* observe are correlated with each other. This implies that agents wrongly assess the informativeness of their signals. They show that correlational neglect generates overconfidence, and they document its impact over political preferences, voter turnout and partisan identification. Levy and Razin (2015, [55]) look at an strategic environment in which the two signals *each* player observes may be correlated. They show that *behavioral voters*, those who misperceive the correlation, may end up making better choices than those who are *rational*, because they may *correctly* change his behavior after observing two apparently independent signals. In this paper we also look at an strategic model, but I introduce correlational neglect as a failure to recognize the correlation among signals within an ideological group. I also introduce bias neglect, as I take information sources to be biased.

On the other hand, this paper contributes to the study of biased media outlets and its impact on political behavior. In the empirical arena, the bias of media outlets and its impact in political behaviour has been widely documented.⁸ In the theoretical arena, Gentzkow and Shapiro (2006, [34]) study a model in which media outlets bias their information in order to build reputation for their audience. Gentzkow, Wong and Zhang (2018, [35]) introduce ideological bias as a form of correlational neglect, under which an agent fails to recognize that his *experiences* are correlated with the information he receives. This may lead people to persistently disagree with each other, and to trust in more biased information sources.

The rest of the paper is organized as follows. Section 5.2 presents the model. Section 5.3 solves the model and performs comparative statics on equilibrium strategies. Section 5.4 studies political extremism and section 5.5 studies social welfare. Section 5.6 concludes and discusses possible directions of further research.

⁸See Glaeser and Wald (2006, [37]), McCarty, Poole and Rosenthal (2016, [65]), Prior (2013, [71]), Martin and Yurukoglu (2017, [60]). In general, it has been documented the rise of bias in media outlets, and its positive effect in political polarization.

5.2 The Model

5.2.1 Basics

I consider a continuum of agents with mass normalized to one. The universe of agents is divided in two groups, U and D , with masses m_U and m_D respectively. There is an unknown state of nature $\theta \in \Theta \equiv \mathbb{R}$. Each agent i in group $h \in \{U, D\}$ chooses an action $a_{h_i} \in A_i \equiv \mathbb{R}$, which is interpreted as its ideal policy. I assume that the loss of agent i in group h is given by

$$g_{h_i}(a, \theta) = -(1 - r)(a_{h_i} - \theta)^2 - r(a_{h_i} - \bar{a})^2 \quad (5.1)$$

where $\bar{a} \equiv \int_{j \in m_U} a_{U_j} dj + \int_{j \in m_D} a_{D_j} dj$, and $r \in (0, 1)$. The loss function has two components. The first component is a standard quadratic loss in the distance between the underlying state and his action. The second component is a quadratic loss in the distance between his action and the average action in the population. This term represents a preference for agreement in the society, and it is the source of strategic complementarity in the actions of the players. The parameter r gives the relative weight between these two components.⁹

5.2.2 Information

The state of nature θ is initially drawn from a uniform improper prior over the real line, and players observe a public signal

$$y = \theta + \eta \quad (5.2)$$

where η is distributed according to $\mathbb{P}_\eta \sim \mathcal{N}(0, \tau^2)$, independent of θ . After observing the public signal, players observe a private signal that is correlated with the signals of individuals in the same group. In particular, I assume that player h_i observes

$$x_{h_i} = \theta + \varepsilon_{h_i} + \phi_h \quad (5.3)$$

where ε_{h_i} is an idiosyncratic shock distributed according to $\mathbb{P}_{\varepsilon_{h_i}} \sim \mathcal{N}(0, \sigma_h^2)$, and ϕ_h is common shock to all members in group h , distributed according to $\mathbb{P}_{\phi_h} \sim \mathcal{N}(\omega_h, \nu_h^2)$. Both random variables are assumed to be independent of all other random variables in the model.

We consider the expectation of ϕ_h to be the bias of the information source of group h . We assume that $\omega_U > 0 > \omega_D$, so that group U receives signals with higher values than those of group D . Note that if ω_h is known to the members of group h , the bias of the information source is irrelevant, in the sense that information is still correctly backed out by the agents. They have to correct for the bias of the source and obtain unbiased estimates of θ .

5.2.3 Ideological Neglect

I assume that agents suffer from *ideological neglect* if they fail to correctly perceive the information they receive. On the one hand, agents may fail to identify the bias in their information source, which I call *bias neglect*. On the other hand, agents may fail to recognise that the signals they receive are correlated with the signals of individuals in the same ideological group. I call this *correlational neglect*.¹⁰

⁹An interesting extension would be to explore the case in which agents value differently the agreement with people in different groups. An extreme case would be that agents in one group only care about agreeing with those in their group.

¹⁰See Ortoleva and Snowberg (2015, [69]).

Formally, I assume that agents in group h misperceive the distribution of ϕ_h , but they are right about the distribution of $\phi_{h'}$. In particular, agent h_i believes that the distribution of ϕ_h is $\tilde{\mathbb{P}}_{\phi_h} \sim \mathcal{N}(a_h \omega_h, b_h \nu_h^2)$, with $a_h, b_h \in [0, 1]$. Consequently, the perceived joint distribution of signals among agents in the same group, conditional on θ , is given by

$$\tilde{\mathbb{P}}_{x_i^h, x_j^h | \theta} \sim \mathcal{N} \left(\begin{pmatrix} a_h \omega_h \\ a_h \omega_h \end{pmatrix}, \begin{pmatrix} \sigma_h^2 + b_h \nu_h^2 & b_h \nu_h^2 \\ b_h \nu_h^2 & \sigma_h^2 + b_h \nu_h^2 \end{pmatrix} \right). \quad (5.4)$$

Hence, there exists *correlational neglect* whenever $b_h < 1$, and *bias neglect* if $a_h < 1$. The true distribution $\mathbb{P}_{x_i^h, x_j^h | \theta}$ is given by expression (5.4) when $a_h = b_h = 1$.

Note that when $b_h < 1$, the variance perceived by h_i is $\sigma_h^2 + b_h \nu_h^2$. Hence, he perceives that his private information is more precise than what it actually is. Consequently, $b_h < 1$ accounts for an over estimation of the precision of your signals (and those of the member of your group) *and* an under estimation of the correlation among signals. A solution to this problem would be to normalize the variance of the private signal to one, as it is done in Gentzkow, Wong and Zhang (2018, [35]). Then, correlational neglect would correspond to an over estimation of *the part* of this variance that comes from the idiosyncratic shock, with no misperception of the total variance.

5.2.4 Strategies and Equilibrium

After observing the public and private signals, each agent updates his belief and chooses an action a_{h_i} to maximize the expectation of equation (5.1). A behavioral strategy for an agent h_i is a mapping from information to actions, that specifies what action to choose given the signals observed. An equilibrium is a profile of such strategies such that no player has incentive to unilaterally deviate.

5.3 Solving the Model

To solve the model I use standard tools to solve games with quadratic payoffs and normal information.¹¹ In this case, equilibrium strategies are linear functions of the signals, so we propose a linear solution to identify coefficients afterwards.

5.3.1 Preliminaries

Before finding the linear equilibrium, I state as lemma the estimates of player h_i of the state of nature θ , signal x_{h_j} and signal $x_{h'_j}$, given the observation of signal x_{h_i} . Since predictions are subject to ideological neglect, I denote the expectations as $\tilde{E}[\cdot | x_{h_i}]$. The result follows from applying standard results of normal distributions and the law of iterated expectations.

Lemma 5.1. *The expectation of agent h_i of state of nature θ , signal x_{h_j} , and signal $x_{h'_j}$, given the observation of signal x_{h_i} , are respectively given by*

- (i) $\tilde{E}[\theta | x_{h_i}] = \beta_h \hat{x}_{h_i} + (1 - \beta_h)y$
- (ii) $\tilde{E}[x_{h_j} | x_{h_i}] = a_h \omega_h + (\beta_h + \gamma_h) \hat{x}_{h_i} + (1 - \beta_h - \gamma_h)y$
- (iii) $\tilde{E}[x_{h'_j} | x_{h_i}] = \omega_{h'} + \beta_h \hat{x}_{h_i} + (1 - \beta_h)y$

where $\hat{x}_{h_i} \equiv x_{h_i} - a_h \omega_h$, $\beta_h \equiv \frac{\tau^2}{\sigma_h^2 + b_h \nu_h^2 + \tau^2}$ and $\gamma_h \equiv \frac{b_h \nu_h^2}{\sigma_h^2 + b_h \nu_h^2 + \tau^2}$.

¹¹ See Morris and Shin (2002, [67]) and Angeletos and Pavan (2007, [3]).

Note that \hat{x}_{h_i} is the signal observed by player h_i when it is corrected by the recognised bias $a_h \omega_h$. Since $\sigma_h^2 + b_h \nu_h^2$ is the perceived variance of the private signal, β_h represents the (perceived) relative variance of the public signal with respect to the (perceived) aggregate variance of information. Additionally, γ_h is the acknowledged variance of the common noise relative to the aggregate variance of information. Finally, $\tilde{E}[\theta|x_{h_i}]$ is a weighted average of private and public information, where the weight is given by the relative precision of the signals.

The estimate of x_{h_j} (by agent h_i) corrects for the recognised bias and puts more weight on the private signal whenever the variance of the group noise (or the variance of the public signal) increases. Since agent h_j is right about the signals $x_{h'_i}$, the estimate of $x_{h'_j}$ is the expectation of θ given x_{h_i} plus a correction for the true bias.

5.3.2 Linear Equilibrium

To find the linear equilibrium, we propose an optimal action that is linear in the signals observed by the agents. That is, a linear solution of the form

$$a_{h_i}^* = \lambda_h^0 + \lambda_h^1 \hat{x}_{h_i} + (1 - \lambda_h^1)y \quad (5.5)$$

We want to find coefficients $(\lambda_h^0, \lambda_h^1)$ so that (5.5) is indeed an equilibrium of the game. After observing signal x_{h_i} , the action that minimises the expected loss in equation (5.1) is given by

$$a_{h_i}^* = (1 - r)\tilde{E}[\theta|x_{h_i}] + r\tilde{E}[\bar{a}|x_{h_i}] \quad (5.6)$$

Then, we obtain an expression for $\tilde{E}[\bar{a}|x_{h_i}]$ by replacing (5.5) into \bar{a} , and replace $\tilde{E}[\bar{a}|x_{h_i}]$ and $\tilde{E}[\theta|x_{h_i}]$ back into (5.6) to obtain

$$a_{h_i}^* = \kappa_h^0 + \kappa_h^1 \hat{x}_{h_i} + (1 - \kappa_h^1)y \quad (5.7)$$

where κ_h^0 and κ_h^1 are functions of $(\lambda_h^0, \lambda_h^1, \lambda_{h'}^0, \lambda_{h'}^1)$.¹² Finally, by equating coefficients in (5.5) and (5.7), we obtain the following proposition.

Proposition 5.1. *The unique linear equilibrium is given by optimal actions $a_{h_i}^* = \lambda_h^0 + \lambda_h^1 \hat{x}_{h_i} + (1 - \lambda_h^1)y$, where*

$$\lambda_h^0 = \frac{rm_{h'}}{1-r} [(1 - a_h)\omega_h \lambda_h^1 rm_h + (1 - a_{h'})\omega_{h'} \lambda_{h'}^1 (1 - rm_{h'})] \quad (5.8)$$

$$\lambda_h^1 = \frac{(1-r)\beta_h}{1 - rm_h(\beta_h + \gamma_h) - \delta_h rm_{h'}\beta_{h'}} \quad (5.9)$$

with $\delta_h \equiv \frac{1 - rm_h \gamma_h}{1 - rm_{h'} \gamma_{h'}}$.

It is easy to check that $0 < \lambda_h^1 < 1$, so λ_h^1 is the weight that an agent in group h puts on his private information, and $(1 - \lambda_h^1)$ is the weight placed on public information.¹³ Note also that the constant in optimal actions corrects for the unrecognized biases that both groups face. In particular, if $a_h = a_{h'} = 1$, then nobody suffers from bias neglect, and $\lambda_h^0 = 0$ for both groups.

¹²See the proof of the next proposition for definitions of κ 's.

¹³Indeed, to check that $rm_h(\beta_h + \gamma_h) - \delta_h rm_{h'}\beta_{h'} < 1$, it is enough to note that $\beta_h + \gamma_h \leq 1$ and that $\delta_h \beta_{h'} \leq \frac{1}{1 - rm_{h'}\gamma_{h'}}$. To conclude that $(1 - r)\beta_h \leq 1 - rm_h(\beta_h + \gamma_h) - \delta_h rm_{h'}\beta_{h'}$, it is enough to consider that $\beta_h + \gamma_h \leq 1$ and to rearrange terms.

5.3.3 Comparative Statics

The weight that an agent puts on his private signal depends not only on the relative variance of the public signal with respect to the variance of his private information, but also on that of agents in the other group. On the one hand, increasing the relative variance of the public signal decreases the relative variance of his private signal, making him to put a higher weight in his private signal. However, an increase in the variance of the public signal also increases the weight that agents in the other group put on their private signals, which adds an ambiguous effect on the equilibrium weights. Overall, as we will see, the effect is positive.

Furthermore, an increase in the idiosyncratic noise encourages the agents to put less weight in their private signals. At the same time, a strategic effect takes place again in which the reduction of the weight on the private signal by people from the other group alters as well the optimal weight. As we will see, the overall effect is negative. This two results are formalized in the following propositions.

Proposition 5.2. *The weight that agent h_i puts on his private signal is increasing in the variance of the public signal, and decreasing in the variance of the idiosyncratic shock of the private signal.*

It is important to assess what happens with the equilibrium weight when, for a fixed level of variance in the private signal, we shift the source of that noise from the variance of the idiosyncratic shock to the variance of the common shock. When increasing the variance of the group noise (while keeping constant the total variance of the private signal), the weight that each agent puts on his private information increases. This is stated in the following proposition.

Proposition 5.3. *For a fixed variance of private signals x_{h_i} , the weight agent h_i puts on his private signal is increasing in the variance of the common shock of his group.*

Proposition (5.3) highlights the strategic effect that group noise generates in the model. The higher the correlation between the agents in a group, the higher the weight they put in their private signals. Agents can better coordinate and trust more their private information when they know that more people is going to learn similar things.

I conclude this section by showing that reducing correlational neglect of one group reduces the weight that agents (from that group) put on their private signals. Intuitively, reducing correlational neglect increases the awareness of potential coordination among agents in a group, but it also increases the perceived variance of the private information. At the end, it is the second effect that prevails. Furthermore, a reduction in the correlational neglect of the other group leads agents to put a lower weight on their private signals.

Proposition 5.4. *The weight agent h_i puts on his private signal is increasing in the correlational neglect of agents in his group, and increasing the correlational neglect of agents in the other group.*

As for the changes in the constant term of the optimal action, the effect of changing τ^2 , σ_h , $\sigma_{h'}$, γ_h and $\gamma_{h'}$ are all ambiguous. The direction of the effect depends not only on the changes in λ_h^1 and $\lambda_{h'}^1$, but also on the size of the bias and the degree of bias neglect in each group. However, what can be seen easily is how the constant changes when the bias neglect of each group changes.

In particular, it is easy to see that $\frac{\partial \lambda_U^0}{\partial a_U} < 0$, so that increasing the awareness of people in group U tend to lower the optimal action chosen by people in that group. That is, people in group U will recognize that they are getting upward bias information, and correct for such bias. The same phenomenon happens with the downward biased group, because $\frac{\partial \lambda_D^0}{\partial a_D} > 0$. Additionally, it can be checked that $\frac{\partial \lambda_U^0}{\partial a_D} > 0$, so that reducing the bias of downward-biased people gives room for increasing the action on the upward-bias group. The argument is analogous for $\frac{\partial \lambda_D^0}{\partial a_U} < 0$.

5.4 Political Extremism

To measure political extremism, I look at the variance of the average action at equilibrium.¹⁴ In this and the following sections, I assume that $\sigma_U^2 = \sigma_D^2 = \sigma^2$, $\nu_U^2 = \nu_D^2 = \nu^2$, $b_U = b_D = b$ and $m_U = m_D = \frac{1}{2}$. Furthermore, I assume that $\omega_U = -\omega_D = \omega$, and that $a_U = a_D = a$. This implies that the weight that agents put in their private signals is the same for both ideological groups, so $\lambda_1 \equiv \lambda_U^1 = \lambda_D^1$ is given by

$$\lambda_1 = \frac{(1-r)\beta}{1 - \frac{r}{2}(2\beta + \gamma)}$$

and the constant in the optimal action for each group is given by

$$\lambda_U^0 = -\lambda_D^0 = -\frac{(1-a)\lambda_1\omega}{2}$$

The next lemma provides an expression for political extremism.¹⁵

Lemma 5.2. *Political extremism, defined as the variance of the average action at equilibrium, is given by*

$$P \equiv \frac{1}{2}\lambda_1^2\nu^2 + (1 - \lambda_1)^2\tau^2 \quad (5.10)$$

From the expression above, we can see that the variance of the idiosyncratic noise affects political extremism only through the weight that agents in each group put to their private information. The reason is that the average idiosyncratic shock along the population is zero with probability one, so it has no effect in the variability of the average action. Expression (5.10) is suitable to understand how changes in the parameters of the model affect political extremism. We summarize these results in the following propositions.

Proposition 5.5. *If $\frac{\sigma^2}{\nu^2} \leq (1-r) - b(1 - \frac{r}{2})$, political extremism is increasing in the variance of the public signal. If $\frac{\sigma^2}{\nu^2} > (1-r) - b(1 - \frac{r}{2})$, there exists a $\hat{\tau}^2$ such that for $\tau^2 < \hat{\tau}^2$ ($\tau^2 > \hat{\tau}^2$) political extremism is increasing (decreasing) in the variance of the public signal.*

The effect of the variance of the public signal on political extremism is in general ambiguous, and it crucially depends on correlational neglect. If correlational neglect is high (i.e. $b \rightarrow 0$) and the variance of the common shock is large, more precise public information will always reduce political extremism. In this case, agents put too much weight in their noisy private signals, and this mistake in group h' is acknowledged by group h . Hence, a reduction in the weight that agents in one group put in their private signals leads individuals in the other group to do the same. Hence, a decrease in the variance of the public signal leads agents to put significantly less weight in their private signals, and the total effect leads to a reduction in political extremism.

If correlational neglect is low (i.e. $b \rightarrow 1$), political extremism may be increasing or decreasing in the variance of the public noise. If τ^2 is small, players pay a lot of attention to the public information, so an increase in τ^2 could only lead to an increase in the variance of the average action. In turn, if τ^2 is high, a small increase in the τ^2 will motivate players to strongly reduce the weight in their private signals, motivated by both the relative increase in precision of the private signal and the relative increase in coordination power (because correlational neglect is low). This leads to a reduction in political extremism, because the total variance of the private signal is less than the variance of the public signal.¹⁶

¹⁴See footnote 5

¹⁵In the appendix we derive an expression for political extremism for the general case, and then impose the conditions above.

¹⁶It is easy to check that $\hat{\tau} > \sigma^2 + b\nu^2$.

Proposition 5.6. *If $b \geq \frac{(1-r)}{4(1-\frac{r}{2})}$, political extremism is increasing in the variance of the common shock. If $0 < b < \frac{(1-r)}{4(1-\frac{r}{2})}$, there exists a $\hat{\nu}^2$ such that for $\nu^2 \geq \hat{\nu}^2$ ($\nu^2 \leq \hat{\nu}^2$) political extremism is decreasing (increasing) in the variance of the common shock.*

When correlational neglect is low enough, political extremism is always increasing in the variance of the common shock. That is, if agents recognize the correlation among their signals, they correctly incorporate it their optimal strategies. Then, an increase in the noise in the private signal could only lead to an increase in political extremism. When correlational neglect and the variance of the common shock are high, everybody is paying a lot of (undue) attention to their private signal, which are highly noisy. If the variance of the common shock increases, players will reduce the weight they put in their noisy private signals, reducing political extremism.

Proposition 5.7. *If $\frac{\sigma^2}{\nu^2} \geq \frac{(1-r)-2b(1-\frac{r}{2})}{2}$, political extremism is increasing in the variance of the idiosyncratic shock and increasing in correlational neglect. If $\frac{\sigma^2}{\nu^2} \leq \frac{(1-r)-2b(1-\frac{r}{2})}{2}$ political extremism is decreasing in the variance of the idiosyncratic shock and decreasing in correlational neglect.*

It is clear that if correlational neglect is low ($b \rightarrow 1$), then political extremism is always increasing in the variance of the idiosyncratic shock of the private signal. In this case, the coordinating power of signals is acknowledged and incorporated into optimal actions, so an increase in the idiosyncratic variance can only increase political extremism. If correlational neglect and the variance of the common shock are high, an increase in the idiosyncratic variance leads to an important reduction in the attention to the (noisy) private information, reducing political extremism.

5.5 Welfare Analysis

In this section I study what happens with social welfare when different parameters of the model change. As in the previous section, I assume that all parameters are the same for both groups. I define welfare as the sum of all agents payoffs at equilibrium. That is, welfare is given by

$$W(a^*, \theta) = \sum_h \int_{m_h} g_{h_i}(a^*, \theta) di \quad (5.11)$$

We are interested in the expectation of (5.11) given θ . Under the above assumptions, I provide an expression for expected welfare conditional on θ in the following lemma.

Lemma 5.3. *Expected welfare is given by*

$$E[W(a, \theta)|\theta] = -\lambda_1(1-a)(1-\frac{r}{2})\omega - (1-r)(1-\lambda_1)^2\tau^2 - \lambda_1^2(\sigma^2 + \nu^2(1-\frac{r}{2})) \quad (5.12)$$

Our first result tells us that increasing the precision of the public signal unambiguously increases social welfare.

Proposition 5.8. *Expected social welfare is decreasing in the variance of the public signal.*

Now we see what happens with welfare when the variance of the idiosyncratic noise in the private signal increases. As we will see, when the variance of the public information is high, increasing the variance of the idiosyncratic noise will reduce welfare, because agents heavily rely on their private signals. However, if the variance of the public information is low enough, the opposite effect takes place. When the τ^2 is low, increasing the variance of the idiosyncratic noise will make agents to put more weight in the public signal, increasing expected welfare.

Proposition 5.9. *If $\tau^2 \geq \frac{(1-\frac{\tau}{2})}{(1-r)}((1-a)\omega + 2\nu^2(1-b))$, then expected welfare is decreasing in the variance of the idiosyncratic noise. If $\tau^2 \leq \frac{(1-\frac{\tau}{2})(1-a)\omega}{(1-r)}$, then expected welfare is increasing in the variance of the idiosyncratic noise.*

Now I show that, for a fixed level of aggregate variance in the private signal, a shift from the variance in the idiosyncratic shock to the variance of the common shock unambiguously decreases welfare. The reason is that for a fixed level of noise in the private signal, increasing ν^2 increase coordination within opposites groups, decreasing social welfare.

Proposition 5.10. *For a fixed total variance in the private signal, expected welfare is decreasing in the variance in the common shock.*

Finally, we conclude this section by looking at the effect that bias and ideological neglect have on social welfare.

Proposition 5.11. *Expected welfare is decreasing in the size of the information bias. Furthermore, expected welfare is decreasing in bias and correlational neglect.*

Proposition 5.11 tells us that the more biased are the sources of information, the worse is the society going to be. Furthermore, it tells us that ideology -as introduced in this paper-, have a detrimental effect on the well being of society. In particular, it is always better when individuals are able to recognize their biases.

5.6 Conclusion

This paper proposes a framework in which ideology is considered as a misperception of reality that agents recognize in others but fail to acknowledge in themselves. Agents receive information that is biased and correlated among individuals within the same ideological group, and it is assumed that agents have ideological neglect. That is, agents suffer from *bias neglect* and *correlational neglect*, while correctly assessing the bias and correlation in the information sources of the other group. Agents also receive information from an unbiased publicly available source.

In this context, I show that the precision of the public and private information have an ambiguous effect in political extremism. A key factor in that relation is the degree of correlational neglect. I show that expected social welfare is increasing in the precision of the public signal and decreasing in the size of the bias of the information sources. I further show that welfare is decreasing in bias and correlational neglect.

The present paper leaves many questions open. First, it is clear that other notions of political extremism could be explored. Under other definitions of political extremism, the role of ideology could be different from that stated in this work. Second, it could be interesting to look at notion of ideology that does not assume perfect acknowledgement of the ideological bias in agents from the other group. In particular, it could be the case that agents in one group may believe that the bias of individuals in the other group is higher that it really is. This may also change the relation between political extremism (and social welfare) and bias and correlational neglect.

Third, ideological neglect may generate *overconfidence* on agents' decisions. They may put too much (or too little) attention to their private information. This effect is not discussed in this paper. Furthermore, it may be interesting to know what happens with overconfidence when the number of private signals that players observe increases. Fourth, the scope of my model could be extended to look not only at the effect of ideological neglect on political extremism, but also on voting turnout and partisan identification. Finally,

I believe that bias and correlational neglect could be measured using data. Hence, it may be possible to assess empirically their impact in political extremism.

5.7 Appendix

5.7.1 Proof of Lemma 5.1

Lemma 5.1. The expectation of agent h_i of state of nature θ , signal x_{h_j} , and signal $x_{h'_j}$, given the observation of signal x_{h_i} , are respectively given by

- (i) $\tilde{E}[\theta|x_{h_i}] = \beta_h \hat{x}_{h_i} + (1 - \beta_h)\mu$
- (ii) $\tilde{E}[x_{h_j}|x_{h_i}] = a_h \omega_h + (\beta_h + \gamma_h) \hat{x}_{h_i} + (1 - \beta_h - \gamma_h)\mu$
- (iii) $\tilde{E}[x_{h'_j}|x_{h_i}] = \omega_{h'} + \beta_h \hat{x}_{h_i} + (1 - \beta_h)\mu$

where $\beta_h = \frac{\tau^2}{\sigma_h^2 + b_h \nu_h^2 + \tau^2}$ and $\gamma_h = \frac{b_h \nu_h^2}{\sigma_h^2 + b_h \nu_h^2 + \tau^2}$.

Proof. First, equation (i) is given by standard Bayesian updating, adjusted to the particular form of misperception in this model. That is, we have that

$$E[\theta|x_{h_i}] = \left(\frac{\tau^2}{\sigma_h^2 + b_h \nu_h^2 + \tau^2} \right) \hat{x}_{h_i} + \left(\frac{\sigma_h^2 + b_h \nu_h^2}{\sigma_h^2 + b_h \nu_h^2 + \tau^2} \right) \mu$$

Then, by standard results in bi-variate normal distributions, we know that

$$\begin{aligned} E[x_{h_j}|x_{h_i}, \theta] &= E[x_{h_j}|\theta] + \left(\frac{b_h \nu_h^2}{\sigma_h^2 + b_h \nu_h^2} \right) (x_{h_i} - E[x_{h_i}|\theta]) \\ &= \left(\frac{b_h \nu_h^2}{\sigma_h^2 + b_h \nu_h^2} \right) E[x_{h_j}|\theta] + \left(\frac{\sigma_h^2}{\sigma_h^2 + b_h \nu_h^2} \right) (\theta + a_h \omega_h) \end{aligned}$$

By the law of iterated expectations, we obtain

$$\begin{aligned} E[E[x_{h_j}|x_{h_i}, \theta]|x_{h_i}] &= \left(\frac{b_h \nu_h^2}{\sigma_h^2 + b_h \nu_h^2} \right) E[x_{h_j}|\theta] + \left(\frac{\sigma_h^2}{\sigma_h^2 + b_h \nu_h^2} \right) (E[\theta|x_{h_i}] + a_h \omega_h) \\ &= a_h \omega_h + \left(\frac{\tau^2 + b_h \nu_h^2}{\sigma_h^2 + b_h \nu_h^2 + \tau^2} \right) \hat{x}_{h_i} + \left(\frac{\sigma_h^2}{\sigma_h^2 + b_h \nu_h^2 + \tau^2} \right) \mu \end{aligned}$$

where the last inequality follows from replacing $E[\theta|x_{h_i}]$ and rearranging terms. This gives equation (ii). Finally, since the best predictor of $x_{h'_j}$ is $\omega_{h'}$ plus $E[\theta|x_{h_i}]$, we obtain equation (iii). \square

5.7.2 Proof of Proposition 5.1

Proposition 5.1. The unique linear equilibrium is given by optimal actions $a_{h_i}^* = \lambda_h^0 + \lambda_h^1 \hat{x}_{h_i} + (1 - \lambda_h^1)y$, where

$$\begin{aligned} \lambda_h^0 &= \frac{r m_{h'}}{1 - r} [(1 - a_h) \omega_h \lambda_h^1 r m_h + (1 - a_{h'}) \omega_{h'} \lambda_{h'}^1 (1 - r m_{h'})] \\ \lambda_h^1 &= \frac{(1 - r) \beta_h}{1 - r m_h (\beta_h + \gamma_h) - \delta_h r m_{h'} \beta_{h'}} \end{aligned}$$

with $\delta_h \equiv \frac{1 - r m_h \gamma_h}{1 - r m_{h'} \gamma_{h'}}$.

Proof. First, we obtain an expression for $\hat{E}[\bar{a}|x_{U_i}]$ in terms of λ_h^1 and λ_h^2 .

$$\begin{aligned} \hat{E}[\bar{a}|x_{U_i}] &= E \left[\int_{j \in m_U} a_{U_j}^* dj + \int_{j \in m_D} a_{D_j}^* dj \mid \hat{x}_{U_i} \right] \\ &= E \left[\sum_h \int_{j \in m_h} (\lambda_h^0 + \lambda_h^1 \hat{x}_{h_i} + (1 - \lambda_h^1)y) dj \mid \hat{x}_{U_i} \right] \\ &= \sum_h m_h (\lambda_h^0 + (1 - \lambda_h^1)y) + \sum_h \int_{j \in m_h} E[\hat{x}_{h_j} | \hat{x}_{U_i}] \end{aligned}$$

Using the results of Lemma 5.1, we obtain

$$\begin{aligned}
\hat{E}[\bar{a}|x_{U_i}] &= \sum_h m_h(\lambda_h^0 + (1 - \lambda_h^1)y) + (m_D\lambda_D^1\beta_U + m_U\lambda_U^1(\beta_U + \gamma_U))\hat{x}_{U_i} \\
&+ (m_D\lambda_D^1(1 - \beta_U) + m_U\lambda_U^1(1 - \beta_U - \gamma_U))y \\
&= (m_U\lambda_U^0 + m_D(\lambda_D^0 + \lambda_D^1(1 - a_D)\omega_D)) + (m_D\lambda_D^1\beta_U + m_U\lambda_U^1(\beta_U + \gamma_U))\hat{x}_{U_i} \\
&+ (1 - m_D\lambda_D^1\beta_U - m_U\lambda_U^1(\beta_U + \gamma_U))y
\end{aligned}$$

Then, we replace this expression in (5.6) to use again Lemma (5.1) to obtain

$$\begin{aligned}
a_{U_i}^* &= r(m_U\lambda_U^0 + m_D(\lambda_D^0 + \lambda_D^1(1 - a_D)\omega_D)) \\
&+ ((1 - r)\beta_U + r(m_D\lambda_D^1\beta_U + m_U\lambda_U^1(\beta_U + \gamma_U)))\hat{x}_{U_i} \\
&+ (1 - ((1 - r)\beta_U + r(m_D\lambda_D^1\beta_U + m_U\lambda_U^1(\beta_U + \gamma_U))))y \\
&= \kappa_U^0 + \kappa_U^1\hat{x}_{U_i} + (1 - \kappa_U^1)y
\end{aligned}$$

where $\kappa_h^0 = r \sum_h m_h(\lambda_h^0 + \lambda_h^1) + r\lambda_{h'}^1(1 - a_{h'})\omega_{h'}$ and $\kappa_h^1 = (1 - r)\beta_h + r\lambda_h^1 m_h(\beta_h + \gamma_h) + r\lambda_{h'}^1 m_{h'}\beta_{h'}$. Doing the same for action $a_{D_i}^*$, and equating coefficients with those in equation (5.6), we obtain the following system of equations.

$$\begin{aligned}
\lambda_U^1 &= (1 - r)\beta_U + r(m_D\lambda_D^1\beta_U + m_U\lambda_U^1(\beta_U + \gamma_U)) \\
\lambda_D^1 &= (1 - r)\beta_D + r(m_U\lambda_U^1\beta_U + m_D\lambda_D^1(\beta_D + \gamma_D)) \\
\lambda_U^0 &= r(m_U\lambda_U^0 + m_D(\lambda_D^0 + \lambda_D^1(1 - a_D)\omega_D)) \\
\lambda_D^0 &= r(m_D\lambda_D^0 + m_U(\lambda_U^0 + \lambda_U^1(1 - a_U)\omega_U))
\end{aligned}$$

Solving the first two equations yields

$$\lambda_U^1 = \frac{(1 - r)\beta_U}{1 - rm_U(\beta_U + \gamma_U) - \delta_U rm_D\beta_D} \quad (5.13)$$

$$\lambda_D^1 = \frac{(1 - r)\beta_D}{1 - rm_D(\beta_D + \gamma_D) - \delta_D rm_U\beta_U} \quad (5.14)$$

where $\delta_U \equiv \frac{1 - rm_U\gamma_U}{1 - rm_D\gamma_D}$, and $\delta_D = \delta_U^{-1}$. Finally, solving the third and fourth equation, we obtain

$$\lambda_U^0 = \frac{rm_D}{1 - r} [(1 - a_U)\omega_U rm_U\lambda_U^1 + (1 - a_D)\omega_D\lambda_D^1(1 - rm_D)] \quad (5.15)$$

$$\lambda_D^0 = \frac{rm_U}{1 - r} [(1 - a_D)\omega_D rm_D\lambda_D^1 + (1 - a_U)\omega_U\lambda_U^1(1 - rm_U)] \quad (5.16)$$

Finally, the equilibrium we have found is unique because $r \in (0, 1)$ and so $\frac{\frac{\partial^2 g_{h_i}}{\partial a_{h_i} \partial \bar{a}}}{\frac{\partial^2 g_{h_i}}{\partial a_{h_i}^2}} = -\frac{1}{2-r} > -1$. See Angelatos and Pavan (2007). □

5.7.3 Proof of Proposition 5.2

Proposition 5.2. The weight that agent h_i puts on his private signal is increasing in the variance of the public signal, and decreasing in the variance of the idiosyncratic shock of the private signal.

Proof. To prove this proposition, we just have to take the derivative of λ_h^1 with respect to τ^2 and $\sigma_{h'}^2$. We start with the first derivative.

$$\begin{aligned}
\frac{\partial \lambda_h^1}{\partial \tau^2} &= \frac{(1 - r) \frac{\partial \beta_h}{\partial \tau^2}}{1 - rm_h(\beta_h + \gamma_h) - \delta_h rm_{h'}\beta_{h'}} \\
&+ \frac{(1 - r)\beta_{h'}r \left(m_h \left(\frac{\partial \beta_h}{\partial \tau^2} + \frac{\partial \gamma_h}{\partial \tau^2} \right) + m_{h'} \left(\frac{\partial \delta_h}{\partial \tau^2} \beta_{h'} + \delta_h \frac{\partial \beta_{h'}}{\partial \tau^2} \right) \right)}{(1 - rm_h(\beta_h + \gamma_h) - \delta_h rm_{h'}\beta_{h'})^2}
\end{aligned}$$

Since $\frac{\partial \beta_h}{\partial \tau^2} > 0$ and $\frac{\partial \beta_h}{\partial \tau^2} + \frac{\partial \gamma_h}{\partial \tau^2} > 0$ we need to check that

$$A \equiv m_{h'} \left(\frac{\partial \delta_h}{\partial \tau^2} \beta_{h'} + \delta_h \frac{\partial \beta_{h'}}{\partial \tau^2} \right) > 0 \quad (5.17)$$

For this, first note that

$$\begin{aligned} \frac{\partial \delta_h}{\partial \tau^2} &= -\frac{rm_h \frac{\partial \gamma_h}{\partial \tau^2}}{1 - rm_{h'} \gamma_{h'}} - \frac{(1 - rm_h \gamma_h) \left(-rm_{h'} \frac{\partial \gamma_{h'}}{\partial \tau^2} \right)}{(1 - rm_{h'} \gamma_{h'})^2} \\ &= \frac{r}{(1 - rm_{h'} \gamma_{h'})} \left(\delta_h m_{h'} \frac{\partial \gamma_{h'}}{\partial \tau^2} - m_h \frac{\partial \gamma_h}{\partial \tau^2} \right) \end{aligned}$$

Replacing the above expression into (5.17), we obtain

$$\begin{aligned} A &= \frac{rm_{h'} \beta_{h'}}{(1 - rm_{h'} \gamma_{h'})} \left(\delta_h m_{h'} \frac{\partial \gamma_{h'}}{\partial \tau^2} - m_h \frac{\partial \gamma_h}{\partial \tau^2} \right) + m_{h'} \delta_h \frac{\partial \beta_{h'}}{\partial \tau^2} \\ &= -\frac{rm_{h'} \beta_{h'}}{(1 - rm_{h'} \gamma_{h'})} \left(m_h \frac{\partial \gamma_h}{\partial \tau^2} \right) + \delta_h m_{h'} \left(\frac{\partial \beta_{h'}}{\partial \tau^2} + \frac{rm_{h'} \beta_{h'}}{(1 - rm_{h'} \gamma_{h'})} \frac{\partial \gamma_{h'}}{\partial \tau^2} \right) \end{aligned}$$

Since $\frac{\partial \gamma_{h'}}{\partial \tau^2} < 0$, we need the second term above to be positive. Furthermore, since $\frac{\partial \beta_{h'}}{\partial \tau^2} = \frac{\sigma_{h'}^2}{\sigma_{h'} + b_{h'} \nu_{h'}^2 + \tau^2} - \frac{\partial \gamma_{h'}}{\partial \tau^2}$, the second term becomes

$$\begin{aligned} B &\equiv \delta_h m_{h'} \left(\frac{\sigma_{h'}^2}{\sigma_{h'} + b_{h'} \nu_{h'}^2 + \tau^2} - \frac{\partial \gamma_{h'}}{\partial \tau^2} + \frac{rm_{h'} \beta_{h'}}{(1 - rm_{h'} \gamma_{h'})} \frac{\partial \gamma_{h'}}{\partial \tau^2} \right) \\ &= \delta_h m_{h'} \left(\frac{\sigma_{h'}^2}{\sigma_{h'} + b_{h'} \nu_{h'}^2 + \tau^2} - \frac{\partial \gamma_{h'}}{\partial \tau^2} \left(1 - \frac{rm_{h'} \beta_{h'}}{(1 - rm_{h'} \gamma_{h'})} \right) \right) \\ &= \delta_h m_{h'} \left(\frac{\sigma_{h'}^2}{\sigma_{h'} + b_{h'} \nu_{h'}^2 + \tau^2} - \frac{\partial \gamma_{h'}}{\partial \tau^2} \left(\frac{1 - rm_{h'} (\beta_{h'} + \gamma_{h'})}{(1 - rm_{h'} \gamma_{h'})} \right) \right) \end{aligned}$$

And since $1 - rm_{h'} (\beta_{h'} + \gamma_{h'}) > 0$, we obtain $B > 0$, which concludes the proof of the first part of the proposition.

Now we want to show that $\frac{\partial \lambda_h^1}{\partial \sigma_h} < 0$.

$$\begin{aligned} \frac{\partial \lambda_h^1}{\partial \sigma_h^2} &= \frac{(1 - r) \frac{\partial \beta_h}{\partial \sigma_h^2}}{1 - rm_h (\beta_h + \gamma_h) - \delta_h rm_{h'} \beta_{h'}} \\ &\quad + \frac{(1 - r) \beta_h r \left(m_h \left(\frac{\partial \beta_h}{\partial \sigma_h^2} + \frac{\partial \gamma_h}{\partial \sigma_h^2} \right) + m_{h'} \beta_{h'} \frac{\partial \delta_h}{\partial \sigma_h^2} \right)}{(1 - rm_h (\beta_h + \gamma_h) - \delta_h rm_{h'} \beta_{h'})^2} \end{aligned}$$

Since $\frac{\partial \beta_h}{\partial \sigma_h^2} < 0$, we need to check that

$$C \equiv m_h \left(\frac{\partial \beta_h}{\partial \sigma_h^2} + \frac{\partial \gamma_h}{\partial \sigma_h^2} \right) + m_{h'} \beta_{h'} \frac{\partial \delta_h}{\partial \sigma_h^2} < 0 \quad (5.18)$$

By noticing that $\frac{\partial \delta_h}{\partial \sigma_h^2} = -\frac{rm_h \frac{\partial \gamma_h}{\partial \sigma_h^2}}{1 - rm_{h'} \gamma_{h'}}$, we obtain

$$\begin{aligned} C &= m_h \left(\frac{\partial \beta_h}{\partial \sigma_h^2} + \frac{\partial \gamma_h}{\partial \sigma_h^2} \right) - m_{h'} \beta_{h'} \frac{rm_h \frac{\partial \gamma_h}{\partial \sigma_h^2}}{1 - rm_{h'} \gamma_{h'}} \\ &= m_h \frac{\partial \beta_h}{\partial \sigma_h^2} + m_h \frac{\partial \gamma_h}{\partial \sigma_h^2} \left(1 - \frac{rm_{h'} \beta_{h'}}{1 - rm_{h'} \gamma_{h'}} \right) \\ &= m_h \frac{\partial \beta_h}{\partial \sigma_h^2} + m_h \frac{\partial \gamma_h}{\partial \sigma_h^2} \left(\frac{1 - rm_{h'} (\beta_{h'} + \gamma_{h'})}{1 - rm_{h'} \gamma_{h'}} \right) \end{aligned}$$

Since $1 - rm_{h'} (\beta_{h'} + \gamma_{h'}) > 0$ and both $\frac{\partial \gamma_h}{\partial \sigma_h^2}$ and $\frac{\partial \beta_h}{\partial \sigma_h^2}$ are negative, we conclude that $C < 0$, which completes the proof. \square

5.7.4 Proof of Proposition 5.3

Proposition 5.3. For a fixed variance of private signals x_{h_i} , the weight agent h_i puts on his private signal is increasing in the variance of the common shock of his group.

Proof. To proof this proposition, we need to take the derivative of λ_h^1 with respect to γ_h . That is,

$$\frac{\partial \lambda_h^1}{\partial \gamma_h} = -\frac{(1-r)\beta_h r \left(m_h + m_{h'} \beta_{h'} \frac{\partial \delta_h}{\partial \gamma_h} \right)}{(1-rm_h(\beta_h + \gamma_h) - \delta_h r m_{h'} \beta_{h'})^2}$$

Hence, we need to check that $D \equiv \left(m_h + m_{h'} \beta_{h'} \frac{\partial \delta_h}{\partial \gamma_h} \right) > 0$. Since $\frac{\partial \delta_h}{\partial \gamma_h} = \frac{-rm_h}{1-rm_{h'}\gamma_{h'}}$, we obtain

$$\begin{aligned} D &= \left(m_h + m_{h'} \beta_{h'} \frac{-rm_h}{1-rm_{h'}\gamma_{h'}} \right) \\ &= m_h \left(1 - \frac{rm_{h'}\beta_{h'}}{1-rm_{h'}\gamma_{h'}} \right) \\ &= m_h \left(\frac{1-rm_{h'}(\beta_{h'} + \gamma_{h'})}{1-rm_{h'}\gamma_{h'}} \right) > 0 \end{aligned}$$

Which concludes the proof. □

5.7.5 Proof of Proposition 5.4

Proposition 5.4. The weight agent h_i puts on his private signal is increasing in the correlational neglect of agents in his group, and increasing the correlational neglect of agents in the other group.

Proof. As before, we take the derivative of λ_h^1 with respect to b_h , and obtain

$$\begin{aligned} \frac{\partial \lambda_h^1}{\partial b_h} &= \frac{(1-r) \frac{\partial \beta_h}{\partial b_h}}{1-rm_h(\beta_h + \gamma_h) - \delta_h r m_{h'} \beta_{h'}} \\ &+ \frac{(1-r)\beta_h r \left(m_h \left(\frac{\partial \beta_h}{\partial b_h} + \frac{\partial \gamma_h}{\partial b_h} \right) + m_{h'} \beta_{h'} \frac{\partial \delta_h}{\partial b_h} \right)}{(1-rm_h(\beta_h + \gamma_h) - \delta_h r m_{h'} \beta_{h'})^2} \end{aligned}$$

It is easy to see that $\frac{\partial \beta_h}{\partial b_h} = -\frac{\beta_h \gamma_h}{b_h}$, $\frac{\partial \gamma_h}{\partial b_h} = \frac{\gamma_h(1-\gamma_h)}{b_h}$, so that $\frac{\partial \beta_h}{\partial b_h} + \frac{\partial \gamma_h}{\partial b_h} = \frac{\gamma_h(1-\gamma_h-\beta_h)}{b_h}$. Furthermore, $\frac{\partial \delta_h}{\partial b_h} = -\frac{rm_h \gamma_h(1-\gamma_h)}{b_h(1-rm_{h'}\gamma_{h'})}$. Hence, we obtain

$$\begin{aligned} \frac{\partial \lambda_h^1}{\partial b_h} &= \frac{(1-r)\beta_h \gamma_h}{b_h(1-rm_h(\beta_h + \gamma_h) - \delta_h r m_{h'} \beta_{h'})} \\ &+ \frac{(1-r)\beta_h r \left(\frac{m_h \gamma_h(1-\gamma_h-\beta_h)}{b_h} - m_{h'} \beta_{h'} \frac{rm_h \gamma_h(1-\gamma_h)}{b_h(1-rm_{h'}\gamma_{h'})} \right)}{(1-rm_h(\beta_h + \gamma_h) - \delta_h r m_{h'} \beta_{h'})^2} \end{aligned}$$

Rearranging terms, we obtain

$$\begin{aligned} \frac{\partial \lambda_h^1}{\partial b_h} &= -\lambda_h^1 \frac{\gamma_h}{b_h} + \frac{(\lambda_h^1)^2 rm_h \gamma_h}{(1-r)\beta_h b_h} \left(1 - \gamma_h - \beta_h - \frac{rm_{h'}\gamma_{h'}(1-\gamma_h)}{b_h(1-rm_{h'}\gamma_{h'})} \right) \\ &= -\frac{\lambda_h^1 \gamma_h}{(1-r)\beta_h b_h} \left((1-r)\beta_h - \lambda_h^1 rm_h + \lambda_h^1 rm_h \left(\gamma_h + \beta_h + \frac{rm_{h'}\beta_{h'}(1-\gamma_h)}{1-rm_{h'}\gamma_{h'}} \right) \right) \end{aligned}$$

Finally, since

$$(1-r)\beta_h - \lambda_h^1 rm_h = (1-r)\beta_h \left(1 - \frac{rm_h}{(1-rm_h(\beta_h + \gamma_h) - \delta_h r m_{h'} \beta_{h'})} \right)$$

and

$$1 - rm_h(\beta_h + \gamma_h) - \delta_h rm_{h'}\beta_{h'} - rm_h \geq 0$$

which implies that $(1 - r)\beta_h - rm_h\lambda_h^1 \geq 0$, so that $\frac{\partial \lambda_h^1}{\partial b_h} \leq 0$.

Then, we take the derivative of λ_h^1 with respect to $b_{h'}$

$$\frac{\partial \lambda_h^1}{\partial b_{h'}} = \frac{(1 - r)r\beta_h m_{h'} \left(\frac{\partial \delta_h}{\partial b_{h'}} \beta_{h'} + \delta_h \frac{\partial \beta_{h'}}{\partial b_{h'}} \right)}{(1 - rm_h(\beta_h + \gamma_h) - \delta_h rm_{h'}\beta_{h'})^2}$$

So we need to check that

$$\frac{\partial \delta_h}{\partial b_{h'}} \beta_{h'} + \delta_h \frac{\partial \beta_{h'}}{\partial b_{h'}} \leq 0$$

Indeed,

$$\frac{\partial \delta_h}{\partial b_{h'}} = \delta_h rm_{h'} \frac{\gamma_{h'}(1 - \gamma_{h'})}{b_{h'}(1 - rm_{h'}\gamma_{h'})}$$

So that we obtain

$$\begin{aligned} \frac{\partial \delta_h}{\partial b_{h'}} \beta_{h'} + \delta_h \frac{\partial \beta_{h'}}{\partial b_{h'}} &= \delta_h rm_{h'} \frac{\gamma_{h'}(1 - \gamma_{h'})}{b_{h'}(1 - rm_{h'}\gamma_{h'})} \beta_{h'} + \delta_h \frac{\beta_{h'}\gamma_{h'}}{b_{h'}} \\ &= -\frac{\delta_{h'}\beta_{h'}\gamma_{h'}}{b_{h'}} \left(1 - \frac{rm_{h'}(1 - \gamma_{h'})}{1 - rm_{h'}\gamma_{h'}} \right) \end{aligned}$$

Finally,

$$\begin{aligned} \left(1 - \frac{rm_{h'}(1 - \gamma_{h'})}{1 - rm_{h'}\gamma_{h'}} \right) &= \frac{1 - rm_{h'}\gamma_{h'} - rm_{h'}(1 - \gamma_{h'})}{1 - rm_{h'}\gamma_{h'}} \\ &= \frac{1 - rm_{h'}}{1 - rm_{h'}\gamma_{h'}} > 0 \end{aligned}$$

So we conclude that $\frac{\partial \lambda_h^1}{\partial b_{h'}} \leq 0$ □

5.7.6 Proof of Lemma 5.2

Here I provide a proof to Lemma 5.2 for the general case, stated as Lemma 5.2b. Proof of Lemma 5.2 is provided below using Lemma 5.2b.

Lemma 5.2b. Political extremism, defined as the variance of the average action at equilibrium, is given by

$$P \equiv m_U^2(\lambda_U^1)^2\nu_U^2 + m_D^2(\lambda_D^1)^2\nu_D^2 + (1 - m_U\lambda_U^1 - m_D\lambda_D^1)^2\tau^2 \quad (5.19)$$

Proof. First, we obtain an expression for $V[\bar{a}|\theta]$. Recall first that $\hat{x}_{h_i} = \theta + (1 - a_h)\omega_h + \epsilon_{h_i} + (\phi_h - \omega_h)$ and $y = \theta + \eta$, so the optimal action can be written as

$$a_{h_i}^* = \theta + \kappa_h + \lambda_h^1(\epsilon_{h_i} + \hat{\phi}_h) + (1 - \lambda_h^1)\eta \quad (5.20)$$

where $\hat{\phi}_h = \phi_h - \omega_h$ and $\kappa_h = \lambda_h^0 + \lambda_h^1(1 - a_h)\omega_h$. Replacing the above expression into \bar{a} , we obtain

$$\begin{aligned} \bar{a} &= \theta + m_U\kappa_U + m_D\kappa_D + \lambda_U^1 \left(\int_{m_U} \epsilon_{U_j} dj + m_U\hat{\phi}_U \right) + m_U(1 - \lambda_U^1)\eta \\ &+ \lambda_D^1 \left(\int_{m_D} \epsilon_{D_j} dj + m_D\hat{\phi}_D \right) + m_D(1 - \lambda_D^1)\eta \end{aligned} \quad (5.21)$$

So that we obtain

$$E[\bar{a}|\theta] = \theta + m_U \kappa_U + m_D \kappa_D \quad (5.22)$$

Then,

$$\begin{aligned} V[\bar{a}|\theta] &= E \left[(\bar{a} - E[\bar{a}|\theta])^2 | \theta \right] \\ &= E \left[\left(\lambda_U^1 \left(\int_{m_U} \epsilon_{U_j} dj + m_U \hat{\phi}_U \right) + m_U (1 - \lambda_U^1) \eta \right. \right. \\ &\quad \left. \left. + \lambda_D^1 \left(\int_{m_D} \epsilon_{D_j} dj + m_D \hat{\phi}_D \right) + m_D (1 - \lambda_D^1) \eta \right)^2 | \theta \right] \\ &= E \left[\left((\lambda_U^1)^2 \left(\int_{m_U} \epsilon_{U_j} dj + m_U \hat{\phi}_U \right)^2 + (\lambda_D^1)^2 \left(\int_{m_D} \epsilon_{D_j} dj + m_D \hat{\phi}_D \right)^2 \right. \right. \\ &\quad \left. \left. + m_U^2 (1 - \lambda_U^1)^2 \eta^2 + m_D^2 (1 - \lambda_D^1)^2 \eta^2 + 2m_U m_D (1 - \lambda_U^1)(1 - \lambda_D^1) \eta^2 + K \right) | \theta \right] \end{aligned}$$

where K are the rest of cross multiplication terms whose expectation is zero. Now, to simplify further, note that

$$\left(\int_{j \in m_h} \epsilon_{h_j} dj + m_h \hat{\phi}_h \right)^2 = \left(\int_{j \in m_h} \epsilon_{h_j} dj \right)^2 + 2m_h \hat{\phi}_h \int_{h_j} \epsilon_{h_j} + m_h^2 \hat{\phi}_h^2$$

Lemma 5.4.

$$E \left[\left(\int_{j \in [0, m_U]} \epsilon_{h_j} dj \right)^2 | \theta \right] = 0$$

Proof. I provide two different proofs. First,

$$\begin{aligned} \left(\int_{j \in [0, m_U]} \epsilon_{h_j} dj \right)^2 &= \left(\lim_{n \rightarrow \infty} \sum_{k=0}^n \epsilon_{\frac{km}{n}} \frac{m_U}{n} \right)^2 \\ &= \left(\lim_{n \rightarrow \infty} \sum_{k=0}^n \epsilon_{\frac{km}{n}} \frac{m_U}{n} \right) \left(\lim_{n \rightarrow \infty} \sum_{k=0}^n \epsilon_{\frac{km}{n}} \frac{m_U}{n} \right) \\ &= \lim_{n \rightarrow \infty} \left(\sum_{k=0}^n \epsilon_{\frac{km}{n}} \frac{m_U}{n} \right) \left(\sum_{k=0}^n \epsilon_{\frac{km}{n}} \frac{m_U}{n} \right) \\ &= \lim_{n \rightarrow \infty} \left(\sum_{k=0}^n \epsilon_{\frac{km}{n}}^2 \frac{m_U^2}{n^2} + 2 \sum_{k=0}^n \sum_{h=0}^k \epsilon_{\frac{km}{n}} \epsilon_{\frac{hm}{n}} \frac{m_U^2}{n^2} \right) \end{aligned}$$

So that

$$\begin{aligned} E \left[\left(\int_{j \in [0, m_U]} \epsilon_{h_j} dj \right)^2 | \theta \right] &= E \left[\lim_{n \rightarrow \infty} \left(\sum_{k=0}^n \epsilon_{\frac{km}{n}}^2 \frac{m_U^2}{n^2} + 2 \sum_{k=0}^n \sum_{h=0}^k \epsilon_{\frac{km}{n}} \epsilon_{\frac{hm}{n}} \frac{m_U^2}{n^2} \right) | \theta \right] \\ &= \lim_{n \rightarrow \infty} \sigma_U^2 \frac{m_U^2}{n} \\ &= 0 \end{aligned}$$

The second proof is as follows. Note that

$$\begin{aligned} \int_{j \in [0, m_U]} \epsilon_{h_j} dj &= \lim_{n \rightarrow \infty} \sum_{k=0}^n \epsilon_{\frac{km}{n}} \frac{m_U}{n} \\ &= \lim_{n \rightarrow \infty} \frac{m_U}{n} \sum_{k=0}^n \epsilon_{\frac{km}{n}} \end{aligned}$$

and $\frac{m_U}{n} \sum_{k=0}^n \epsilon_{km}$ is the sample mean, so by the law of large numbers,

$$\int_{j \in [0, m_U]} \epsilon_{h_j} dj = 0$$

□

Consequently,

$$\begin{aligned} E \left[\left(\int_{h_j} \epsilon_{h_j} dh_j + m_h \hat{\phi}_h \right)^2 \middle| \theta \right] &= m_h^2 E \left[\hat{\phi}_h^2 \middle| \theta \right] \\ &= m_h^2 \nu_h^2 \end{aligned}$$

Hence,

$$\begin{aligned} V[\bar{a}|\theta] &= m_U^2 [(\lambda_U^1)^2 \nu_U^2 + (1 - \lambda_U^1)^2 \tau^2] + m_D^2 [(\lambda_D^1)^2 \nu_D^2 + (1 - \lambda_D^1)^2 \tau^2] \\ &\quad + 2m_U m_D (1 - \lambda_U^1)(1 - \lambda_D^1) \tau^2 \\ &= m_U^2 (\lambda_U^1)^2 \nu_U^2 + m_D^2 (\lambda_D^1)^2 \nu_D^2 + (1 - m_U \lambda_U^1 - m_D \lambda_D^1)^2 \tau^2 \end{aligned} \quad (5.23)$$

Which concludes the proof. □

5.7.7 Proof of Lemma 5.2

Political extremism, defined as the variance of the average action at equilibrium, is given by

$$P \equiv \frac{1}{2} \lambda_1^2 \nu^2 + (1 - \lambda_1)^2 \tau^2$$

Proof. Trivial by replacing $\sigma_U^2 = \sigma_D^2 = \sigma^2$, $\nu_U^2 = \nu_D^2 = \nu^2$, $b_U = b_D = b$ and $m_U = m_D = \frac{1}{2}$ in equation (5.19). □

5.7.8 Proof of Proposition 5.5

Proposition 5.5 If $\frac{\sigma^2}{\nu^2} \leq ((1 - r) - b(1 - \frac{r}{2}))$, political extremism is increasing in the variance of the public signal. If $\frac{\sigma^2}{\nu^2} > ((1 - r) - b(1 - \frac{r}{2}))$, there exists a $\hat{\tau}^2$ such that for $\tau^2 < \hat{\tau}^2$ ($\tau^2 > \hat{\tau}^2$) political extremism is increasing (decreasing) in the variance of the public signal.

Proof. We take the derivative of expression (5.10) with respect to τ^2

$$\frac{\partial P}{\partial \tau^2} = \frac{\partial \lambda_1}{\partial \tau^2} \left(\lambda_1 \nu^2 - 2(1 - \lambda_1) \tau^2 + \frac{(1 - \lambda_1)^2}{\frac{\partial \lambda_1}{\partial \tau^2}} \right)$$

But

$$\begin{aligned} \lambda_1 &= \frac{(1 - r)\beta}{1 - \frac{r}{2}(2\beta + \gamma)} = \frac{(1 - r)\tau^2}{\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2} \\ 1 - \lambda_1 &= \frac{(\sigma^2 + b\nu^2(1 - \frac{r}{2}))}{(\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2)} \end{aligned}$$

So

$$\frac{\partial \lambda_1}{\partial \tau^2} = \frac{(1 - r)(\sigma^2 + b\nu^2(1 - \frac{r}{2}))}{(\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2)^2}$$

and

$$\frac{(1 - \lambda_1)^2}{\frac{\partial \lambda_1}{\partial \tau^2}} = \frac{\sigma^2 + b\nu^2(1 - \frac{r}{2})}{(1 - r)}$$

Hence,

$$\begin{aligned} \frac{\partial P}{\partial \tau^2} &= K \left((1 - r)^2 \tau^2 \nu^2 - 2\tau^2(1 - r)(\sigma^2 + b\nu^2(1 - \frac{r}{2})) \right. \\ &\quad \left. + (\sigma^2 + b\nu^2(1 - \frac{r}{2}))(\sigma^2 + b\nu^2(1 - \frac{r}{2}) + \tau^2(1 - r)) \right) \end{aligned}$$

with $K \equiv \frac{\frac{\partial \lambda_1}{\partial \tau^2}}{(\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2)(1 - r)}$. Consequently, $\frac{\partial P}{\partial \tau^2} \geq 0$ if and only if

$$\begin{aligned} 0 &\leq \left((1 - r)^2 \tau^2 \nu^2 - 2\tau^2(1 - r)(\sigma^2 + b\nu^2(1 - \frac{r}{2})) \right. \\ &\quad \left. + (\sigma^2 + b\nu^2(1 - \frac{r}{2}))(\sigma^2 + b\nu^2(1 - \frac{r}{2}) + \tau^2(1 - r)) \right) \\ \left(\sigma^2 + b\nu^2(1 - \frac{r}{2}) \right)^2 &\geq \tau^2(1 - r)(\sigma^2 + b\nu^2(1 - \frac{r}{2})) - (1 - r)^2 \nu^2 \tau^2 \\ \left(\sigma^2 + b\nu^2(1 - \frac{r}{2}) \right)^2 &\geq \tau^2(1 - r)(\sigma^2 + \nu^2(b(1 - \frac{r}{2}) - (1 - r))) \end{aligned}$$

Consequently, if $\sigma^2 + \nu^2(b(1 - \frac{r}{2}) - (1 - r)) \leq 0$, $\frac{\partial P}{\partial \tau^2} \geq 0$ for all τ^2 . If $\sigma^2 + \nu^2(b(1 - \frac{r}{2}) - (1 - r)) \geq 0$, $\frac{\partial P}{\partial \tau^2} \geq 0$ if and only if

$$\tau^2 \leq \frac{(\sigma^2 + b\nu^2(1 - \frac{r}{2}))^2}{(1 - r)(\sigma^2 + \nu^2(b(1 - \frac{r}{2}) - (1 - r)))}$$

□

5.7.9 Proof of Proposition 5.6

Proposition 5.6. If $b \geq \frac{(1-r)}{4(1-\frac{r}{2})}$, political extremism is increasing in the variance of the common shock. If $0 < b < \frac{(1-r)}{4(1-\frac{r}{2})}$, there exists a $\hat{\nu}^2$ such that for $\nu^2 \geq \hat{\nu}^2$ ($\nu^2 \leq \hat{\nu}^2$) political extremism is decreasing (increasing) in the variance of the common shock.

Proof. We take the derivative of expression (5.10) with respect to ν^2

$$\frac{\partial P}{\partial \tau^2} = \frac{\partial \lambda_1}{\partial \nu^2} \left(\lambda_1 \nu^2 - 2(1 - \lambda_1)\tau^2 + \frac{\lambda_1^2}{2\frac{\partial \lambda_1}{\partial \nu^2}} \right)$$

Again, use

$$\begin{aligned} \lambda_1 &= \frac{(1 - r)\tau^2}{\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2} \\ 1 - \lambda_1 &= \frac{(\sigma^2 + b\nu^2(1 - \frac{r}{2}))}{(\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2)} \\ \text{and} \\ \frac{\partial \lambda_1}{\partial \nu^2} &= -\frac{(1 - r)\tau^2 b(1 - \frac{r}{2})}{(\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2)^2} \end{aligned}$$

to obtain

$$\begin{aligned} \frac{\partial P}{\partial \tau^2} &= \frac{\partial \lambda_1}{\partial \nu^2} \left(\frac{(1 - r)\tau^2 \nu^2}{\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2} - \frac{2\tau^2(\sigma^2 + b\nu^2(1 - \frac{r}{2}))}{\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2} - \frac{(1 - r)\tau^2}{2b(1 - \frac{r}{2})} \right) \\ &= L \left((1 - r)\nu^2 2b(1 - \frac{r}{2}) - 2(\sigma^2 + b\nu^2(1 - \frac{r}{2}))2b(1 - \frac{r}{2}) \right. \\ &\quad \left. - (1 - r) \left(\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2 \right) \right) \end{aligned}$$

where

$$L \equiv \frac{\frac{\partial \lambda_1}{\partial \nu^2} \tau^2}{2b(1 - \frac{r}{2})(\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2)}$$

Since $\frac{\partial \lambda_1}{\partial \nu^2} \leq 0$, for $\frac{\partial P}{\partial \tau^2} \leq 0$, we need

$$\begin{aligned} 0 \leq & L \left((1 - r)\nu^2 2b(1 - \frac{r}{2}) - 2(\sigma^2 + b\nu^2(1 - \frac{r}{2})) 2b(1 - \frac{r}{2}) \right. \\ & \left. - (1 - r) \left(\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2 \right) \right) \end{aligned}$$

which happens when

$$\nu^2 b \left(1 - \frac{r}{2}\right) \left((1 - r) - 4b \left(1 - \frac{r}{2}\right) \right) \geq 4b\sigma^2 \left(1 - \frac{r}{2}\right) + (1 - r)(\sigma^2 + \tau^2(1 - r))$$

Hence,

$$(1 - r) - 4b \left(1 - \frac{r}{2}\right) \leq 0 \Leftrightarrow b \geq \frac{1 - r}{4 \left(1 - \frac{r}{2}\right)} \Rightarrow \frac{\partial P}{\partial \tau^2} \geq 0$$

If $0 < b < \frac{1 - r}{4 \left(1 - \frac{r}{2}\right)}$, then for

$$\nu^2 \geq \hat{\nu}^2 \equiv \frac{4b\sigma^2 \left(1 - \frac{r}{2}\right) + (1 - r)(\sigma^2 + \tau^2(1 - r))}{b \left(1 - \frac{r}{2}\right) \left((1 - r) - 4b \left(1 - \frac{r}{2}\right) \right)}$$

we have $\frac{\partial P}{\partial \tau^2} \leq 0$, and when $\nu^2 \leq \hat{\nu}^2$, $\frac{\partial P}{\partial \tau^2} \geq 0$. □

5.7.10 Proof of Proposition 5.7

Proposition 5.7. If $\frac{\sigma^2}{\nu^2} \geq \frac{(1 - r) - 2b(1 - \frac{r}{2})}{2}$, political extremism is increasing in the variance of the idiosyncratic shock and increasing in correlational neglect. If $\frac{\sigma^2}{\nu^2} \leq \frac{(1 - r) - 2b(1 - \frac{r}{2})}{2}$ political extremism is decreasing in the variance of the idiosyncratic shock and decreasing in correlational neglect.

Proof. We take the derivative of expression (5.10) with respect to σ^2

$$\begin{aligned} \frac{\partial P}{\partial \sigma^2} &= \frac{\partial \lambda_1}{\partial \sigma^2} (\lambda_1 \nu^2 - 2(1 - \lambda_1)\tau^2) \\ &= \frac{\tau^2 \frac{\partial \lambda_1}{\partial \sigma^2}}{\sigma^2 + b\nu^2(1 - \frac{r}{2}) + (1 - r)\tau^2} \left((1 - r)\nu^2 - 2 \left(\sigma^2 + b\nu^2 \left(1 - \frac{r}{2}\right) \right) \right) \end{aligned}$$

Since $\frac{\partial \lambda_1}{\partial \sigma^2} \leq 0$,

$$\frac{\sigma^2}{\nu^2} \geq \frac{(1 - r) - 2b(1 - \frac{r}{2})}{2} \Rightarrow \frac{\partial P}{\partial \sigma^2} \geq 0$$

and

$$\frac{\sigma^2}{\nu^2} \leq \frac{(1 - r) - 2b(1 - \frac{r}{2})}{2} \Rightarrow \frac{\partial P}{\partial \sigma^2} \leq 0$$

Finally, since

$$\frac{\partial P}{\partial b} = \frac{\partial \lambda_1}{\partial b} (\lambda_1 \nu^2 - 2(1 - \lambda_1)\tau^2) \tag{5.24}$$

and $\frac{\partial \lambda_1}{\partial b} \leq 0$, the same reasoning over b concludes the proof. □

5.7.11 Proof of Lemma 5.3

Here I provide a proof to Lemma 5.3 for the general case, stated as Lemma 5.3b. Proof of Lemma 5.3 is provided below using Lemma 5.3b.

Lemma 5.3b. Expected welfare is given by

$$E[W(a, \theta)|\theta] = r \left(V[\bar{a}|\theta] + (m_U \kappa_U + m_D \kappa_D)^2 \right) - \sum_h (m_h \kappa_h^2 + m_h V[a_{h_i}^*|\theta]) \quad (5.25)$$

where $\kappa_h = \lambda_h^0 + \lambda_h^1(1 - a_h)\omega_h$ and $V[a_h|\theta] = (\lambda_h^1)^2(\sigma_h^2 + \nu_h^2) + (1 - \lambda_h^1)^2\tau^2$.

Proof. We first rewrite the definition of welfare as follows

$$\begin{aligned} W(a, \theta) &= \sum_h \int_{m_h} g_{h_i}(a, \theta) dh_i \\ &= \int_{U_j} g_{U_j}(a, \theta) dU_j + \int_{D_j} g_{D_j}(a, \theta) dD_j \\ &= -(1-r) \left(\int_{U_j} (a_{U_j} - \theta)^2 dU_j + \int_{D_j} (a_{D_j} - \theta)^2 dD_j \right) \\ &\quad - r \left(\int_{U_j} (a_{U_j} - \bar{a})^2 dU_j + \int_{D_j} (a_{D_j} - \bar{a})^2 dD_j \right) \\ &= -(1-r) \left(\int_{U_j} a_{U_j}^2 dU_j + \int_{D_j} a_{D_j}^2 dD_j - 2\theta\bar{a} + \theta^2 \right) \\ &\quad - r \left(\int_{U_j} a_{U_j}^2 dU_j + \int_{D_j} a_{D_j}^2 dD_j - \bar{a}^2 \right) \end{aligned} \quad (5.26)$$

Hence, to find an expression for $E[W(a, \theta)|\theta]$, we need to obtain first expressions for (i) $E[\bar{a}|\theta]$, (ii) $E[\bar{a}^2|\theta]$ and (iii) $E[(a_{h_i}^*)^2|\theta]$, for $h = U, D$.

(i) Expression for $E[\bar{a}|\theta]$.

This expression was already obtain in lemma 5.2.

(ii) Expression for $E[\bar{a}^2|\theta]$.

Using equation (5.21), we obtain

$$\begin{aligned} \bar{a}^2 &= (\theta + m_U \kappa_U + m_D \kappa_D)^2 + (\lambda_U^1)^2 \left(\int_{U_j} \epsilon_{U_j} dU_j + m_U \hat{\phi}_U \right)^2 + m_U^2 (1 - \lambda_U^1)^2 \eta^2 \\ &\quad + (\lambda_D^1)^2 \left(\int_{D_j} \epsilon_{D_j} dD_j + m_D \hat{\phi}_D \right)^2 + m_D^2 (1 - \lambda_D^1)^2 \eta^2 + 2m_U m_D (1 - \lambda_U^1)(1 - \lambda_D^1) \eta^2 + K \end{aligned}$$

where K are the rest of cross multiplication terms whose expectation is zero. Now, to simplify further, recall that in lemma 5.2 we proved that

$$\begin{aligned} E \left[\left(\int_{h_j} \epsilon_{h_j} dh_j + m_h \hat{\phi}_h \right)^2 \middle| \theta \right] &= m_h^2 E \left[\hat{\phi}_h^2 \middle| \theta \right] \\ &= m_h^2 \nu_h^2 \end{aligned}$$

Hence,

$$\begin{aligned}
E[\bar{a}^2|\theta] &= (\theta + m_U\kappa_U + m_D\kappa_D)^2 + m_U^2[(\lambda_U^1)^2\nu_U^2 + (1 - \lambda_U^1)^2\tau^2] \\
&+ m_D^2[(\lambda_D^1)^2\nu_D^2 + (1 - \lambda_D^1)^2\tau^2] + 2m_Um_D(1 - \lambda_U^1)(1 - \lambda_D^1)\tau^2 \\
&= (\theta + m_U\kappa_U + m_D\kappa_D)^2 + m_U^2(\lambda_U^1)^2\nu_U^2 + m_D^2(\lambda_D^1)^2\nu_D^2 \\
&+ \tau^2(1 - m_U\lambda_U^1 - m_D\lambda_D^1) \\
&= (\theta + m_U\kappa_U + m_D\kappa_D)^2 + V[\bar{a}|\theta]
\end{aligned} \tag{5.27}$$

(iii) Expression for $E[(a_{h_i}^*)^2|\theta]$.

Using equation (5.20), it is easy to obtain

$$\begin{aligned}
E[(a_{h_i}^*)^2|\theta] &= E[(\theta + \kappa_h + \lambda_h^1(\epsilon_{h_i} + \hat{\phi}_h) + (1 - \lambda_h^1)\eta)^2|\theta] \\
&= E\left[(\theta + \kappa_h)^2 + (\lambda_h^1)^2(\epsilon_{h_i} + \hat{\phi}_h)^2 + (1 - \lambda_h^1)^2\eta^2\right. \\
&+ 2(\theta + \kappa_h)\lambda_h^1(\epsilon_{h_i} + \hat{\phi}_h) + 2(\theta + \kappa_h)(1 - \lambda_h^1)\eta \\
&+ \left.2\lambda_h^1(\epsilon_{h_i} + \hat{\phi}_h)(1 - \lambda_h^1)\eta \mid \theta\right] \\
&= (\theta + \kappa_h)^2 + (\lambda_h^1)^2(\sigma_h^2 + \nu_h^2) + (1 - \lambda_h^1)^2\tau^2
\end{aligned} \tag{5.28}$$

Note also that

$$\begin{aligned}
V[a_{h_i}^*|\theta] &= E\left[(a_{h_i}^* - E[a_{h_i}^*|\theta])^2|\theta\right] \\
&= E\left[\left(\lambda_h^1(\epsilon_{h_i} + \hat{\phi}_h) + (1 - \lambda_h^1)\eta\right)^2|\theta\right] \\
&= (\lambda_h^1)^2(\sigma_h^2 + \nu_h^2) + (1 - \lambda_h^1)^2\tau^2
\end{aligned} \tag{5.29}$$

Finally, we replace our expressions in equation (5.26). First, let

$$\begin{aligned}
A &\equiv \left(\int_{U_j} a_{U_j}^2 dU_j + \int_{D_j} a_{D_j}^2 dD_j - 2\theta\bar{a} + \theta^2\right) \\
B &\equiv \left(\int_{U_j} a_{U_j}^2 dU_j + \int_{D_j} a_{D_j}^2 dD_j - \bar{a}^2\right)
\end{aligned}$$

So that

$$\begin{aligned}
E[A|\theta] &= m_U((\theta + \kappa_U)^2 + V[a_{U_i}^*|\theta]) + m_D((\theta + \kappa_D)^2 + V[a_{D_i}^*|\theta]) \\
&- 2\theta(\theta + m_U\kappa_U + m_D\kappa_D) + \theta^2 \\
&= m_U\kappa_U^2 + m_D\kappa_D^2 + m_UV[a_{U_i}^*|\theta] + m_DV[a_{D_i}^*|\theta]
\end{aligned}$$

and

$$\begin{aligned}
E[B|\theta] &= m_U((\theta + \kappa_U)^2 + V[a_{U_i}^*|\theta]) + m_D((\theta + \kappa_D)^2 + V[a_{D_i}^*|\theta]) - E[\bar{a}^2|\theta] \\
&= m_U((\theta + \kappa_U)^2 + V[a_{U_i}^*|\theta]) + m_D((\theta + \kappa_D)^2 + V[a_{D_i}^*|\theta]) \\
&- (\theta + m_U\kappa_U + m_D\kappa_D)^2 - V[\bar{a}|\theta] \\
&= m_U\kappa_U^2 + m_D\kappa_D^2 + m_UV[a_{U_i}^*|\theta] + m_DV[a_{D_i}^*|\theta] \\
&- (V[\bar{a}|\theta] + (m_U\kappa_U + m_D\kappa_D)^2)
\end{aligned}$$

Finally,

$$\begin{aligned}
E[W(a, \theta)|\theta] &= -(1 - r)E[A|\theta] - rE[B|\theta] \\
&= -(m_U\kappa_U^2 + m_D\kappa_D^2 + m_UV[a_{U_i}^*|\theta] + m_DV[a_{D_i}^*|\theta]) \\
&+ r(V[\bar{a}|\theta] + (m_U\kappa_U + m_D\kappa_D)^2)
\end{aligned}$$

Which corresponds to equation (5.25). □

5.7.12 Proof of Lemma 5.3.

Lemma 5.3. Expected welfare is given by

$$E[W(a, \theta)|\theta] = -\lambda_1(1-a)\left(1 - \frac{r}{2}\right)\omega - (1-r)(1-\lambda_1)^2\tau^2 - \lambda_1^2(\sigma^2 + \nu^2\left(1 - \frac{r}{2}\right))$$

Proof. We then apply the conditions $\sigma_u^2 = \sigma^d = \sigma^2$, $\nu_u^2 = \nu_d^2 = \nu^2$, $b_h = b_d = b$, $m_u = m_d = \frac{1}{2}$, $a_u = a_d = a$ and $\omega_u = -\omega_d = \omega$ to equation (5.25) to obtain that expected welfare is given by

$$E[W(a, \theta)|\theta] = -\lambda_1(1-a)\left(1 - \frac{r}{2}\right)\omega - (1-r)(1-\lambda_1)^2\tau^2 - \lambda_1^2(\sigma^2 + \nu^2\left(1 - \frac{r}{2}\right))$$

□

5.7.13 Proof of Proposition 5.8

Proposition 5.8. Expected social welfare is decreasing in the variance of the public signal.

Proof. We take the derivative of expression (5.12) with respect to τ^2 and rearrange terms.

$$\begin{aligned} \frac{\partial E[W(a, \theta)|\theta]}{\partial \tau^2} &= \frac{\partial \lambda_1}{\partial \tau^2} \left(-(1-a)\left(1 - \frac{r}{2}\right) + 2(1-r)(1-\lambda_1)\tau^2 - 2\lambda_1(\sigma^2 + \nu^2\left(1 - \frac{r}{2}\right)) \right) \\ &\quad - (1-r)(1-\lambda_1)^2 \end{aligned}$$

But since $1 - \lambda_1 = \frac{1 - \frac{r\gamma}{2} - \beta}{1 - \frac{r}{2}(2\beta + \gamma)}$, the terms $A \equiv 2(1-r)(1-\lambda_1)\tau^2 - 2\lambda_1(\sigma^2 + \nu^2\left(1 - \frac{r}{2}\right))$ becomes

$$\begin{aligned} A &= \frac{2(1-r)}{1 - \frac{r}{2}(2\beta + \gamma)} \left(\left(1 - \frac{r\gamma}{2} - \beta\right)\tau^2 - \beta(\sigma^2 + \nu^2\left(1 - \frac{r}{2}\right)) \right) \\ &= \frac{2(1-r)}{1 - \frac{r}{2}(2\beta + \gamma)} \left(\frac{2(\sigma^2 + b\nu^2 + \tau^2) - rb\nu^2 - 2\tau^2}{2(\sigma^2 + b\nu^2 + \tau^2)}\tau^2 - \frac{2\tau^2}{2(\sigma^2 + b\nu^2 + \tau^2)}(\sigma^2 + \nu^2\left(1 - \frac{r}{2}\right)) \right) \\ &= C \left(2(\sigma^2 + b\nu^2 + \tau^2) - rb\nu^2 - 2\tau^2 - 2(\sigma^2 + \nu^2\left(1 - \frac{r}{2}\right)) \right) \end{aligned}$$

where

$$\begin{aligned} C &= \frac{2(1-r)}{1 - \frac{r}{2}(2\beta + \gamma)} \frac{\tau^2}{2(\sigma^2 + b\nu^2 + \tau^2)} \\ &= \lambda_1 \end{aligned}$$

Hence, obtain

$$\begin{aligned} A &= \lambda_1(2b\nu^2 - rb\nu^2 - 2\nu^2\left(1 - \frac{r}{2}\right)) \\ &= -2\lambda_1\nu^2\left(1 - \frac{r}{2}\right)(1-b) \\ &\leq 0 \end{aligned}$$

Hence, we conclude that $\frac{\partial E[W(a, \theta)|\theta]}{\partial \tau^2} \leq 0$. □

5.7.14 Proof of Proposition 5.9

Proposition 5.9. If $\tau^2 \geq \frac{(1-\frac{r}{2})}{(1-r)}((1-a)\omega + 2\nu^2(1-b))$, then expected welfare is decreasing in the variance of the idiosyncratic noise. If $\tau^2 \leq \frac{(1-\frac{r}{2})(1-a)\omega}{(1-r)}$, then expected welfare is increasing in the variance of the idiosyncratic noise.

Proof. We take the derivative of expected welfare with respect to σ^2

$$\frac{\partial E[W(a, \theta)|\theta]}{\partial \tau^2} = \frac{\partial \lambda_1}{\partial \sigma^2} \left(-(1-a)(1-\frac{r}{2}) + 2(1-r)(1-\lambda_1)\tau^2 - 2\lambda_1(\sigma^2 + \nu^2(1-\frac{r}{2})) \right) - \lambda_1^2$$

As in proposition 5.8, we obtain

$$\frac{\partial E[W(a, \theta)|\theta]}{\partial \sigma^2} = \frac{\partial \lambda_1}{\partial \sigma^2} \left(-(1-a)(1-\frac{r}{2}) - 2\lambda_1\nu^2(1-\frac{r}{2})(1-b) \right) - \lambda_1^2$$

But

$$\frac{\partial \lambda_1}{\partial \sigma^2} = \frac{(1-r)\frac{\partial \beta}{\partial \sigma^2}}{1-\frac{r}{2}(2\beta+\gamma)} + \frac{(1-r)\beta\frac{r}{2}\left(2\frac{\partial \beta}{\partial \sigma^2} + \frac{\partial \gamma}{\partial \sigma^2}\right)}{\left(1-\frac{r}{2}(2\beta+\gamma)\right)^2}$$

and

$$\begin{aligned} \frac{\partial \beta}{\partial \sigma^2} &= -\frac{\tau^2}{(\sigma^2 + b\nu^2 + \tau^2)^2} = -\frac{\beta^2}{\tau^2} \\ \frac{\partial \gamma}{\partial \sigma^2} &= -\frac{b\nu^2}{(\sigma^2 + b\nu^2 + \tau^2)^2} = -\frac{\gamma\beta}{\tau^2} \end{aligned}$$

Hence,

$$\begin{aligned} \frac{\partial \lambda_1}{\partial \sigma^2} &= -\frac{\lambda_1^2(1-\frac{r}{2}(2\beta+\gamma))}{(1-r)\tau^2} - \frac{\lambda_1^2\frac{r}{2}(2\beta+\gamma)}{(1-r)\tau^2} \\ &= -\frac{\lambda_1^2}{(1-r)\tau^2} \left(1-\frac{r}{2}(2\beta+\gamma) + \frac{r}{2}(2\beta+\gamma) \right) \\ &= -\frac{\lambda_1^2}{(1-r)\tau^2} \end{aligned}$$

Consequently, we obtain

$$\frac{\partial E[W(a, \theta)|\theta]}{\partial \sigma^2} = -\frac{\lambda_1^2}{(1-r)\tau^2} \left(-(1-a)(1-\frac{r}{2}) - 2\lambda_1\nu^2(1-\frac{r}{2})(1-b) + (1-r)\tau^2 \right)$$

Finally, since $\lambda_1 \leq 1$, $\tau^2 \geq \frac{(1-\frac{r}{2})}{(1-r)}((1-a)\omega + 2\nu^2(1-b))$ implies that $\frac{\partial E[W(a, \theta)|\theta]}{\partial \sigma^2} \leq 0$. Since $\lambda_1 \geq 0$, $\tau^2 \leq \frac{(1-\frac{r}{2})(1-a)\omega}{(1-r)}$ implies that $\frac{\partial E[W(a, \theta)|\theta]}{\partial \sigma^2} \geq 0$, which concludes the proof. \square

5.7.15 Proof of Proposition 5.10

Proposition 5.10. For a fixed level of variance in the private signal, increasing the variance in the common shock decreases welfare.

Proof.

$$\frac{\partial E[W(a, \theta)|\theta]}{\partial \gamma^2} \Big|_{d\sigma^2 = -d\nu^2} = \frac{\partial \lambda_1}{\partial \gamma} \left(-(1-a)(1-\frac{r}{2}) - 2\lambda_1\nu^2(1-\frac{r}{2})(1-b) \right) - \lambda_1^2(1-\frac{r}{2})\frac{\partial \nu^2}{\partial \gamma}$$

But

$$\frac{\partial \lambda_1}{\partial \gamma} = \frac{(1-r)\beta\frac{r}{2}}{(\sigma^2 + b\nu^2 + \tau^2)} \geq 0$$

and

$$\frac{\partial \nu^2}{\partial \gamma} \geq 0$$

so we conclude that

$$\frac{\partial E[W(a, \theta)|\theta]}{\partial \gamma^2} \Big|_{d\sigma^2 = -d\nu^2} \leq 0$$

\square

5.7.16 Proof of Proposition 5.11

Proposition 5.11. Expected welfare is decreasing in the size of the bias of the informational sources. Furthermore, Expected welfare is decreasing in bias and correlational neglect.

Proof. First, it is easy to see that

$$\begin{aligned}\frac{\partial E[W(a, \theta)|\theta]}{\partial \omega} &= -\lambda_1(1-a)\left(1 - \frac{r}{2}\right) \leq 0 \\ \frac{\partial E[W(a, \theta)|\theta]}{\partial a} &= -\lambda_1\omega\left(1 - \frac{r}{2}\right) \geq 0\end{aligned}$$

So expected welfare is decreasing in the size of the bias of the information sources, and decreasing in bias neglect (as it is increasing in the acknowledgement of the neglect a).

Then I take the derivative with respect to b .

$$\frac{\partial E[W(a, \theta)|\theta]}{\partial b} = \frac{\partial \lambda_1}{\partial b} \left(-(1-a)\left(1 - \frac{r}{2}\right) - 2\lambda_1\nu^2\left(1 - \frac{r}{2}\right)(1-b) \right)$$

But

$$\frac{\partial \lambda_1}{\partial b} = \frac{(1-r)\frac{\partial \beta}{\partial b}}{1 - \frac{r}{2}(2\beta + \gamma)} + \frac{(1-r)\beta\frac{r}{2}\left(2\frac{\partial \beta}{\partial b} + \frac{\partial \gamma}{\partial b}\right)}{\left(1 - \frac{r}{2}(2\beta + \gamma)\right)^2}$$

and

$$\begin{aligned}\frac{\partial \beta}{\partial b} &= -\frac{\tau^2\nu^2}{(\sigma^2 + b\nu^2 + \tau^2)^2} = -\frac{\beta\gamma}{b} \\ \frac{\partial \gamma}{\partial \sigma^2} &= -\frac{\nu^2(\sigma^2 + \tau^2)}{(\sigma^2 + b\nu^2 + \tau^2)^2} = \frac{\gamma(1-\gamma)}{b}\end{aligned}$$

So we obtain

$$\begin{aligned}\frac{\partial \lambda_1}{\partial b} &= -\frac{\lambda_1\gamma}{b} + \frac{(1-r)\beta\frac{r}{2}\frac{\gamma}{b}(1-\gamma-2\beta)}{\left(1 - \frac{r}{2}(2\beta + \gamma)\right)^2} \\ &= -\frac{\lambda_1\gamma}{b} \left(1 - \frac{\frac{r}{2}(1-\gamma-2\beta)}{1 - \frac{r}{2}(2\beta + \gamma)} \right) \\ &= -\frac{\lambda_1\gamma}{b} \frac{\left(1 - \frac{r}{2}\right)}{1 - \frac{r}{2}(2\beta + \gamma)} \leq 0\end{aligned}$$

So we conclude that $\frac{\partial E[W(a, \theta)|\theta]}{\partial b} \geq 0$. □

Bibliography

- [1] Ricardo Alonso and Odilon Câmara. Persuading voters. *Available at SSRN 2688969*, 2015.
- [2] Ricardo Alonso and Odilon Camara. Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory*, 165:672–706, 2016.
- [3] George-Marios Angeletos and Alessandro Pavan. Efficient use of information and social value of information. *Econometrica*, 75(4):1103–1142, 2007.
- [4] Robert Aumann and Michael Maschler. Repeated games of incomplete information: a survey of recent results. in Report of the U.S. Arms Control and Disarmament Agency ST-116, Washington, D.C., Chapter III, pp.287-403, 1967.
- [5] Robert Aumann and Michael Maschler. Repeated games of incomplete information: the zero sum extensive case. in Report of the U.S. Arms Control and Disarmament Agency ST-116, Washington, D.C., Chapter III, pp.37-116, 1968.
- [6] Robert J Aumann. Mixed and behavior strategies in infinite extensive games. Technical report, 1964.
- [7] Robert J Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica: Journal of the Econometric Society*, pages 1–18, 1987.
- [8] Robert J Aumann, Y Katznelson, Roy Radner, Robert W Rosenthal, and B Weiss. Approximate purification of mixed strategies. *Mathematics of Operations Research*, 8(3):327–341, 1983.
- [9] Robert J Aumann, Michael Maschler, and Richard E Stearns. *Repeated games with incomplete information*. MIT press, 1995.
- [10] Robert J Aumann et al. Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics*, 1(1):67–96, 1974.
- [11] David Austen-Smith and John R Wright. Competitive lobbying for a legislator’s vote. *Social Choice and Welfare*, 9(3):229–257, 1992.
- [12] Erik J Balder. Comments on purification in continuum games. *International Journal of Game Theory*, 37(1):73–92, 2008.
- [13] Bruno Bassan, Olivier Gossner, Marco Scarsini, and Shmuel Zamir. Positive value of information in games. *International Journal of Game Theory*, 32(1):17–31, 2003.
- [14] Roland Bénabou. Joseph schumpeter lecture. *Journal of the European Economic Association*, 6 (2-3):321–352, 2008.

- [15] Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.
- [16] Dirk Bergemann and Stephen Morris. Information design, bayesian persuasion, and bayes correlated equilibrium. *American Economic Review*, 106(5):586–91, 2016.
- [17] Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.
- [18] Dirk Bergemann, Tibor Heumann, and Stephen Morris. Information and volatility. *Journal of Economic Theory*, 158:427–465, 2015.
- [19] Alexander W Bloedel and Ilya R Segal. Persuasion with rational inattention. *Available at SSRN 3164033*, 2018.
- [20] Isabelle Brocas and Juan D Carrillo. Influence through ignorance. *The RAND Journal of Economics*, 38(4):931–947, 2007.
- [21] Andrew Caplin and Mark Dean. Behavioral implications of rational inattention with shannon entropy. Technical report, National Bureau of Economic Research, 2013.
- [22] Andrew Caplin and Mark Dean. Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203, 2015.
- [23] Andrew Caplin and Daniel Martin. A testable theory of imperfect perception. *The Economic Journal*, 125(582):184–202, 2014.
- [24] Andrew Caplin, Mark Dean, and John Leahy. Rationally inattentive behavior: Characterizing and generalizing shannon entropy. Technical report, National Bureau of Economic Research, 2017.
- [25] Christopher Cotton. Should we tax or cap political contributions? a lobbying model with policy favors and access. *Journal of Public Economics*, 93(7-8):831–842, 2009.
- [26] Christopher S Cotton and Arnaud Dellis. Informational lobbying and agenda distortion. *The Journal of Law, Economics, and Organization*, 32(4):762–793, 2016.
- [27] Brendan Daley and Erik Snowberg. Even if it is not bribery: the case for campaign finance reform. *The Journal of Law, Economics, & Organization*, 27(2):324–349, 2011.
- [28] Peter M DeMarzo, Dimitri Vayanos, and Jeffrey Zwiebel. Persuasion bias, social influence, and unidimensional opinions. *The Quarterly journal of economics*, 118(3):909–968, 2003.
- [29] Tommaso Denti. Games with unrestricted information acquisition. Technical report, Technical report. Mimeo, 2015.
- [30] Joan Esteban and Debraj Ray. Inequality, lobbying, and resource allocation. *American Economic Review*, 96(1):257–279, 2006.
- [31] Matthew Gentzkow and Emir Kamenica. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- [32] Matthew Gentzkow and Emir Kamenica. Costly persuasion. *The American Economic Review*, 104(5):457–462, 2014.

- [33] Matthew Gentzkow and Emir Kamenica. Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, 104:411–429, 2017.
- [34] Matthew Gentzkow and Jesse M Shapiro. Media bias and reputation. *Journal of political Economy*, 114(2):280–316, 2006.
- [35] Matthew Gentzkow, Michael B Wong, and Allen T Zhang. Ideological bias and trust in information sources. *Unpublished manuscript*, 2018.
- [36] Edward L Glaeser and Cass R Sunstein. Extremism and social learning. *Journal of Legal Analysis*, 1(1):263–324, 2009.
- [37] Edward L Glaeser and Bryce A Ward. Myths and realities of american political geography. *Journal of Economic Perspectives*, 20(2):119–144, 2006.
- [38] Olivier Gossner. Comparison of information structures. *Games and Economic Behavior*, 30(1): 44–63, 2000.
- [39] Michael Greinecker and Konrad Podczeck. Purification and roulette wheels. *Economic Theory*, 58(2):255–272, 2015.
- [40] John C Harsanyi. Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International journal of game theory*, 2(1):1–23, 1973.
- [41] Jonas Hedlund. Bayesian persuasion by a privately informed sender. *Journal of Economic Theory*, 167:229–268, 2017.
- [42] Jennifer L Hochschild. *What's fair?: American beliefs about distributive justice*. Harvard University Press, 1981.
- [43] Jennifer L Hochschild. *Facing up to the American dream: Race, class, and the soul of the nation*, volume 51. Princeton University Press, 1996.
- [44] John T Jost, Brenda Major, et al. *The psychology of legitimacy: Emerging perspectives on ideology, justice, and intergroup relations*. Cambridge University Press, 2001.
- [45] Shizuo Kakutani et al. A generalization of brouwer's fixed point theorem. *Duke mathematical journal*, 8(3):457–459, 1941.
- [46] M Ali Khan and Kali P Rath. On games with incomplete information and the dvoretzky–wald–wolfowitz theorem with countable partitions. *Journal of Mathematical Economics*, 45(12):830–837, 2009.
- [47] M Ali Khan, Kali P Rath, and Yeneng Sun. On a private information game without pure strategy equilibria. *Journal of Mathematical Economics*, 31(3):341–359, 1999.
- [48] Frédéric Koessler, Marie Laclau, and Tristan Tomala. Interactive Information Design. Pse working papers, HAL, May 2018.
- [49] Anton Kolotilin, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. Persuasion of a privately informed receiver. *Econometrica*, 85(6):1949–1964, 2017.
- [50] Timur Kuran. *Private truths, public lies*. Harvard University Press, 1997.

- [51] Robert E Lane. The fear of equality. *The American Political Science Review*, 53(1):35–51, 1959.
- [52] Maël Le Treust and Tristan Tomala. Persuasion with limited communication capacity. *Journal of Economic Theory*, 184:104940, 2019.
- [53] EL Lehmann. Comparing location experiments. *The Annals of Statistics*, pages 521–533, 1988.
- [54] Wolfgang Leininger. A generalization of the ‘maximum theorem’. *Economics Letters*, 15(3-4): 309–313, 1984.
- [55] Gilat Levy and Ronny Razin. Correlation neglect, voting behavior, and information aggregation. *American Economic Review*, 105(4):1634–45, 2015.
- [56] Elliot Lipnowski and Laurent Mathevet. Disclosure to a psychological audience. *American Economic Journal: Microeconomics*, 10(4):67–93, 2018.
- [57] Peter Loeb, Yeneng Sun, et al. Purification of measure-valued maps. *Illinois Journal of Mathematics*, 50(1-4):747–762, 2006.
- [58] Daniel Martin. Rational inattention in games: experimental evidence. 2016.
- [59] Daniel Martin. Strategic pricing with rational inattention to quality. *Games and Economic Behavior*, 104:131–145, 2017.
- [60] Gregory J Martin and Ali Yurukoglu. Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9):2565–99, 2017.
- [61] Filip Mat vejka. Rationally inattentive seller: Sales and discrete pricing. *The Review of Economic Studies*, 83(3):1125–1155, 2016.
- [62] Filip Matejka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *The American Economic Review*, 105(1):272–298, 2014.
- [63] Laurent Mathevet, Jacopo Perego, and Ina Taneva. On information design in games. *Journal of Political Economy*, 128(4):1370–1404, 2020.
- [64] Ludmila Matyskova. Bayesian persuasion with costly information acquisition. *CERGE-EI Working Paper Series*, (614), 2018.
- [65] Nolan McCarty, Keith T Poole, and Howard Rosenthal. *Polarized America: The dance of ideology and unequal riches*. mit Press, 2016.
- [66] Paul R Milgrom and Robert J Weber. Distributional strategies for games with incomplete information. *Mathematics of operations research*, 10(4):619–632, 1985.
- [67] Stephen Morris and Hyun Song Shin. Social value of public information. *The American Economic Review*, pages 1521–34, 2002.
- [68] Abraham Neyman. The positive value of information. *Games and Economic Behavior*, 3(3):350–355, 1991.
- [69] Pietro Ortoleva and Erik Snowberg. Overconfidence in political behavior. *American Economic Review*, 105(2):504–35, 2015.

- [70] Konrad Podczeck. On purification of measure-valued maps. *Economic Theory*, 38(2):399–418, 2009.
- [71] Markus Prior. Media and political polarization. *Annual Review of Political Science*, 16:101–127, 2013.
- [72] Roy Radner and Robert W Rosenthal. Private information and pure-strategy equilibria. *Mathematics of Operations Research*, 7(3):401–409, 1982.
- [73] Eric Rasmusen. Lobbying when the decisionmaker can acquire independent information. *Public Choice*, 77(4):899–913, 1993.
- [74] Doron Ravid. Bargaining with rational inattention. 2017.
- [75] Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of political Economy*, 118(5): 949–987, 2010.
- [76] Ariel Rubinstein. Comments on the interpretation of game theory. *Econometrica: Journal of the Econometric Society*, pages 909–924, 1991.
- [77] Christopher A Sims. Implications of rational inattention. *Journal of monetary Economics*, 50(3): 665–690, 2003.
- [78] Christopher A Sims. Rational inattention: Beyond the linear-quadratic case. *The American economic review*, 96(2):158–163, 2006.
- [79] Jakub Steiner, Colin Stewart, and Filip Mat vejka. Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica*, 85(2):521–553, 2017.
- [80] Ina Taneva. Information design. *American Economic Journal: Microeconomics*, 11(4):151–85, 2019.
- [81] Eagleton Terry. *Ideology: an introduction*. Verso, London, 1991.
- [82] Jianwei Wang and Yongchao Zhang. Purification, saturation and the exact law of large numbers. *Economic Theory*, 50(3):527–545, 2012.
- [83] Ming Yang. Coordination with flexible information acquisition. *Journal of Economic Theory*, 158: 721–738, 2015.

Titre : Essais Dans les Jeux Bayésiens

Mots clés : Jeux Bayésiens, Acquisition d'information coûteuse, Inattention Rationnelle, Structures d'information endogènes, Persuasion Bayésienne.

Résumé : Cette thèse de doctorat propose une série d'essais sur les jeux bayésiens. Il concentre son attention sur le rôle de l'information dans les résultats de ces jeux, depuis les stratégies et l'équilibre jusqu'aux implications pour le bien-être. D'une manière générale, les jeux bayésiens sont des situations stratégiques dans lesquelles il existe des informations incomplètes par rapport à un état pertinent lors de la détermination des paiements. Les joueurs ont une croyance a priori sur le paramètre inconnu, et ils reçoivent des informations sur l'état avant de choisir leurs actions. Les joueurs sont appelés Bayesian car ils mettent à jour leurs informations en utilisant la règle de Bayes. Traditionnellement, les informations reçues par les joueurs sont supposées exogènes et correctement comprises.

Cette thèse de doctorat explore les implications et les conséquences de l'abandon de ces hypothèses.

Les trois principaux chapitres de cette thèse abordent deux questions différentes liées à l'information. D'une part, on ne suppose plus que les joueurs reçoivent des informations données de manière exogène sur l'état inconnu, mais qu'ils doivent les collecter tout en engageant un coût. Cette hypothèse est à l'origine des contributions et des résultats des chapitres 3 et 4. Par contre, au chapitre 5, il n'est plus supposé que les informations distribuées entre les acteurs sont bien comprises. Autrement dit, je considère que les joueurs reçoivent des biais et des informations corrélées et qu'ils n'en sont pas pleinement conscients.

Title: Essays in Bayesian Games

Keywords: Bayesian Games, Costly Information Acquisition, Rational Inattention, Endogenous Information structures, Bayesian Persuasion

Abstract: This PhD dissertation provides a series of essays in Bayesian games. It centers its attention on the role of information in the outcomes of such games, from strategies and equilibrium to welfare implications. Broadly speaking, Bayesian games are strategic situations in which there is incomplete information regarding a payoff relevant state. Players have a prior belief about the unknown parameter, and they receive information about the state before choosing their actions. Players are called *Bayesian* because they update their information using Bayes rule. Traditionally, the information players receive is assumed to be exogenous and correctly understood. This PhD dissertation explores

the implications and consequences of giving up these assumptions.

The three core chapters of this thesis tackle two different issues related to information. On the one hand, it is no longer assumed that players receive exogenously given information about the unknown state, but that they have to gather it while incurring in a cost. This assumption drives the contributions and results of chapters 4 and 3. On the other hand, in chapter 5 it is no longer assumed that information distributed among the players is well understood. That is, I consider that players receive bias and correlated information and that they are not fully aware of it.