CERGE
Center for Economics Research and Graduate Education
Charles University

Essays on Costly Information Acquisition in Economics

Vladimír Novák

Dissertation

Prague, July 2020

Vladimír Novák

Essays on Costly Information Acquisition in Economics

Dissertation

Prague, July 2020

Dissertation Committee

Filip Matějka (CERGE-EI; chair)

Jan Zápal (CERGE-EI)

Gérard Roland (University of California, Berkeley)

Michal Kejak (CERGE-EI)

Referees

Mark Dean (Columbia University)

Giorgia Romagnoli (University of Amsterdam)

Table of Contents

Abstract **vii**

Abstrakt **ix**

Acknowledgments **xi**

Introduction **1**

1 A Note on Optimal Experimentation Under Risk Aversion **5**
1.1 Introduction 5
1.2 A bandit model with non-linear utility 6
1.3 The effects of risk aversion 9
1.4 Discussion 13
1.A Appendix A: $u_2(x)/u_1(x)$ is strictly decreasing in x 16
1.B Appendix B: Generalizing Proposition 1 17
1.C Appendix C: A two-period model 19

2 The Status Quo and Belief Polarization of Inattentive Agents: Theory and Experiment **21**
2.1 Introduction 21
2.2 Example 25
2.3 The model 27
2.3.1 Description of the setup 28
2.3.2 Agent's problem 28
2.3.3 Description of beliefs evolution 31
2.3.4 Updating in the opposite direction from the realized value 32
2.3.5 Belief polarization 35
2.3.6 Beliefs' convergence and divergence of beliefs updated in the same direction 36
2.4 Over-optimism and Polarization: Intuition and Implications 37
2.5 Comparative statics 42
2.6 Experimental design 45
2.6.1 Overview of the experimental design 46

2.6.2 Task 1 - Colorblind advisor game . . . 46
2.6.3 Task 2 - Imprecise advisor game . . . 47
2.6.4 Task 3 - Card color prediction game . . . 49
2.6.5 Task 4 - Ball color prediction game . . . 49
2.6.6 Questionnaire . . . 50
2.6.7 Procedure . . . 50
2.6.8 Subject understanding . . . 50
2.7 Experimental investigation . . . 51
2.7.1 Formal setup . . . 51
2.7.2 Bridging theory and experiment . . . 53
2.7.3 Hypothesis . . . 54
2.8 Experimental results . . . 55
2.8.1 Optimal response to the change in the status quo . . . 56
2.8.2 State pooling and preference for certainty . . . 58
2.8.3 Action selection . . . 62
2.8.4 Belief elicitation . . . 64
2.8.5 Observed belief polarization . . . 65
2.8.6 Compression effect in the willingness to pay . . . 67
2.9 Conclusion . . . 70
2.A Appendix A: Derivation of formula 2.5 . . . 73
2.B Appendix B: Proof of proposition 1 . . . 74
2.C Appendix C: Proof of proposition 2 . . . 76

3 Estimating Models with Rationally Inattentive Agents 79
3.1 Introduction . . . 79
3.2 The Model . . . 82
3.2.1 Assumptions . . . 83
3.2.2 The representative household . . . 83
3.2.3 The financial intermediary . . . 84
3.2.4 The monetary authority . . . 85
3.2.5 Firms . . . 85
3.2.6 Price-setting equation and law of motions . . . 86
3.3 Rational inattention model . . . 88
3.3.1 Quantifying information flow . . . 89
3.3.2 The objective . . . 91
3.3.3 Rational inattention Kalman filter formulation . . . 92
3.3.4 Optimal signal selection . . . 94
3.4 Imperfect common knowledge model . . . 96
3.5 Calvo pricing model . . . 98
3.6 Empirical Analysis . . . 98
3.6.1 Data . . . 99
3.6.2 VAR analysis . . . 99
3.6.3 Prior and posterior distributions . . . 99
3.6.4 Results comparison . . . 102
3.7 Concluding Remarks . . . 105
3.A Appendix A: Derivation of the log-linearized firm's price-setting equation 106

3.B Appendix B: Analogous objective of rational inattention model 110
3.C Appendix C: Profit-maximizing price for the rational inattention model . 112
3.D Appendix D: Solving the ICKM 113

Bibliography **117**

Abstract

In the first chapter, we solve the two-armed bandit problem when decision-makers are risk-averse. We show – counterintuitively – that a more risk-averse decision-maker might be more willing to take risky actions. This finding relates to the fact that pulling the risky arm in bandit models produces information on the environment – thereby reducing the risk that a decision-maker will face in the future. Thus, we suggest there is reason for caution when inferring risk preferences from observed actions: in a bandit setup, observing a greater appetite for risky actions can be indicative of more risk aversion, not less.

In the second chapter, we characterize when it is rational to acquire information leading to belief polarization, in situations that involve a choice between the implementation of a new policy with an uncertain outcome and the preservation of the status quo. Specifically, we model the agent to be rationally inattentive: any information about the new policy can be acquired before the choice is made, but doing so is costly. We show how the choice of information, and thus the belief formation, depends on the agent-specific value of the status quo. Importantly, beliefs can then, in expectations, update away from the realized truth. This is due to endogenous information acquisition because the agent chooses only to learn whether the uncertain payoff is higher or lower than the payoff of the status quo. Consequently, two agents with the same prior beliefs about a new policy might become polarized if they differ in the valuations of the status quo. We show that the lower cost of information leads to more severe polarization. We conduct a novel experiment to test and confirm our predicted information acquisition strategy and its dependence on the value of the status quo. In our setting with multiple states, we also replicate the well-known preference for certainty and verify the occurrence of belief polarization.

In the third chapter, we present a likelihood evaluation of a DSGE model with price-setting firms that select properties of their signals subject to a limited attention constraint. We compare the performance of a rational inattention DSGE model (RIM), with an imperfect common knowledge model (ICKM) and a model with price stickiness à la Calvo. We demonstrate that the rational inattention model matches the data better than the Calvo model and reproduces the persistence more easily than the ICKM model. This result occurs because (i) RI firms pay attention to a higher number of lags of fundamentals than is assumed in the ICKM models, and (ii) the full information method selects the different degree of strategic complementarity in various models.

Abstrakt

V první kapitole zkoumáme model optimálního experimentování s hráči, kteří mají averzi k riziku. Ukazujeme, že, navzdory intuici, hráč s větší averzí k riziku může optimálně volit riskantnější experiment. K tomuto efektu dochází, protože riskantnější experiment přináší informace, které snižují budoucí riziko. Větší apetit k rizikovějšímu experimentování tedy může indikovat větší, nikoliv nižší, averzi vůči riziku. Empirické a experimentální studie, které by ignorovaly tento efekt, mohou produkovat vychýlené odhady míry averze vůči riziku.

Ve druhé kapitole charakterizujeme, kdy je racionální získávat informace vedoucí k názorové polarizaci, v situacích, které zahrnují volbu mezi implementací nové politiky s nejistým výsledkem a zachováním současného stavu. Konkrétně modelujeme agenta, který je racionálně nepozorný: veškeré informace o nové politice lze získat ještě před provedením výběru, ale je to nákladné. Ukazujeme, jak výběr informací, a tím i tvorba názorů, závisí na jeho subjektivní hodnotě statusu quo. Důležité je, že po získaní informace se aktualizovaný názor může vzdalovat od skutečnosti. To je způsobeno endogenní volbou jakou informaci získat, protože agent se rozhodne pouze zjistit, zda nová politika povede k lepšímu výsledku než současný stav. V důsledku toho se dva agenti se stejným názorem na novou politikou po obdržení informace mohou polarizovat, pokud se liší ve spokojenosti se současným stavem. Ukazujeme, že nižší náklady na informace vedou k významnější polarizaci. Provádíme také nový laboratorní experiment, který testuje a potvrzuje naši předpokládanou strategii získávání informací a její závislost na hodnotě současného stavu. V našem prostředí s více stavy světa také replikujeme známé preference pro jistotu a ověřujeme výskyt polarizace.

V třetí kapitole představujeme odhadování DSGE modelu s firmami určujícími ceny, které se vybírají jaké informace chtějí získat, ale jsou omezené v kapacitě zpracování těchto informací. Porovnáváme přesnost predikcí DSGE modelu s racionálně nepozornými firmami (RIM), s modelem nedokonalých sdílených znalostí (ICKM) a modelem s cenovou přizpůsobivostí à la Calvo. Ukazujeme, ze RIM odpovídá datům lépe než Calvo model a snadněji replikuje zpoždění a perzistenci než ICKM model. K tomuto výsledku dochází, protože (i) RI firmy věnují pozornost většímu počtu historických hodnot o fundamentech, než se předpokládá v ICKM, (ii) estimace vybere jiný stupeň strategické komplementarity v různých modelech.

Acknowledgments

I would like to express my special thanks to my supervisor Filip Matějka for all the support, help and invaluable advice for research and life offered me throughout my doctoral studies. I am grateful to other members of my dissertation committee Jan Zápal, Gérard Roland and Michal Kejak, who were supportive and helpful throughout. I would also like to thank Jakub Steiner for our fascinating discussions, Ole Jann and Ctirad Slavík for support during the job market.

I am indebted to many other inspiring people, especially Mark Dean who provided me with excellent supervision and support when it was most needed during my stay at Columbia University and to a fantastic group of people associated with the Cognition and Decision Lab who warmly welcomed me among themselves. I am very thankful to my co-authors, Andrei Matveenko, Silvio Ravaioli, Tim Willems and Godfrey Keller. I hope that our collaboration will continue in the future. I also want to thank my fellow students Michal Šoltés, Klára Svitáková, Sergei Mikhalishev and Ľudmila Matýsková, who were always there for me, CERGE-EI SAO for help with administrative duties, Andrea Downing and Paul Whitaker who carefully and patiently edited all my writings over the years, and the whole CERGE-EI community. For their love, and helping me to keep things in perspective, I warmly thank Zuzka, my brother Marek and my close friends.

Financial support received from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreements No. 678081), support by the Charles University in Prague (GA UK projects No. 197216 and No. 616219), and by the H2020-MSCA-RISE project GEMCLIME-2020 GA. (no 681228) are gratefully acknowledged.

All errors remaining in this text are my responsibility.

Prague, Czech Republic Vladimír
July 2020

Introduction

In the information-rich world, attention is a scarce resource. The unifying theme of all three chapters of this dissertation centers on endogenous information acquisition problems and their implications. In the first chapter, we investigate the impact of risk-aversion on the sequential allocation of an 'informational' resource in the two-armed bandit problem. The other two chapters apply the theory of rational inattention (RI), which is a disciplined model of how cognitively-limited people simplify and summarize available information. The key assumption is that all information is available, but decision-makers cannot pay full attention to it, though they can choose to pay more attention to more important things. In the second chapter, we adopt RI theory to study, both theoretically and experimentally, how the choice of information, and thus the belief formation depends on the status quo and can lead to belief polarization of rational people. The final chapter takes steps to enable the Bayesian estimation of the rational inattention DSGE models and compares it with the imperfect common knowledge model and Calvo model.

In the first chapter: "A Note on Optimal Experimentation under Risk Aversion" (joint work with Tim Willems and Godfrey Keller), which was published in the Journal of Economic Theory. We extend the standard exponential bandit model by allowing for risk aversion on behalf of the decision maker and study how a rational, risk-averse decision maker solves the two-armed bandit problem of having to sequentially allocate its 'informational' resource between a safe alternative that yields a known reward, and a risky one that generates an unknown payoff. In doing so, we uncover the previously overlooked result that a more risk-averse decision maker might be more willing to pull the risky arm than a less risk-averse colleague. The reason for this counterintuitive result relates to the notion that risk in bandit models can be reduced through experimentation with the risky arm. Our findings imply that it is not obvious to infer risk preferences from observed actions: when there is scope for experimentation, observing a greater appetite for risky actions might actually be indicative of more risk aversion, not less.

In the second chapter: "The Status Quo and Beliefs Polarization of Inattentive Agents: Theory and Experiment" (co-authored with Andrei Matveenko and Silvio Ravaioli), we study a binary choice problem with uncertainty and show how rational people can become more polarized in expectations as they learn more. In particular, the decision makers' choice is between preservation of the status quo and a new policy with multiple possible outcomes. Consider the 2016 Brexit referendum as an illustrative example - the consequences of choosing to leave the European Union are uncertain, whereas the status quo is more clear. We analyze this question both theoretically and with a novel laboratory experiment. We show how the choice of information, and thus the belief formation, depends on the value of the status quo. Importantly, beliefs can then, in expectations, update away from the realized truth. This is due to the acquisition of endogenous information, because the agent chooses only to learn whether the uncertain payoff is higher or lower than the payoff of the status quo. Consequently, two agents with the same prior beliefs about the new policy might become polarized if they differ in the value of the status quo. We show that the lower cost of information makes the polarization more severe. This paper also provides a disciplined model proposing how preference for the skewed information might depend on the value of the status quo, and thus provide an important channel that is absent from the research studying whether people prefer negatively or positively skewed information. We experimentally test and confirm our theoretically predicted information acquisition strategy and its dependence on the value of the status quo. We show that the probability of choosing an advisor increases with the instrumental value of the corresponding information structure. Importantly, we verify the consequent belief polarization in expectations. In our setting with multiple states, we also replicate the well-known preference for certainty and verify the occurrence of belief polarization.

In the final chapter: "Estimating Models with Rationally Inattentive Agents", we estimate a DSGE model with rationally inattentive price-setting firms. We compare the performance of a rational inattention DSGE model (RIM), with an imperfect common knowledge model (ICKM) and a model with price stickiness à la Calvo. Importantly, this comparison helps us to identify what parameter values would be selected for the rational inattention model and how well it matches the data in contrast to other models. It also sheds light on restrictiveness and implications of the ICKM assumed signal form versus the optimally selected signal, given the information capacity constraint. In contrast with previous studies, we do not assume any particular exogenously given signal form or signals' independence, but by modelling firms as rationally inattentive, we allow firms to choose the optimal signals about the state variables optimally under the limited attention constraint.The substantive contribution of this study is in showing that the RI model matches the data better than the Calvo model, in particular by reproducing the persistence in the data more easily. Furthermore, the RI model reproduces the long-term mean-reversion better than the ICKM model, whereas the ICKM model seems to perform better in matching the short-run momentum of the hump-shaped response of output and inflation to monetary disturbances. Our likelihood analysis also emphasizes the role of the strategic complementarity in price-setting in the specifications of these models.

Chapter 1

A Note on Optimal Experimentation Under Risk Aversion

Co-authored with Godfrey Keller (University of Oxford) and Tim Willems (International Monetary Fund)[1].

1.1 Introduction

This paper analyzes how a rational, risk-averse decision maker solves the two-armed bandit problem of having to choose between a safe alternative that yields a known reward, and a risky one that generates an unknown payoff.

At first sight, it seems intuitive that decision makers who are more risk averse will be less willing to take the risky action. Indeed, an earlier paper (Chancelier, De Lara, and De Palma 2009) arrives at such a conclusion. However, we show that there exists a previously overlooked part of the parameter space, where this result is overturned.

Our model is based upon the exponential bandit model of Keller, Rady, and Cripps (2005). Following Roberts and Weitzman (1981) and Bolton and Harris (1999), who in turn built upon the seminal work of Rothschild (1974b, Rothschild (1974a), it uses a continuous-time framework. We extend the standard model by allowing for risk aversion

[1] The views expressed in this paper are those of the authors and should not be attributed to the International Monetary Fund, its Executive Board, or its management.

on behalf of the decision maker. In doing so, we uncover the previously overlooked result that a more risk-averse decision maker might be *more* willing to pull the risky arm than a less risk-averse colleague.

The reason for this counterintuitive result relates to the notion that risk in bandit models can be reduced through experimentation with the risky arm. It is most likely to arise in settings where information arrives at a high frequency, which makes our finding of specific relevance to the machine learning literature (where reinforcement learning algorithms have the bandit problem at their core; see Sutton and Barto (1998)).

It is furthermore important to understand how risk aversion of decision makers affects the decisions they make. Willingness to take risks has been linked to the success of entrepreneurs (Cantillon 1775; Knight 1921; Kihlstrom and Laffont 1979; Herranz, Krasa, and Villamil 2015), while it has also been studied in a principal-agent setup – for example analyzing decision making by CEOs (Bandiera et al. 2011) and politicians (Lilienfeld 2012). A common narrative that can be found in this literature is that more risk-averse individuals can be expected to take less-risky actions. The point of this paper is to show that the introduction of learning and experimentation can overturn this wisdom: appointing a more risk-averse decision maker is no guarantee for the implementation of less risky actions.

Finally, our findings imply that it is not obvious to infer risk preferences from observed actions: when there is scope for experimentation, observing a greater appetite for risky actions might actually be indicative of *more* risk aversion, not less. Studies which do not take this into account may produce biased estimates.

1.2 A bandit model with non-linear utility

In this section, we employ the exponential bandit model of Keller, Rady, and Cripps (2005) but extend it by allowing for a decision maker (henceforth 'DM') that need not be risk-neutral. For ease of exposition, we focus on the one-player case.[2]

[2]The results in this section and the next carry over straightforwardly to an N-agent cooperative setup.

Time $t \in [0, \infty)$ is continuous and the discount rate is $r > 0$. The player is facing a two-armed bandit problem, and at time t can allocate an amount $k_t \in \{0, 1\}$ of his 'informational' resource to the risky arm R, and thus $1 - k_t$ to the safe arm S.

The safe arm provides lump-sum payoffs of $s > 0$ (with the value of s fixed and known to the player, hence why this arm is called 'safe') according to a Poisson process with parameter 1 (which is also known). So if the player uses the safe arm over an interval $[t, t + dt)$, he receives $s\,dt$ in expectation.

The source of risk in the other arm lies in the fact that its type θ, and hence the size of its payoff, is unknown to the agent at $t = 0$. He knows that the arm is either 'good' ($\theta = 1$) or 'bad' ($\theta = 0$). At time t, the player holds a belief p_t that the risky arm is good. The DM's learning process on the risky arm's type is obstructed by the presence of noise in the associated payoff stream. When the arm is good, it yields lump-sum payoffs of h according to a Poisson process with parameter $\lambda > 0$ (both h and λ are fixed and known by the player). When it is bad, it never pays off. Consequently, if the player uses the risky arm over an interval $[t, t + dt)$, he receives $p_t \lambda h\, dt$ in expectation (where the expectation is taken over both the unknown state of the world θ and the probabilistic arrival of the lump-sums).

The player evaluates the lump-sums using a utility function u, and so the expected increase in his utility is $[(1 - k_t)u(s) + k_t p_t \lambda u(h)]\, dt$. We will assume that $u(0) = 0$ for the following reason. Consider a stream of zero payoffs that arrive according to a Poisson process with parameter ν. The total expected discounted utility from this stream, expressed in per-period terms is $\mathbb{E}\left[\int_0^\infty re^{-rt}\,\nu u(0)\, dt\right] = \nu u(0)$, and it is natural to require that this be independent of the rate ν – it should not matter at what frequency nothing is paid out; this translates into a requirement that $u(0) = 0$.[3] Also, to make the problem meaningful, we require that the player strictly prefers R, if it is good, to S, and strictly prefers S to R, if it is bad. Consequently, we assume that:

ASSUMPTION 1. $0 = u(0) < u(s) < \lambda u(h)$

As more information arrives over time, the belief p_t is revised according to Bayes' rule.

[3] We thank a referee for this justification of $u(0)$ being zero, and for noting that this requirement implies that two such utility functions represent the same preferences if and only if one is a monotone linear transformation of the other.

When the DM plays the risky arm but no lump-sum arrives, his belief that the risky arm is good is revised downward:

$$dp_t = -\lambda p_t(1-p_t)\,dt. \tag{1.1}$$

On the other hand, if a lump-sum h does arrive, the belief p_t jumps to 1. The objective of the DM is to choose $\{k_t\}_{t\geq 0}$ so as to maximize the total expected discounted utility, expressed in per-period terms:

$$\mathbb{E}\left[\int_0^{\infty} re^{-rt}\left[(1-k_t)u(s)+k_t p_t\lambda u(h)\right]\,dt\right],$$

where the expectation is over the processes $\{k_t\}$ and $\{p_t\}$,[4] and with beliefs being the state variable. The solution procedure is analogous to that in Keller, Rady, and Cripps (2005) – the only difference being the presence of the utility function. As in Keller, Rady, and Cripps (2005), the Principle of Optimality implies that the value function V satisfies:

$$V(p) = \max_{k\in\{0,1\}}\left\{r\left[(1-k)u(s)+kp\lambda u(h)\right]\,dt + e^{-r\,dt}\mathbb{E}\left[V(p+dp)|p,k\right]\right\}.$$

To eliminate the expectations operator, observe that with subjective probability $pk\lambda\,dt$ a lump-sum h arrives, revealing to the DM that the risky arm is of the good type. In that case, the value function jumps to $V(1)=\lambda u(h)$. With complementary probability $1-pk\lambda\,dt$, no lump-sum arrives; then, application of Bayes' rule (1.1) enables us to write $V(p+dp)\approx V(p)+V'(p)\,dp = V(p)-k\lambda p(1-p)V'(p)\,dt$. Combining this with $1-r\,dt$, the approximation to $e^{-r\,dt}$, leads to the following Bellman equation:

$$V(p) = \max_{k\in\{0,1\}}\left\{(1-k)u(s)+kp\lambda u(h)+kp\lambda\left[\lambda u(h)-V(p)-(1-p)V'(p)\right]/r\right\}.$$

As the maximand is linear in k, the DM would never optimally choose an interior allocation even were it allowed. If the DM chooses $k=0$, then $V(p)=u(s)$. If he chooses $k=1$, then V satisfies the following first-order ordinary differential equation:

$$\lambda p(1-p)V'(p)+(r+\lambda p)V(p) = (r+\lambda)\lambda u(h)p,$$

[4]The total expected discounted utility given in the main text is equivalent to:

$$\mathbb{E}\left[\int_0^{\infty} re^{-rt}\left[(1-k_t)u(s)\,dN_{1,t}+k_t p_t u(h)\,dN_{\lambda,t}\right]\right],$$

where $N_{\ell,t}$ is a standard Poisson process with intensity ℓ, since $N_{\ell,t}-\ell t$ is a martingale.

whose solution is given by:

$$V(p) = \lambda u(h)p + C(1-p)\left(\frac{1-p}{p}\right)^{r/\lambda},$$

where C is the constant of integration.

This solution has the exact same structure as that in Keller, Rady, and Cripps (2005). It therefore inherits the feature that there exists a cut-off belief p^* above which it is optimal for the DM to play the risky arm R, while playing the safe arm S becomes optimal when the DM's belief $p \leq p^*$. By imposing value matching ($V^*(p^*) = u(s)$) and smooth pasting ($(V^*)'(p^*) = 0$), we can derive the cut-off belief[5] as:

$$p^* = \frac{(r/\lambda)u(s)}{(r/\lambda + 1)\left[\lambda u(h) - u(s)\right] + (r/\lambda)u(s)}. \tag{1.2}$$

Now consider two DMs (indexed by i), with the utility function of DM_i being u_i. The difference between their cut-off beliefs $p_2^* - p_1^*$ satisfies:

$$\begin{aligned}\text{sgn}(p_2^* - p_1^*) &= \text{sgn}\left(u_2(s)\left[\lambda u_1(h) - u_1(s)\right] - u_1(s)\left[\lambda u_2(h) - u_2(s)\right]\right) \\ &= \text{sgn}\left(u_2(s)u_1(h) - u_1(s)u_2(h)\right) \\ &= \text{sgn}\left(\frac{u_2(s)}{u_1(s)} - \frac{u_2(h)}{u_1(h)}\right). \end{aligned} \tag{1.3}$$

Note that if DM_1 and DM_2 are both risk-neutral, then the right-hand side of the above equation is zero and hence $p_2^* = p_1^*$.

1.3 The effects of risk aversion

For the remainder of this article, we assume that the DM's utility function over payoffs u is increasing and concave, and recall that $u(0) = 0$. This assumption captures the notion of risk aversion in the sense that when our DM compares two streams of lump-sum payoffs, in each interval $[t, t + dt)$ he is facing a lottery over payoff increments, and when one of these lotteries second-order stochastically dominates the other, he prefers the less risky of the two.

[5] Note that this belief is invariant to a monotone linear transformation of u.

Consider one stream $\mathcal{P}_{h,\lambda}$ that delivers lump-sum payoffs h according to a Poisson process with parameter λ, and another one $\mathcal{P}_{c,1}$ that delivers lump-sum payoffs c according to a Poisson process with parameter 1, with $c = \lambda h$. In $[t, t + dt)$, the probability of no payoff from $\mathcal{P}_{h,\lambda}$ equals $1 - \lambda\, dt$, while the probability of no payoff from $\mathcal{P}_{c,1}$ equals $1 - dt$. Also note that the lottery with the larger probability of no payoff is a mean-preserving spread of the other. Consequently, the lottery with the smaller probability of no payoff second-order stochastically dominates the other, and $\mathcal{P}_{h,\lambda}$ is preferred to $\mathcal{P}_{c,1}$ iff $\lambda \geq 1$. In terms of the total expected discounted utility from the two streams, which are $\lambda u(h)$ and $u(c) = u(\lambda h)$, we note that $\lambda u(h) \geq u(\lambda h)$ iff $\lambda \geq 1$.

This ordering of 'lumpy' payoff streams manifests itself as follows: other things being equal, a DM prefers a stream of modest payoffs that arrive with a high expected frequency to a stream of larger payoffs that are expected to arrive infrequently. This favoring of 'less risky' payoff streams and of 'smaller payoffs at higher expected frequency' are simply two manifestations of the same preference. Consequently, we couch most of our discussion below in terms of higher (expected) frequency rather than in terms of second-order stochastic dominance or lower risk.

To analyze the effects of risk aversion, let DM_2 be more risk averse than DM_1. In particular, the more risk-averse DM_2 has an increasing, concave utility function u_2 which is a concave transformation of u_1, the utility function of the less risk averse DM_1.

In Appendix A, we show that $u_2(x)/u_1(x)$ is strictly decreasing in x; this, together with reference back to equation (1.3), leads to our main result:

Proposition 1. *The ordering of the cut-off beliefs for DM_1 and DM_2 is as follows: (a) when $h > s$, $p_2^* > p_1^*$; (b) when $h < s$, $p_2^* < p_1^*$; (c) finally, when $h = s$, $p_2^* = p_1^*$.*[6,7]

Part (a) of Proposition 1 implies that the more risk-averse DM needs a more optimistic belief on the quality of the risky arm to become willing to play R. In case (b) however,

[6] Parts (b) and (c) of the proposition require λ to be high enough so that Assumption 1 is not violated. If it were violated, the DM would never pull the risky arm (even if it was known to be of good quality).

[7] Following suggestions by a referee, Appendix B contains a generalization of this proposition by considering an infinitesimal increase in risk aversion, employing a Pratt (1964) representation.

the more risk-averse DM has the ***lower*** threshold p^* – implying that he will play R at more pessimistic beliefs relative to the less risk-averse DM.

To gain intuition for Proposition 1, start with part (c). When $h = s$, the safe arm gives rise to the exact same payoff as a good-quality risky arm (only at a different frequency given that $\lambda > 1$). As a result, $h/s = u_i(h)/u_i(s) = 1$ for $i = 1, 2$ and all payoff-related terms disappear from the cut-off formula (1.2). Both collapse to:

$$p_1^* = p_2^* = \frac{r/\lambda}{r + \lambda - 1} \tag{1.4}$$

From (1.4), one can see that pushing λ up (making the risky arm more attractive), lowers the cut-off belief (thus increasing the DM's willingness to try the risky arm). Crucially, however, when $h = s$, the cut-off belief falls at the same rate for all DMs irrespective of their degree of risk aversion (because $u_i(h)/u_i(s) = 1$ for $i = 1, 2$ and transformations of payoffs no longer affect the cut-off location), thereby keeping $p_1^* = p_2^*$.

This no longer holds true when we increase h slightly to $h' > s$. Again, we have made the risky arm more attractive, in response to which both p_1^* and p_2^* fall (see equation (1.2)). But since marginal utility of a more risk-averse DM decreases at a faster rate when the payoff rises, he gains fewer utils from the increase in h than his less risk-averse counterpart. As a result, the less risk-averse DM's cut-off p_1^* falls by more than the more risk-averse DM's cut-off p_2^* – putting us in case (a) of Proposition 1, where the less risk-averse DM is more willing to pull the risky arm.

Further understanding of the difference between parts (a) and (b) of Proposition 1 can be gained by taking learning incentives into account and by realizing that our DM solves two fundamental trade-offs:

1. Choosing between an arm of known quality (the safe one, S) and an arm of unknown quality (the risky one, R), where information on the latter's quality can be gathered through experimentation. (This is the learning dimension of the problem.)

2. Choosing between an arm that provides a relatively frequent stream of modest payoffs, and an arm that provides a less frequent stream of larger payoffs. (This is the dimension along which curvature in the utility function plays a role.)

When $h < s$, the risky arm's payoff frequency λ has to be rather high by Assumption 1 (otherwise the DM will always choose S and the problem is not meaningful). A high λ implies that pulling the risky arm is relatively informative: if the arm is of the good type, a payoff h should be observed soon; if not, the belief about the nature of the risky arm will quickly be revised downward by equation (1.1) – ending the experimentation process once the belief p falls below the cut-off p_i^*. So when λ is high, uncertainty about the quality of the risky arm (captured under point 1) is likely to be short-lived. Consequently, the consideration under point 2 becomes more important. Along this dimension, a more risk-averse DM prefers a frequent stream of modest payoffs to an infrequent stream of larger payoffs.[8] In the case where $h < s$ (but λ is high enough to meet Assumption 1), arm R is the one that offers a relatively frequent stream of modest payoffs (provided the arm is of good quality). The more risk-averse DM does not like the fact that this arm is risky (its quality is initially unknown and may turn out to be low, in which case it will never pay) but when λ is high, this risk is likely to be resolved soon and therefore of subordinate importance.

It is thus the trade-off between these two forces that determines a DM's decision to pull R or S. On the one hand, a more risk-averse DM is drawn towards the safe arm (the fact that the risky arm is of unknown quality introduces extra uncertainty in its payoff stream, which he dislikes). But, on the other hand, the DM realizes that pulling the risky arm enables him to reduce risk (which a risk-averse DM particularly likes). When λ is large, there is a high 'informational return' to pulling the risky arm, as there is a good chance that pulling R eliminates risk (which happens when a payoff h is observed, no matter how small h is). Subsequently, the DM is able to enjoy the (higher) utility stream provided by R in a world that no longer exhibits uncertainty on the nature of R. This explains the counterintuitive part of our result that a more risk-averse DM might be more willing to pull the risky arm than a less risk-averse DM.[9,10]

[8]Because of the difference in the concavity of the utility functions, the value of any increase in h is lower for the more risk-averse DM (due to decreasing marginal utility, which a risk-neutral DM for example does not experience).

[9]Proposition 1 continues to hold in the more general framework of Keller and Rady (2010): in their setup, even bad arms generate occasional payoffs equal to h – only at a lower frequency than good arms. More specifically, a good arm pays off according to a Poisson process with parameter λ_H, while this parameter equals λ_L for a bad arm (with $\lambda_H > \lambda_L$). Setting $\lambda_L = 0$ puts us back into the framework of Keller, Rady, and Cripps (2005) and simplifies the algebra considerably.

[10]This can be rephrased in terms of entropy reduction: the entropy of a Poisson distribution is increas-

In Appendix C we show that our result also arises in a simple two-period setup. It similarly carries over to an infinite-horizon, discrete-time version of the model.

1.4 Discussion

The counterintuitive part (b) of Proposition 1 is seemingly at odds with the result of Chancelier, De Lara, and De Palma (2009), who conclude that more risk-averse DMs are always more likely to pull the safe arm in bandit problems. Closer inspection of Theorem 1 in Chancelier, De Lara, and De Palma (2009), however, reveals that the assumption made there restricts payoffs in such a way that it only covers case (a) of our Proposition 1.[11] There, we obtain the same result. Since the continuous-time framework employed in this paper makes the existence of different regimes more transparent, it becomes apparent that there is a part of the parameter space (with $h < s$ and λ sufficiently high) in which the intuitive result does not arise.

Instances where λ is high (which means that the risky arm pays out frequently, conditional on it being 'good') are particularly likely to occur in online settings. There, information abounds and arrives at a high frequency. Reinforcement learning for example has the bandit problem at its core, there often referred to as the 'exploration vs. exploitation trade-off' (Sutton and Barto 1998). Such algorithms are, among other things, used to customize webpage advertisements to user-preferences. At each page visit, the algorithm faces a choice between, say, showing a well-understood ad which is known to generate infrequent per-click payoffs of considerable size s (e.g. an ad for expensive watches), or show an ad for a new product (which comes with lower per-click payoffs, $h < s$). Suppose that the market for the associated product is very competitive and a priori it is not known whether the brand behind the advertisement will become popular (this is the source of risk in the problem). If the brand does take off, it is expected to generate frequent clicks

ing in its parameter λ, as a result of which the expected entropy reduction (= uncertainty reduction = information production) is higher when the λ of the risky arm is higher. This makes it more attractive for a risk-averse DM to pull that arm.

[11] To see this, note that Theorem 1 of Chancelier, De Lara, and De Palma (2009) can be rewritten in our notation/model as: "Assume that there exists a concave increasing function $\varphi : \mathbb{R} \to \mathbb{R}$ such that $u_2(s) \geq \varphi(u_1(s))$ and $u_2(h) \leq \varphi(u_1(h))$ ...". Starting from a situation where $h > s$ (our case (a)), it is not possible to respect the concave increasing function φ and move to a situation in which $h < s$ (our case (b)) while satisfying the assumption that $u_2(s) \geq \varphi(u_1(s))$ and $u_2(h) \leq \varphi(u_1(h))$.

(high λ) – improving upon the payoff generated by the safe ad (in this case, 'the risky arm is good'). If the new brand does not take off ('the risky arm is bad'), clicks on the ad will be infrequent (low λ) and the webpage would have been better off by sticking with the old ad. In such a setup, our result demonstrates that equipping the machine learning algorithm with a more risk averse objective function might lead to a greater appetite for the risky arm (the unknown ad).

Alternatively, our counterintuitive result can arise in a labor market setup. Consider an interpretation of the two-armed bandit model which captures the career choice between becoming a worker or becoming an entrepreneur. When going down the latter route, the DM will face greater uncertainty about his long-run payoffs – at least initially (less so after he has learned the popularity of his product).[12] Our DM is uncertain on – say – his organizational talent, which can be either high or low. If it is high, he will obtain a higher utility level as an entrepreneur; if it is low, his firm will never take off and he is better off as a worker. Following the seminal paper by Kihlstrom and Laffont (1979), most earlier papers featuring this choice have started from the (widely accepted) premise that more risk-averse individuals will choose to become workers (the safer option, which immediately gives greater clarity on long-run payoffs), while the less risk-averse ones will choose to start a business. In this literature, the narrative of the 'risk tolerant entrepreneur' has been proposed as a parsimonious and plausible fix to the puzzling observation that entrepreneurs tend to earn less and bear more risk than salaried workers (Hamilton 2000; Moskowitz and Vissing-Jorgensen 2002).

But by taking learning and experimentation dynamics into account, this paper demonstrates that this popular narrative does not necessarily hold true. It opens up the possibility that more risk-averse individuals might be *more* willing to start with the riskier action (setting up a business), if taking such a risk produces a sufficient amount of information on their organizational talent.[13] This theoretical ambiguity could explain why previous empirical studies have reported mixed results regarding the effect of risk aversion on the

[12] See Kerr, Nanda, and Rhodes-Kropf (2014) and Manso (2015) for examples of this interpretation.

[13] i.e., if the information arrival rate λ is high enough. Bonatti and Hörner (2017) apply their bandit model to the labor market and argue that the information arrival rate λ is increasing in the amount of effort exerted by the DM, which seems intuitive. This suggests that high-effort exerting, relatively risk-averse DMs are more likely to display a greater preference for the risky arm than their less risk-averse counterparts (especially those who are not inclined to exert much effort).

decision to become an entrepreneur.[14] In a principal-agent setup, it furthermore implies that appointing a more risk-averse agent is no guarantee that the principal will see more 'safe' actions implemented.

More generally, our findings illustrate that it is not straightforward to infer risk preferences from observed actions when the setup is dynamic and offers scope for experimentation. Risk-aversion estimates obtained from game shows (such as *Deal or No Deal*[15]), which neglect the point made by this paper, might suffer from a serious bias (not only along the quantitative dimension, but even along the qualitative one).

[14] Compare Schiller and Crewson (1997), who report mixed results themselves, Barsky et al. (1997) who find no significant effect – Andersen et al. (2014) also falls in this category; and Cramer et al. (2002) , who do find a significant negative effect of risk aversion on the probability of becoming self-employed, but conclude that they are not able to make statements on causality.

[15] In this game show, which is similar to the Pandora's Box problem of Weitzman (1979) , a participant is typically presented with 26 suitcases – one of which becomes 'his' at the start of the game. Each case contains a monetary prize, the value of which is hidden to the participant (but he knows that the distribution of prizes is uniform over 26 pre-specified amounts). Subsequently, a game unfolds in which the participant has to open all remaining suitcases in a sequential manner. After each round, the show makes a cash offer to the participant, which he has to accept or reject. If he rejects ('no deal'), the game proceeds to the next round – until the participant makes a deal or all suitcases are opened (and the participant is left with the prize in 'his' case). Studies that try to elicit risk preferences from this game show include Post et al. (2008) and De Roos and Sarafidis (2009). See Andersen et al. (2008) for an extensive overview of studies inferring risk aversion from behavior in (dynamic) game shows.

1.A Appendix A: $u_2(x)/u_1(x)$ is strictly decreasing in x

LEMMA 1. Let $u_1 : [0,\infty) \to [0,\infty)$ be a strictly increasing function with $u_1(0) = 0$; let $\varphi : [0,\infty) \to [0,\infty)$ be an increasing and concave function, strictly so on some interval that includes the origin, with $\varphi(0) = 0$; finally, let $u_2 = \varphi \circ u_1$. Then $u_2(x)/u_1(x)$ is strictly decreasing in x.

PROOF. Measure u_1 on the horizontal axis and u_2 on the vertical axis. The graph of φ, the mapping from u_1 to u_2, is then a curve through the origin that is increasing and concave, strictly so on some interval that includes the origin, and the ratio $u_2(x)/u_1(x)$ is the slope of the chord from the origin to the point $(u_1(x), u_2(x))$. As we increase x, and thus $u_1(x)$, this slope strictly decreases and therefore $u_2(x)/u_1(x)$ is strictly decreasing in x.

1.B Appendix B: Generalizing Proposition 1

Using the representation of Pratt (1964), we can write $u(x) = \int_0^x \exp\left(\int_0^y -r(z)dz\right) dy$, where $r(x) = -u''(x)/u'(x)$ is the measure of absolute risk aversion associated with u. Let us consider an increase in the measure of absolute risk aversion by $\epsilon > 0$. This brings us to $u(x;\epsilon) = \int_0^x \exp\left(\int_0^y -(r(z)+\epsilon)\,dz\right) dy$.

From (1.2), we can rewrite the threshold belief associated with $u(x;\epsilon)$ as:

$$p^*(\epsilon) = \frac{ru(s;\epsilon)}{(r+\lambda)\lambda u(h;\epsilon) - \lambda u(s;\epsilon)}$$

and differentiate it with respect to ϵ to find that:

$$\begin{aligned}\operatorname{sgn}\left(\frac{\partial p^*}{\partial \epsilon}\right) &= \operatorname{sgn}\Big(u(h;\epsilon)\,\partial u(s;\epsilon)/\partial\epsilon - u(s;\epsilon)\,\partial u(h;\epsilon)/\partial\epsilon\Big)\\ &= \operatorname{sgn}\left(\frac{u(h;\epsilon)}{\partial u(h;\epsilon)/\partial\epsilon} - \frac{u(s;\epsilon)}{\partial u(s;\epsilon)/\partial\epsilon}\right)\end{aligned}$$

Writing $u_\epsilon(x;\epsilon)$ for the partial derivative with respect to the parameter ϵ, our aim is to show that $u(x;\epsilon)/u_\epsilon(x;\epsilon)$ is increasing $x > 0$, in which case we are back to Proposition 1 formulated in the main text. The following Lemma (which is a generalization of Lemma 1 in Appendix A) establishes this result and thereby generalizes the main result of our paper for an infinitesimal increase in risk aversion of any (twice-differentiable) Bernoulli utility function u.

To this end, define $f(y;\epsilon) = \int_0^y -(r(z)+\epsilon)\,dz$, giving

$$u(x;\epsilon) = \int_0^x \exp\left(f(y;\epsilon)\right) dy,$$

and $u_\epsilon(x;\epsilon) = \int_0^x f_\epsilon(y;\epsilon)\exp\left(f(y;\epsilon)\right) dy$. Noting that $f_\epsilon(y;\epsilon) = -y$, we see that

$$u_\epsilon(x;\epsilon) = -\int_0^x y\exp\left(f(y;\epsilon)\right) dy.$$

For future reference, we note that $u'(x;\epsilon) = \exp\left(f(x;\epsilon)\right)$, and $u'_\epsilon(x;\epsilon) = -x\exp\left(f(x;\epsilon)\right)$, where the prime denotes the derivative with respect to the variable x.

LEMMA 2. Given $u_\epsilon(x;\epsilon) < 0$ when $x > 0$, $u(x;\epsilon)/u_\epsilon(x;\epsilon)$ is increasing in $x \in [0,\infty)$.

PROOF.

$$\begin{aligned}
\operatorname{sgn}\Big(u(x;\epsilon)/u_\epsilon(x;\epsilon)\Big)' &= \operatorname{sgn}\Big(u_\epsilon(x;\epsilon)\,u'(x;\epsilon) - u(x;\epsilon)\,u'_\epsilon(x;\epsilon)\Big) \\
&= \operatorname{sgn}\Big(-\exp\left(f(x;\epsilon)\right)\textstyle\int_0^x y\exp\left(f(y;\epsilon)\right)dy + x\exp\left(f(x;\epsilon)\right)\int_0^x \exp\left(f(y;\epsilon)\right)dy\Big) \\
&= \operatorname{sgn}\Big(\exp\left(f(x;\epsilon)\right)\textstyle\int_0^x (x-y)\exp\left(f(y;\epsilon)\right)dy\Big) \\
&\geq 0
\end{aligned}$$

with the inequality being strict when $x > 0$.

1.C Appendix C: A two-period model

Here, we show that our result also holds in a discrete-time, two-period setup. Without loss of generality, we abstract from discounting.

In each period, the safe arm S pays out a lump-sum s with probability $\frac{1}{2}$, whereas the risky arm R pays out a lump-sum h with probability $\frac{1}{2}\gamma$ if it is of good quality, while a bad risky arm never pays off.[16] As in the main text, we use p to denote the DM's belief that R is of good quality.

At this stage, we rephrase Assumption 1 as follows:

ASSUMPTION C1. $0 = u(0) < \frac{1}{2}u(s) < \frac{1}{2}\gamma u(h)$.

In this simple setup, one can analyze the expected utilities resulting from the four possible strategies:

1. Playing the safe arm in both periods (SS) yields a total expected utility equal to $u(s)$.
2. Playing RR yields a total subjective expected utility equal to $\gamma u(h)p$.
3. Playing SR yields a total subjective expected utility equal to $\frac{1}{2}u(s)+\frac{1}{2}\gamma u(h)p$. This is dominated by SS if p is low, and dominated by RR if p is high.
4. Playing RS conditionally, i.e. only sticking with R after a success in period 1, yields a total subjective expected utility equal to $\frac{1}{2}\gamma u(h)p+\left[\frac{1}{2}\gamma u(h)\left(\frac{1}{2}\gamma p\right)+\frac{1}{2}u(s)\left(1-\frac{1}{2}\gamma p\right)\right] = \frac{1}{2}u(s)+\frac{1}{2}\gamma\left[u(h)+\frac{1}{2}\gamma u(h)-\frac{1}{2}u(s)\right]p$.

Equating expected utility from SS and RS gives a lower cut-off belief:

$$p^{\ell} = \frac{\frac{1}{2}u(s)}{\frac{1}{2}\gamma\left[\left(1+\frac{1}{2}\gamma\right)u(h)-\frac{1}{2}u(s)\right]},$$

while equating expected utility from RR and RS gives an upper cut-off belief:

[16] Our maintained assumption is that a good risky arm is preferred to the safe arm. A necessary condition for the counterintuitive part of Proposition 1 from the main text is that the lump-sum h can nevertheless be smaller than the lump-sum s. Consequently we need to allow for the possibility that the probability of a payoff from R is greater than the probability of a payoff from S, and hence the requirement that the probability of a payoff from S is < 1.

$$p^u = \frac{\frac{1}{2}u(s)}{\frac{1}{2}\gamma\left[\left(1-\frac{1}{2}\gamma\right)u(h)+\frac{1}{2}u(s)\right]}.$$

When the DM's belief $p < p^\ell$, it is optimal for him to play S in both periods. Similarly, when $p > p^u$ it is optimal to play R in both periods (even if no lump-sum arrived in period 1). For intermediate beliefs, i.e. when $p^\ell < p < p^u$, it is optimal to play R in the first period and switch to S in period 2 if no lump-sum arrived.

Defining utility functions u_1 and u_2 as in the main body of the paper (with u_2 exhibiting greater risk aversion), it is straightforward to show that the sign of the difference in cut-offs again satisfies:

$$\operatorname{sgn}(p_2^\ell - p_1^\ell) = \operatorname{sgn}(p_2^u - p_1^u) = \operatorname{sgn}\left(\frac{u_2(s)}{u_1(s)} - \frac{u_2(h)}{u_1(h)}\right).$$

In this case, the counterintuitive part of the parameter space opens up when $\gamma > 1$ (but notice that because probabilities cannot be greater than 1, we also need that $\gamma \leq 2$). When $\gamma > 1$, the risky arm is expected to pay out more frequently than the safe arm and R second-order stochastically dominates S. This makes the risky arm more attractive to the more risk averse DM_2.

By following similar steps to the discrete-time formulation model of Heidhues, Rady, and Strack (2015) of the infinite-horizon Poisson bandit model, one can verify that the same logic continues to apply in that setup.

Chapter 2

The Status Quo and Belief Polarization of Inattentive Agents: Theory and Experiment

Co-authored with Andrei Matveenko (University of Copenhagen) and Silvio Ravaioli (Columbia University).

2.1 Introduction

The fundamental characteristic of many democratic votes is that they simplify a complex choice into a binary problem of choice between the status quo policy with a known outcome and a new policy with multiple possible outcomes. Consider the 2016 Brexit referendum as an illustrative example – the consequences of choosing to leave the European Union are uncertain and could be: no-deal Brexit, Brexit with a deal, or soft Brexit, among others. At the same time, public opinion surveys suggest that British society a few months after the Brexit referendum was even more polarized than on the referendum day, with voters becoming more commited to the vote in the referendum (Smith 2019)[1]. Generally, an increase in opinion polarization is a widely discussed phenomenon in both the academic literature[2] and public discourse. However, the question of how binary choices between alternatives with multiple possible outcomes might be connected with belief polarization remains unanswered. In this paper we present a model that de-

[1]Similarly, Ciani et al. (2019), using a pan-European dataset show empirically that expectations about pension reform do not converge as a result of announcements or implementations.

[2] See, for instance Poole and Rosenthal 1984; McCarty, Poole, and Rosenthal 2008; Gentzkow, Shapiro, and Taddy 2016.

scribes what information a citizen chooses to obtain before voting, and as a result, we show that the relative position of the status quo determines the evolution of opinions, and may lead to systematic polarization even between agents with the same prior beliefs. We test and confirm these results in a lab experiment setting.

We model the agent to be rationally inattentive, following Sims (1998a) and Sims (2003b), which allows us to account for endogenous information acquisition without imposing any exogenously given biases. Since information is plentiful but attention is scarce, the agent chooses to learn the essential pieces of information for her decision problem, that is, whether the payoff of a new policy is higher or lower than the agent's valuation of the status quo. We show that under some conditions the agent systematically updates her belief in the wrong direction with respect to the true realization, even in expectations. When we include multiple agents, this can lead to belief polarization in expectations, even for the agents with the same prior expected belief about the payoff of the new policy.

We start our paper with a simple illustrative example. The information acquisition strategies in the example are very limited, but they provide an intuition for our main result. In Section 2.3 we consider a static decision problem with a rationally inattentive agent, n states and two actions, in the manner of Matějka and McKay (2015). However, in contrast with Matějka and McKay (2015), our focus is on the evolution of beliefs. The primary indicator that we consider for the belief evolution over the optimal signals is the change in the mean of beliefs about the payoff of the new policy. In order to judge how the belief about the expected payoff from the new policy changes after the signal is received and the option is chosen, we take the position of an external observer. The observer knows which state of the world was realized and thus what the realized state-dependent payoff of the new policy is. We can study whether the posterior expected belief about the value of the new policy is moving from the prior expected belief away from or towards the actual payoff of the new policy.

We show that the sign of the change in the mean of beliefs is the same as the sign of the difference between the realized outcome of the new policy and the value of the status quo (Proposition 1). This result simplifies the considerations of the complex endogenous beliefs evolution process. At the same time, it also demonstrates the link between the value of the status quo and the updating process. Due to costly information acquisition, the rationally inattentive agent chooses only the information necessary to disentangle whether to select the status quo or the new policy. Thus, the agent endogenously divides the states that determine the payoff of the new policy into categories that separate the states into two groups (based on the payoff of the new policy being higher or lower than the status quo), and chooses to acquire information only to disentangle from which of these two categories the realized outcome is. We refer to this endogenous classification behavior of the agent as the *state pooling effect*.

As a result of the state pooling effect, we observe that the agent might update her expected posterior belief about the payoff of the new policy in the opposite direction from the realized payoff of the new policy. We formulate a simple criterion to identify when the agent updates in the opposite direction from the realized value of the new policy (Theorem 1). We analyze the magnitude of the distance between the expected posterior belief and the prior belief. We show that the higher the realized value of the new policy, the greater the distance is (Proposition 2). Thanks to this monotonicity result we can identify the set of payoffs of the new policy for which the agent updates in the opposite direction from the true payoff: it lies between the prior expected payoff of the new policy and the payoff of the status quo.

The above results suggest that two agents might become polarized when they either differ in their prior expected beliefs about the payoff of the new policy or in the valuation of the status quo (Theorem 2). With our results for the general case in hand, we focus in Section 2.4 on the case with three states, in order to illustrate in detail implications of the previously stated results, and to show that the behavior resembling over-optimism[3] and over-pessimism can appear as a result of the agent's inattentiveness. We also consider the impact of the marginal cost of information on belief polarization. In Section 2.5 we demonstrate, using a numerical example, that when information is cheaper, the agent's set of prior beliefs, in which she does not acquire any information, shrinks, and for other prior beliefs she chooses to learn more. Hence, when the cost of information is lower, the polarization of agents is more severe.

[3]The important consequences of over-optimism can be seen, for instance, in Beaudry and Willems (2018), who show that recessions, fiscal problems, and Balance of Payment-difficulties are more likely to arise in countries where past growth expectations have been overly optimistic.

This paper contributes to the literature on belief polarization. Modelling the agent as rationally inattentive relieves us of the need to assume exogenously given biases[4] or bi-modality of preferences (Dixit and Weibull 2007), which are common in the preceding literature. This in turn allows us to move in a new direction away from findings that the beliefs of Bayesian agents would converge over time and that they will almost surely assign probability 1 to a true state (Savage 1954; Blackwell and Dubins 1962). The closest theoretical paper to ours is Nimark and Sundaresan (2019), which studies the question of how inattentiveness can lead to persistent belief polarization. In our model, agents with the same prior beliefs might be polarized in expectation, whereas in Nimark and Sundaresan (2019) two agents with the same prior beliefs will always choose the same signal structures and can become polarized only when they receive sufficiently different signal realizations. We consider a multiple states environment with full flexibility in the shape of the received signal, which allows us to discover a state pooling effect that cannot arise in the environment of Nimark and Sundaresan (2019)[5].

This paper also adds to the rational inattention literature[6] by studying the evolution of beliefs alongside the manifestation of the crucial implications of incorporating the safe option into the choice set. Our findings give reason for caution for empirical and experimental work inferring preferences for information. In particular, this paper provides a disciplined model that suggests how the preference for the skewed information might depend on the value of the status quo and thus provides an important channel that is missing in the research on whether people prefer negatively or positively skewed information, e.g. Masatlioglu, Orhun, and Raymond (2017).

In Section 2.6, we report an experiment designed to test our theoretical results; in particular the state pooling effect and the presence of belief polarization in expectations. The subjects are presented with a binary choice, and can acquire instrumentally valuable information from advisors before making their decision. In the first task, all advisors are

[4] Gerber and Green (1999) review the literature that invokes some biases in learning or perception in order to modify Bayesian updating.

[5] They focus on the persistence of the polarization and thus study a two state environment where the agent might receive binary signals and the information structure of the agent is characterized by error probabilities.

[6] A survey of the literature on rational inattention is provided in Mackowiak, Matějka, and Wiederholt (2018). For a posterior-based approach see Caplin and Dean (2015) and a dynamic discrete choice model is presented in Steiner, Stewart, and Matějka (2017).

evaluated separately and provide degenerate signal structures. Subjects report their willingness to renounce the signal. In each trial of task 2 only two advisors are displayed, and subjects make a binary choice between signal structures presented by them. In contrast with Charness, Oprea, and Yuksel (2018), we focus on variation in the value of the status quo and not on the variation in prior beliefs. After subjects have made their information choices, we elicit their beliefs about the likelihood of a signal, given the advisor and likelihood of each state as a function of the signal provided by the chosen information structure. A key feature of our design, distinctive from the previous literature, is that we consider the environment with more than two states.

We find that subjects do react to the value of the status quo, as predicted by our theoretical model, and that they display preference for state pooling information structures. We also show that the probability of choosing an advisor increases with the instrumental value of the corresponding information structure. Importantly, we verify the consequent belief polarization in expectations. As a byproduct of our design, we replicate the results of Ambuehl and Li (2018) in a setting with three possible states. When subjects are asked to report their subjective valuation for information structures, they display compression in the evaluations. They underreact to increasing instrumental value of information and display a strong preference for advisors that provide certainty in the posterior beliefs.

2.2 Example

In this section we illustrate the logic of our result in a simple example. In this example the information acquisition strategy is highly restricted, but it still allows us to demonstrate the main results of the paper: the state pooling effect and belief updating in the opposite direction from the realized value.

Payoffs. There is a single payoff-maximizing agent. The agent faces a choice between a currently implemented policy (status quo) and a new policy. The currently implemented policy brings her a payoff $R \in \mathbb{R}$. The outcome of the new policy is uncertain and depends on the realized state of the world. There are three possible realizations of the state of the world. In the state $s = i$, $i \in \{1, 2, 3\}$, the new policy brings a payoff $v_i \in \mathbb{R}$. The probabilities with which the states are realized are g_1 where $0 \leq g_1 \leq 2/3$, $g_2 = 1/3$ and

$g_3 = 2/3 - g_1$ for states $s = 1, 2, 3$, respectively.

Signals. The agent has an opportunity to learn about the realized state of the world, but her learning opportunities are restricted. Let us assume that she can ask "yes/no" questions, each for a cost $\alpha > 0$. We restrict the agent's questions to the form "Is the realized state i, or is it not?", where $i \in \{1, 2, 3\}$. Formally, her initial partition of the states of the world is $\{1, 2, 3\}$. By asking one question , she can arrive to any partition of the form $(\{i, j\}, \{k\})$, where $i, j, k \in \{1, 2, 3\}$ and $i \neq j \neq k$. By asking two questions, she arrives to partition $(\{1\}, \{2\}, \{3\})$.

For example, the partition $(\{1, 3\}, \{2\})$ means that the agent knows that the realized state belongs either to $\{1, 3\}$ or to $\{2\}$. That is, the agent can choose the following partitions: $(\{1, 2, 3\})$ – with no cost, $(\{1, 2\}, \{3\})$, $(\{1, 3\}, \{2\})$, $(\{2, 3\}, \{1\})$ – with cost α and $(\{1\}, \{2\}, \{3\})$ – with cost 2α. We can now identify the optimal learning strategy of an agent.

Actions. First, let us show that if the agent decides to ask at least one question, then she will not ask the second one. Second, let us show that if the agent asks one question, the optimal partitions are: $(\{1\}, \{2, 3\})$, if $v_1 < R < v_2 < v_3$, and $(\{1, 2\}, \{3\})$, if $v_1 < v_2 < R < v_3$. In Section 2.4 we show that a similar partitioning arises as optimal behavior of the rationally inattentive agent and we call it a *state pooling effect.*

In the remainder of this section, we consider the situation $v_1 < R < v_2 < v_3$; the situation $v_1 < v_2 < R < v_3$ can be considered analogically. The partitions $(\{1, 2\}, \{3\})$ and $(\{1, 3\}, \{2\})$ cannot be optimal because, for example, if state 1 is realized, the agent in both cases remains uncertain about the optimal action, while if the agent partitions states into categories $(\{1\}, \{2, 3\})$, she either learns that the state is 1, and thus chooses the status quo, or learns that the state is 2 or 3, and thus chooses the new policy. Moreover, further learning would not change the agent's action, and, since learning is costly, the agent will not do it.

Belief. How does the agent decide whether to participate in costly learning, that is, whether to ask at least one question? Her prior expected belief about the outcome is

$$V_{NL} = \max\{R, g_1 v_1 + \frac{1}{3}v_2 + \left(\frac{2}{3} - g_1\right) v_3\}.$$

The expected posterior belief about the outcome is

$$V_L = g_1 R + \frac{1}{3}v_2 + \left(\frac{2}{3} - g_1\right) v_3.$$

The agent compares V_{NL} and $V_L - \alpha$ and decides whether to participate in the costly learning.

Direction of updating. If the realized state of the world is $s = 2$, then the expected posterior belief about the value of the new policy is

$$\mathbb{E}_i[\mathbb{E}(v|i)|s = 2)] = \frac{1/3}{1 - g_1}v_2 + \frac{2/3 - g_1}{1 - g_1}v_3.$$

Interestingly, if $V_L - \alpha > V_{NL} > v_2$, in the state of the world 2 the following inequality holds:

$$\mathbb{E}_i[\mathbb{E}(v|i)|s = 2)] > V_{NL} > v_2.$$

The last inequality implies that in the state $s = 2$ the agent's conditional expected posterior belief about the value of the new policy is higher than the agent's prior belief. At the same time, the agent's prior belief about the value of the new policy is higher than v_2. That is, the agent updates her belief about the value of the new policy in the opposite direction from the realized value v_2.

2.3 The model

In this section we describe the general case of the agent's decision problem, introduce a methodology for assessment of beliefs' evolution, and present the main theoretical results. The structure of this section is as follows. In Subsections 2.3.1 and 2.3.2 we describe the agent's problem, which is a special case from Matějka and McKay (2015). Subsections 2.3.3 and 2.3.4 present our main results about the evolution of the beliefs given the true state of the world. Subsection 2.3.5 discusses beliefs' polarization of rationally inattentive agents.

2.3.1 Description of the setup

A single agent faces a problem of discrete choice between two options. The first option, which we refer to as *a new policy,* provides a payoff $v_s \in \mathbb{R}$ that depends on the realized state of the world $s \in S = \{1, \ldots, n\}$, where $n \in \mathbb{N}$. When we say that the state is realized, we mean that it became possible for the potential outcome of the new policy to be evaluated. The states are labeled in ascending order $v_1 < v_2 < ... < v_n$. The second option, which we refer to as *a status quo policy,* yields a known fixed payoff $R \in \mathbb{R}$. We assume that $v_1 < R < v_n$ in order to exclude trivial cases.[7] That is, there exists a unique $k \in \{1, 2, ..., n-1\}$, such that $v_k \leq R < v_{k+1}$.

The agent is uncertain which state of the world s is going to be realized and we denote her prior belief as a vector of probabilities $\mathbf{g} = [g_1 \ \ g_2 \ \ ... \ \ g_n]^T$, where $\mathcal{P}(s = j) = g_j$, $\forall j \in S$; $\sum_{j=1}^{n} g_j = 1$ and $g_j > 0$, $\forall j \in S$. We model the agent to be rationally inattentive in the fashion of Sims (2003b). The agent wishes to select the option with the highest payoff. Prior to making the decision, she has a possibility to acquire some information about the actual value of the new policy, which is modeled as receiving a signal $x \in \mathbb{R}$. The distribution of the signals, $f(x, s) \in P(\mathbb{R} \times S)$, where $P(\mathbb{R} \times S)$ is the set of all probability distributions on $\mathbb{R} \times S$, is subject to the agent's choice. Upon receiving a signal, the agent updates her belief using Bayes rule. However, observing a signal is costly and we assume the cost to be proportional to the expected reduction in entropy[8] between the agent's prior and posterior beliefs. For detailed treatment of entropy, see, for example, Cover and Thomas (2012). Upon receiving a signal, the agent chooses an action, and her choice rule is modeled as $\sigma(x) : x \rightarrow \{\text{new policy}, \text{status quo}\}$. Given the updated belief, the agent chooses the action with the highest expected payoff. The timing of the decision problem is depicted in Figure 2.1.

2.3.2 Agent's problem

According to Lemma 1 from Matějka and McKay (2015), the choice behavior of the rationally inattentive agent can be found as a solution to a simpler maximization problem that

[7]If $R \leq v_1$, the safe option is weakly dominated by the risky option, and if $R \geq v_n$ the risky option is weakly dominated by the safe option. In both of these cases the agent does not have incentives to acquire information about the realization of the state of the world.

[8]The entropy $H(Z)$ of a discrete random variable Z with support $\mathcal{Z}$ and probability mass function $\mathcal{P}(z) = Pr\{Z = z\}, z \in \mathcal{Z}$ is defined by $H(Z) = -\sum_{z \in \mathcal{Z}} p(z) \log p(z)$.

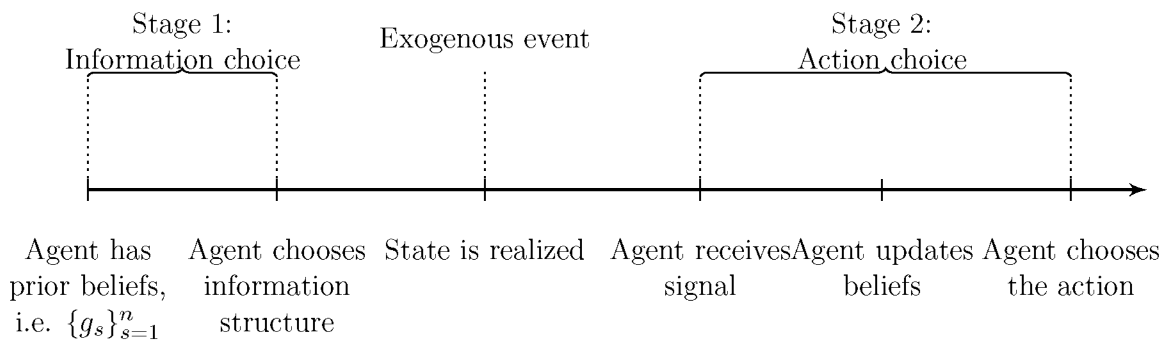

Figure 2.1: Timing of the events in the problem. The decision problem consists of two stages: an information strategy selection stage and a standard choice under uncertainty stage.

is stated in terms of state-contingent choice probabilities alone. The information strategy is characterized by the collection of conditional probabilities of choosing option i in state of the world s : $\mathcal{P} = \{\mathcal{P}(i|s) \big| \ i = 1, 2; \ s \in S\}$, where $i \in \{\text{new policy, status quo}\} = \{1, 2\}$ denotes the option and s is the state. The agent's problem is to find an information strategy maximizing the expected utility less the information cost. That is, the agent solves:

$$\max_{\{\mathcal{P}(i|s)|i=1,2;\ s\in S\}} \left\{ \sum_{s=1}^{n} \left(v_s \mathcal{P}\left(i = 1|s\right) + R\mathcal{P}(i = 2|s) \right) g_s - \lambda\kappa \right\}, \tag{2.1}$$

subject to

$$\forall i: \ \mathcal{P}(i|s) \geq 0 \qquad \forall s \in S \ , \tag{2.2}$$

$$\sum_{i=1}^{2} \mathcal{P}(i|s) = 1 \qquad \forall s \in S \ , \tag{2.3}$$

and where

$$\kappa = \underbrace{-\sum_{i=1}^{2} \mathcal{P}(i) \log \mathcal{P}(i)}_{\text{prior uncertainty}} - \sum_{s=1}^{n} \left(\underbrace{-\left(\sum_{i=1}^{2} \mathcal{P}(i|s) \log \mathcal{P}(i|s) \right) g_s}_{\text{posterior uncertainty in state } s} \right). \tag{2.4}$$

$\mathcal{P}(i)$ is the unconditional probability that the option i will be chosen and is defined as

$$\mathcal{P}(i) = \sum_{s=1}^{n} \mathcal{P}(i|s) g_s, \qquad i = 1, 2.$$

Here κ denotes the expected reduction in entropy between the prior and the posterior beliefs about the choice outcome, $\lambda \geq 0$ is the unit cost of information, and thus, $\lambda\kappa$ reflects the cost of generating signals with different precision.

Matějka and McKay (2015) study a general case of the static problem described above, and show that at the optimum the probabilities with which the agent chooses the options follow the modified multinomial logit formula.This result translates into our setting in the following way:

Lemma 1. *Conditional on the realized state of the world $s \in S$, the probability of choosing a new policy for $\lambda > 0$ is implicitly defined by:*

$$\mathcal{P}(i = 1|s) = \frac{\mathcal{P}(i = 1)e^{\frac{v_s}{\lambda}}}{\mathcal{P}(i = 1)e^{\frac{v_s}{\lambda}} + (1 - \mathcal{P}(i = 1))e^{\frac{R}{\lambda}}},$$

the probability of choosing the status quo is

$$\mathcal{P}(i = 2|s) = \frac{(1 - \mathcal{P}(i = 1))e^{\frac{R}{\lambda}}}{\mathcal{P}(i = 1)e^{\frac{v_s}{\lambda}} + (1 - \mathcal{P}(i = 1))e^{\frac{R}{\lambda}}},$$

where $\mathcal{P}(i = 1)$ is the unconditional probability of choosing a new policy.

When $\lambda = 0$, the agent chooses from two available options: the new policy or the status quo, the option with the highest value with probability one.

Proof. Lemma 1 is a direct consequence of Theorem 1 from Matějka and McKay (2015). □

Non-learning areas. An important feature of the solution to the agent's problem is that for the given vector of payoffs of the new policy $(v_1, ..., v_n)$, the value of the status quo R and the marginal cost of information λ there exist prior beliefs of the agent for which she decides not to acquire any information. In this case, we say that the agent is in a *non-learning area.* Once the agent is in a non-learning area, she makes her decision based on her prior beliefs only. That is, when the agent is in a non-learning area, if $\mathbb{E}v = \sum_{s=1}^{n} v_s g_s > R$ then the agent chooses the new policy with certainty, if $\mathbb{E}v < R$, then the agent chooses the status quo with certainty, and if $\mathbb{E}v = R$ then the agent is indifferent between the two policies. Let us assume, without loss of generality, that in the latter case, the agent would decide to keep the status quo. Given this assumption, the unconditional choice probabilities of the agent who is in a non-learning area are either 0 or 1. If the agent's prior is such that she decides to acquire at least some information, we say that the agent is in a *learning area.* For such prior beliefs, the unconditional choice probabilities lie in the open interval $(0, 1)$.

2.3.3 Description of beliefs evolution

The uncertainty in this model is about the realized state of the world and thus about the actual payoff of the new policy. Without the information acquisition stage of the problem the agent would choose the option based on the comparison of the status quo payoff R with the agent's prior expected payoff from the new policy being

$$\mathbb{E}v = \sum_{s=1}^{n} v_s g_s.$$

In order to judge how this expected payoff from the new policy changes after the signal is received and the option is chosen, we take the position of an external observer. The observer knows that a realized state of the world is s^* and is interested in the agent's posterior expected belief about the payoff of the new policy v given the realized state s^*. Note that the agent's posterior belief is given by the signal she receives and thus the observer not only wants to know what the expected posterior belief is for a given signal, but is interested in the expected posterior belief about the new policy, on average across all possible signals the agent may receive. Since there is a one to one mapping between the selected information structure and consequently chosen action, the posterior expected payoff of the new policy is

$$\mathbb{E}_i[\mathbb{E}(v|i)|s^*] = \sum_{i=1}^{2} \left(\sum_{s=1}^{n} v_s \mathcal{P}(s|i) \right) \mathcal{P}(i|s^*),$$

where option $i \in \{1, 2\} = \{\text{new policy}, \text{status quo}\}$. For the optimally behaving rationally inattentive agent, who is solving the problem (2.1)–(2.4), the posterior expected belief can be rewritten as[9]:

$$\mathbb{E}_i[\mathbb{E}(v|i)|s^*] = \sum_{s=1}^{n} v_s g_s \frac{\mathcal{P}(i=1|s^*)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1|s^*))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}. \tag{2.5}$$

The primary indicator for the expected belief evolution over the optimal signals that we consider is the change in the mean of beliefs about the payoff of the new policy that can be defined as $\Delta(s^*) = \mathbb{E}_i[\mathbb{E}(v|i)|s^*] - \mathbb{E}v$. In particular, we are interested in the sign of $\Delta(s^*)$, which informs us whether the posterior expected belief in state s^* is moving from $\mathbb{E}v$ towards v_1 or v_n or stays equal to $\mathbb{E}v$.

Proposition 1. *Given that the agent is in a learning area and that the realized state of the world is s^*, the sign of the change in the mean of beliefs about the payoff of the new policy $\Delta(s^*)$ is the same as the sign of $(v_{s^*} - R)$.*

Proof. The proof is presented in Appendix 2.B.

□

Proposition 1 significantly simplifies the considerations of the beliefs evolution. At the same time, it also demonstrates the important link between the value of the status quo R and the updating process.

2.3.4 Updating in the opposite direction from the realized value

In this section, we show the impact of the value of the status quo R on the opinion polarization of inattentive agents. For the rest of this section we assume that the agent is in the learning area, i.e. $0 < \mathcal{P}(i = 1) < 1$.

Definition 1. We say that the agent is updating in the *opposite direction from the realized value* v_{s^*} in the state $s^* \in S$, if the condition $(v_{s^*} - \mathbb{E}v) \cdot \Delta(s^*) = (v_{s^*} - \mathbb{E}v)(\mathbb{E}_i[\mathbb{E}(v|i)|s^*] - \mathbb{E}v) < 0$ is satisfied.

In the following theorem we provide conditions for the presence of the states in which the agent is updating in the opposite direction from the realized value of the new policy.

[9]The derivation of Formula (2.5) is in Appendix 2.A.

Theorem 1. *If and only if, in a state $s^* \in S$ holds that $(v_{s^*} - R)(\mathbb{E}v - v_{s^*}) > 0$, then, in this state of the world, the agent is updating in the opposite direction from the realized value v_{s^*}.*

Proof. We need to show that if $(v_{s^*} - R)(\mathbb{E}v - v_{s^*}) > 0$ is satisfied, then $(v_{s^*} - \mathbb{E}v)(\mathbb{E}_i[\mathbb{E}(v|i)|s^*] - \mathbb{E}v) < 0$ is also satisfied. By Proposition 1, the sign of $v_{s^*} - R$ is the same as the sign of $\mathbb{E}_i[\mathbb{E}(v|i)|s^*] - \mathbb{E}v$, which implies the condition needed. The proof in the opposite direction is analogical. □

Theorem 1 provides a simple criterion for updating in the opposite direction from the realized value of the new policy. The criterion does not require solving the agent's problem and is formulated in the primitives of the model only. Intuition for the result presented in Theorem 1 is as follows. Due to costly information acquisition, the rationally inattentive agent chooses only the necessary information in order to disentangle whether to select the status quo or the new policy. This leads to the **state pooling effect**, when the agent endogenously divides the states into two categories (those in which the payoff of the new policy is higher than that of the status quo and those in which it is lower) and chooses only information that helps to disentangle which of these two categories the realized state s^* is from. Namely, as Proposition 1 states, for all the states s in which $v_s > R$ the expected posterior belief about the value of the new policy is higher than the prior belief; thus all such states are pooled into one category. Similarly, all the states s for which $v_s < R$ are pooled into another category.

It is important to notice that the agent's expected posterior belief is changing for different realized states even when such states are from the same category. Notice that so far we have shown that updating in the opposite direction from the realized value of the new policy can occur in some realized state s^*. A natural question arises: How is the difference between the prior and the posterior expected payoff from the new policy $\Delta(s^*)$ influenced by a different realized true state s^*? The answer to this question is provided by the following proposition.

Proposition 2. *The change in the mean of beliefs $\Delta(s^*)$ is an increasing function of s^*.*

Proof. The proof is presented in Appendix 2.C. □

Proposition 2 significantly helps us to identify the set of states $W \subset S$ in which the agent updates in the opposite direction from the realized payoffs v_{s^*} in corresponding states. First, it follows from Proposition 1 that $\Delta(s^* = 1) < 0$ and $\Delta(s^* = n) > 0$. These findings, together with Proposition 2, imply that $\Delta(s^*)$ reaches its minimum in state 1, its maximum in state n and $\Delta(s^*) = 0$ occurs in between. We remind the reader, at this point, that we have defined state k such that $v_k \leq R < v_{k+1}$. Thus, due to Proposition 1, $\Delta(k) \leq 0$ and $\Delta(k+1) > 0$. The last two inequalities together with Proposition 2, imply that for all $s^* \leq k$ the change of the mean of beliefs $\Delta(s^*) \leq 0$ and that for all $s^* \geq k+1$ holds that $\Delta(s^*) \geq 0$. We know that in states for which the condition $(\mathbb{E}v - v_{s^*}) \cdot \Delta(s^*) > 0$ is satisfied, the agent updates in the opposite direction from the realized payoff of the new policy. Let us assume that the agent's prior expected value of the new policy is $\mathbb{E}v > R$. Then one can see that the agent is updating towards the true payoff of the new policy for all states where $\Delta(s^*)$ is negative. However, updating in the opposite direction from the realized value occurs for all states that have payoffs smaller than $\mathbb{E}v$ and at the same time higher than R (see Figure 2.2).

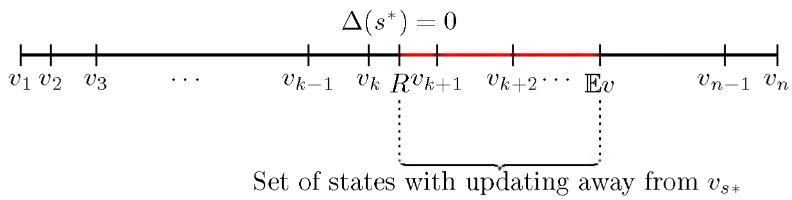

Figure 2.2: Set of states W (the red part of the line) for $\mathbb{E}v > R$, where the agent updates in the opposite direction from the realized value of the new policy

If we assume that $\mathbb{E}v < R$ then updating in the opposite direction from the realized value would happen in all states s^* for which it holds that $R > v_s^* > \mathbb{E}v$. We can write that

$$W = \begin{cases} \{s \mid R < v_s < \mathbb{E}v\}, & \text{if } \mathbb{E}v > R, \\ \{s \mid \mathbb{E}v < v_s < R\}, & \text{otherwise.} \end{cases}$$

It is worth noticing that the set of states where the agent updates in the opposite direction from the realized value are those where payoffs are neither very high nor very low. This has significant implications for predictions when the inattentive agents become polarized. We can also see that the number of states in the set W is determined by the status quo payoff R and by the prior expected value of the new policy.

2.3.5 Belief polarization

Let us now consider a situation with two agents $j = 1, 2$. The agents are facing the binary choice described in Section 2.3.2; however, they might have (i) different preferences about the same status quo policy R^j and, (ii), different prior beliefs about the value of the same new policy $\mathbb{E}^j v$. The expected posterior belief of agent j about the value of the new policy, conditional on the realized state s^*, is denoted by $\mathbb{E}_i^j[\mathbb{E}(v|i)|s^*]$. The difference between the expected posterior beliefs of agent j in the state s^* and the prior beliefs of agent j is denoted by $\Delta_j(s^*)$, $\Delta_j(s^*) = \mathbb{E}_i^j[\mathbb{E}(v|i)|s^*] - \mathbb{E}^j v$.

Definition 2. We say that two agents $j = 1, 2$, characterized by the pair $(R^j, \mathbb{E}^j v)$ and choosing between actions $i = \{1, 2\}$, become *polarized in the state* $s^* \in S$ when the following two conditions are satisfied

1. $|\mathbb{E}_i^1[\mathbb{E}(v|i)|s^*] - \mathbb{E}_i^2[\mathbb{E}(v|i)|s^*]| > |\mathbb{E}^1 v - \mathbb{E}^2 v|$.

2. $\Delta_1(s^*) \cdot \Delta_2(s^*) < 0$.

The first condition ensures that the expected posterior beliefs in the state s^* of two agents are further apart than the expected prior beliefs, whereas the second ensures that they update in opposite directions in the state s^*. In the following theorem we provide conditions for the presence of states of the world in which the agents become polarized.

Theorem 2. *Let us assume that there are two agents $j = 1, 2$, characterized by the pair $(R^j, \mathbb{E}^j v)$. If in state of the world $s^* \in S$ the conditions $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^1) > 0$ and $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^2) < 0$ hold, then in expectations the two agents or two groups of atomistic agents become polarized in this state of the world.*

Proof. Without loss of generality, let us assume that $\mathbb{E}^1 v > \mathbb{E}^2 v$. For the condition $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^1) > 0$ to be satisfied, it is necessary that $v_{s^*} > R^1$. Proposition 1 states that in this case $\Delta_1(s^*) > 0$. Analogously, the second condition $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^2) < 0$ holds when $v_{s^*} < R^2$, which further implies that $\Delta_2(s^*) < 0$. That is, two agents $j = 1, 2$ update in different directions and the expected posterior beliefs are farther away from each other than the priors are. Both conditions from Definition 2 are satisfied and the agents, indeed, become polarized in state s^*. □

2.3.6 Beliefs' convergence and divergence of beliefs updated in the same direction

In order to draw the whole picture of all possible situations (directions of belief updating), we describe our framework's predictions on when the beliefs of two agents converge to each other and at the same time move closer to the true value of the new policy.

Definition 3. We say that two agents $j = 1, 2$, characterized by the pair $(R^j, \mathbb{E}^j v)$ and choosing between actions $i = \{1, 2\}$, *converge in their beliefs in the state* $s^* \in S$ when the following two conditions are satisfied

1. $|\mathbb{E}_i^1[\mathbb{E}(v|i)|s^*] - \mathbb{E}_i^2[\mathbb{E}(v|i)|s^*]| < |\mathbb{E}^1 v - \mathbb{E}^2 v|$.

2. $\Delta_1(s^*) \cdot \Delta_2(s^*) > 0$.

Theorem 3. *Let us assume that there are two agents $j = 1, 2$, characterized by the pair $(R^j, \mathbb{E}^j v)$. If in state of the world $s^* \in S$ the conditions $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^1) < 0$ and $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^2) > 0$ hold, then the two agents converge in their beliefs in this state of the world.*

Proof. Without loss of generality, let us assume that $\mathbb{E}^1 v > \mathbb{E}^2 v$. For the condition $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^1) < 0$ to be satisfied, it is necessary that $v_{s^*} < R^1$. Proposition 1 states that in this case $\Delta_1(s^*) < 0$. Analogously, the second condition $(\mathbb{E}^1 v - \mathbb{E}^2 v)(v_{s^*} - R^2) < 0$ holds when $v_{s^*} < R^2$, which further implies that $\Delta_2(s^*) < 0$. That is, two agents $j = 1, 2$ update in different directions and the expected posterior beliefs are closer to each other than the priors are. Both conditions from Definition 3 are satisfied and the agents, indeed, converge in their beliefs in the state s^*. □

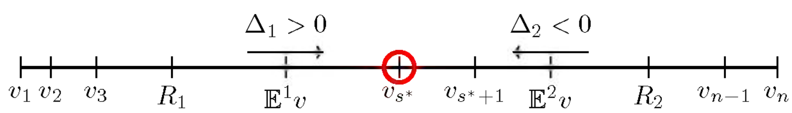

Figure 2.3: Illustration of the situation when the agents' beliefs converge to each other and to the true value of the new policy.

The situation when two agents converge in their beliefs occurs when agents have different prior expectations of the new policy and when they value the status quo differently. In addition, for both agents, it has to hold that their prior expected value from the new policy $\mathbb{E}^j v$ and valuation of the status quo R^j for the same agent j are close to each other, i.e.

both $\mathbb{E}^j v$ and R^j are smaller or bigger than v_{s^*}. This situation is illustrated in Figure 2.3.

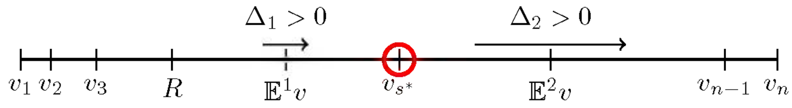

Figure 2.4: Illustration of the situation when the agents diverge in their belief updated in the same direction.

Until this point we have considered only the polarization when two agents are updating in opposite directions. However, as Figure 2.4 illustrates, there is also a possibility that posterior expected values of two agents are further away from each other than their prior expected values, while both agents update in the same direction.

Definition 4. We say that two agents $j = 1, 2$, characterized by the pair $(R^j, \mathbb{E}^j v)$ and choosing between actions $i = \{1, 2\}$, *diverge in their belief updated in the same direction* when in the state $s^* \in S$ the following two conditions are satisfied

1. $|\mathbb{E}_i^1[\mathbb{E}(v|i)|s^*] - \mathbb{E}_i^2[\mathbb{E}(v|i)|s^*]| > |\mathbb{E}^1 v - \mathbb{E}^2 v|$,
2. $\Delta_1(s^*) \cdot \Delta_2(s^*) > 0$

In Section 2.5, we consider an example which demonstrates that such divergence can occur, but this situation is too complex to be studied analytically.

2.4 Over-optimism and Polarization: Intuition and Implications

In order to understand in detail implications of the previously stated results and to draw a connection with behavioral phenomena such as over-optimism, we now focus on the case with three states. We assume that a rationally inattentive agent is choosing between a new policy that takes values $v_1 < v_2 < v_3$ in the states of the world $s = 1, 2, 3$, respectively; and keeping the status quo that has a payoff $v_1 < R < v_3$, independently of the realized state of the world. The agent has a prior expectation of the value of the new policy $\mathbb{E}v = v_1 g_1 + v_2 g_2 + v_3 g_3$.

In Proposition 1 we have shown that for fixed state s^*, the sign of the change in the mean of beliefs $\Delta(s^*) = \mathbb{E}_i[\mathbb{E}(v|i)|s^*] - \mathbb{E}v$ is the same as the sign of $v_{s^*} - R$. When we consider the true realization of the state to be $s^* = 1$ (the worst payoff of the new policy), the agent on average shifts her belief about the value of the new policy down ($\Delta(s^* = 1) < 0$ because $v_1 - R < 0$), for any $\mathbb{E}v$ and R that are inside the interval (v_1, v_3). There is no surprise here: the value of the new policy is the lowest possible v_1 and the agent on average shifts her expectation of this option's payoff down, towards the true value. Similarly, when $s^* = 3$, implementing the new policy would lead towards the highest possible value v_3 and the agent correctly shifts the expected posterior belief closer to v_3 (because $\Delta(s^* = 3) > 0$).

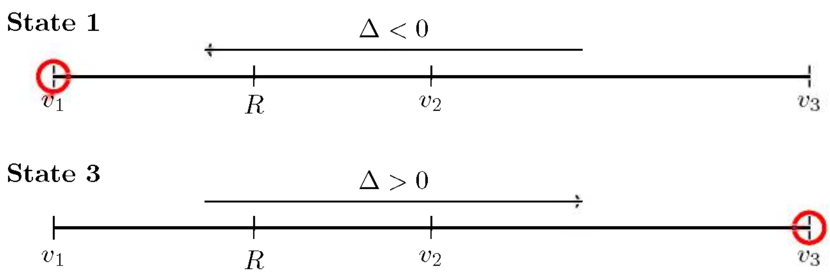

Figure 2.5: Updating when extreme states are realized. The true state is highlighted by the red circle.

Updating in the opposite direction from the realized value and the state pooling effect

A more interesting situation occurs when intermediate state $s^* = 2$ is realized. First, without loss of generality, we assume that prior belief about the new policy is such that $\mathbb{E}v < v_2$ and we fix it. We consider different possible valuations of the status quo R. If $R < v_2$ holds, then $\Delta(s^* = 2) > 0$, that is, after receiving the signal the agent updates her average posterior belief about the payoff of the new policy towards the true realized payoff v_2. However, when $R > v_2$ then $\Delta(s^* = 2) < 0$ meaning that the agent updates her expected belief to the left, i.e. away from the true payoff of the new policy. Note that this is not possible with Bayesian updating and exogenous Gaussian signals. Both these cases, when $s^* = 2$ is realized, are depicted in Figure 2.6. In these scenarios, the agent is rather pessimistic about the new policy $\mathbb{E}v < v_2$. In the first case, when $R < v_2$, the agent on average understands that the impact of the new policy is beneficial. The reason is that she knows that keeping the status quo would lead to a relatively bad outcome and thus when the realized value of the risky option is relatively high, she correctly increases her expected belief about the probabilities that the new policy can lead to better outcomes (v_2 and v_3).

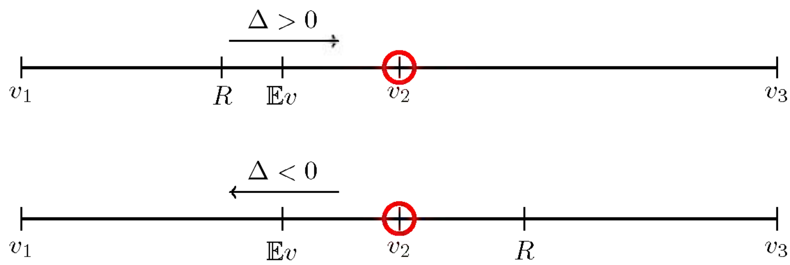

Figure 2.6: Updating for the changing status quo when state $s^* = 2$ is realized.

When $R > v_2$ the agent shifts her expectation of the new policy down, closer to the outcome v_1, away from the true payoff v_2, i.e. she updates in the *opposite direction from the realized value.* One could expect that this result is just a consequence of confirmatory learning. However, we would like to emphasize that in the problem described in this paper, the updating in the opposite direction from the true payoff is a consequence of a different mechanism. Specifically, the agent chooses between the new policy and preserving the status quo. Initially she expects the payoff of the new policy not to be very high ($\mathbb{E}v \in (v_1, v_2)$) and at the same time she perceives the payoff of the status quo as quite high ($R \in (v_2, v_3)$). The agent would prefer to choose the status quo in the realizations of the state $s^* = 1, 2$. Hence, to some extent, she acquires information that would allow her to disentangle whether state 3 is realized. She acquires information, and on average she understands that the realization of the state is indeed not $s^* = 3$, but since, to some extent, she is not interested which one of the other two states is realized exactly, both her posterior probabilities of states 1 and 2 rise. This is the *state pooling effect* mentioned in Section 2.3.4, that is, the agent endogenously pools states into categories. In this example, one category is composed of states 1 and 2; and the second category of state 3. Consequently, the direction of updating of the expected belief about the value of the risky option depends on the category to which the realized state belongs. This may result in the presense of updating in the opposite direction from the realized value of the risky option.

Symmetry, over-optimism and over-pessimism

In the example presented in this section so far we have emphasized the role of the status quo, which is not usually considered in the papers studying polarization, by keeping the prior expected belief $\mathbb{E}v$ fixed and varying the value of the status quo R. The analysis, however, could also be done for the fixed R and varying prior expected value of the new policy $\mathbb{E}v$. The whole effect works symmetrically; that is, in the previously discussed example with $s^* = 2$ the agent would also be updating in the opposite direction from the realized value when the R and $\mathbb{E}v$ are interchanged. Let us consider two situations where $s^* = 2$: the first with $\mathbb{E}v < v_2 < R$ and the second with $R < v_2 < \mathbb{E}v$. In the former situation the agent has a low prior expected value from the new policy and then updates towards v_1. In the latter situation the prior expected value is quite high and then it is updated upwards towards v_3. Stated differently, in the situation when $\mathbb{E}v < v_2 < R$, the agent is pessimistic about the new policy and consequently becomes even more pessimistic. In the second case, when $R < v_2 < \mathbb{E}v$, the opposite is true. The agent is optimistic and becomes over-optimistic about the outcome of the new policy.

Polarization

Suppose that there are two types of agents that differ only in how they value the current situation R^j, $j = 1, 2$. It is often the case that different people do not necessarily need to have different expectations about future policy, but they disagree about the favorability of the current policy. This is especially common for disputes connected with globalization, migration, robotization, climate change etc. It might be possible that those who currently benefit from the current situation and those who, for instance, lost their jobs due to globalization would have different valuations of the status quo. Let us assume that the first group (blue in Figure 2.7) opposes the current policy and the second group (red in Figure 2.7) benefits highly from the current situation , i.e. $R^1 < v_2 < R^2$. According to Proposition 1, the members of group 1 on average update their belief up ($\Delta_1(s^* = 2) > 0$) and the members of group 2 on average update down ($\Delta_2(s^* = 2) < 0$). This situation is illustrated in Figure 2.7. The prior expected belief is the same for both groups (top figure). The lower figure illustrates the posterior expected beliefs for groups 1 and 2. We can observe that posterior expected beliefs for these two groups move further apart from each other, which documents the polarization situation created only by the difference in evaluation of the current policy.

Agent not acquiring any information

In all our previously stated results, we assume that the agent is in the learning region.

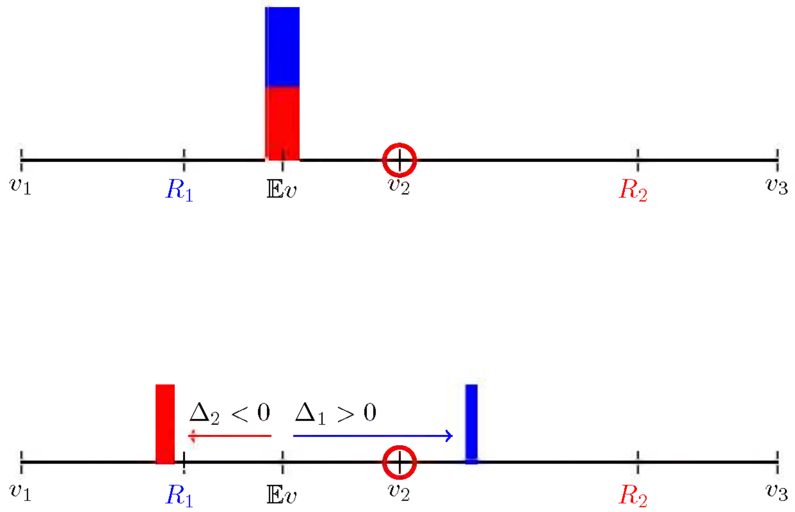

Figure 2.7: Polarization in the state $s^* = 2$ of the two groups with different status quo. Blue and red bars represent two groups of agents that differ in their valuations of the status quo, respectively. The fact that the bars are narrower at the lower figure illustrates that both groups of agents decrease their uncertainty; however, they might choose to reduce the uncertainty by different levels.

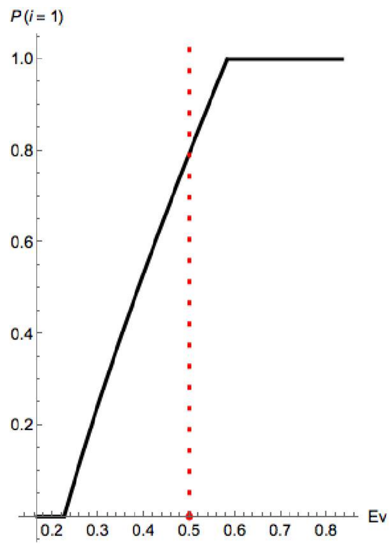

Figure 2.8: $\mathcal{P}(i = 1)$ as function of $\mathbb{E}v$, $\lambda = \frac{1}{4}$, $R = \frac{3}{8}$.

However, we know that the agent is not acquiring any information when $P(i = 1) = \{0, 1\}$. Because the agent can choose only from two options, there exist only the two following cases when the agent is not acquiring any information. The first case is when the agent's valuation of the status quo is in close proximity either to the lowest possible payoff of the new policy v_1 or to the highest possible outcome from the new policy v_3. Due to the aforementioned symmetry between $\mathbb{E}v$ and R, the second case when the agent is not acquiring any information corresponds to the situation when her prior expected belief is in close proximity to v_1 and v_3, i.e. the agent assigns relatively high prior probability to one of the extreme states. The plot of $P(i = 1)$ for varying $\mathbb{E}v$ and fixed λ and R is depicted in Figure 2.8. This behavior is expectable. For instance, when the agent a priori believes that the new policy is extremely good, while acquiring the information is costly, she would choose not to obtain any information. How the non-learning regions change with the varying cost of information is studied in the following section.

2.5 Comparative statics

Previous sections investigate the conditions determining the direction of the mean belief updating and when the polarization of agents occurs. In this section we explore the influence of the model parameters (marginal cost of information, value of status quo) on the magnitude of the change in the mean of beliefs. Specifically, we are interested in the following questions. How much does the expected posterior belief about the value of the new policy differ from the prior expected value? What is the role of the cost of information? Does the model predict that the agents become more polarized in a situation with a higher marginal cost of information? Does the actual valuation of the status quo have an influence on the value of $\Delta(s^*)$ or does it have an influence only on whether the agent is updating towards or away from the realized value of the new policy? We answer these questions using numerical solutions. Therefore, we take advantage of the example with three states and two actions. This problem is a simple benchmark and its solution exhibits the basic features of solutions to the problems with n states and 2 actions. The solution we analyze in this section is a symbolic solution.

In the scenarios under consideration we use several different values of the status quo R and of marginal costs of information λ. All parameter values are summarized in Table 2.1.

v_1	v_2	v_3	g_1	g_2	g_3	R_1	R_2	λ_1	λ_2	λ_3	λ_4
0	$\frac{1}{2}$	1	$g \in (0, \frac{2}{3})$	$\frac{1}{3}$	$\frac{2}{3} - g$	$\frac{3}{8}$	$\frac{5}{8}$	$\frac{1}{8}$	$\frac{1}{4}$	$\frac{3}{8}$	1

Table 2.1: Parameters used in this section

Note that keeping prior probability of state 2, g_2, fixed, we can vary prior probability of state 1, g only between $(0, \frac{2}{3})$. Also, $\mathbb{E}v$ can vary only from $\frac{1}{6}$ to $\frac{5}{6}$. To solve the problem (2.1)-(2.4) it is necessary to find the unconditional probabilities $\mathcal{P}(i = 1)$ and $\mathcal{P}(i = 0)$, that we then use for finding the conditional probabilities from Lemma 1. First, for a given set of parameters, the unconditional probability $\mathcal{P}(i = 1)$ as a function of $\mathbb{E}v$ for different values of λ is shown in Figure 2.9.

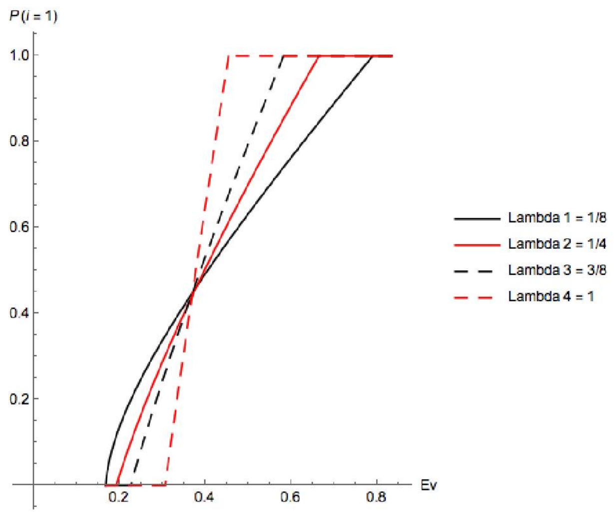

Figure 2.9: $\mathcal{P}(i = 1)$ as function of $\mathbb{E}v$ for different λ, R_1

For $\mathbb{E}v$ close to $\frac{1}{6}$ and $\frac{5}{6}$, the agent does not process any information and chooses with certainty the status quo and the new policy, respectively. With increasing marginal cost of information, the area in which she chooses with certainty grows. In the middle region, the agent acquires information, and the unconditional probability of selecting the new policy is an increasing function of the prior expected value of the new policy. With an increase in marginal cost of information λ, the small changes in $\mathbb{E}v$ translate into bigger changes in $\mathcal{P}(i = 1)$.

In order to observe how the change in the mean of beliefs is influenced by the parameters, see Figure 2.10, which depicts $\mathbb{E}_i[\mathbb{E}(v|i)|s^* = 2]$ as a function of $\mathbb{E}v$ for different levels of R and λ. In line with Proposition 1, different R change the direction of updating. Moreover, for this example, the role of the marginal cost of information is clear from this figure. The lower the marginal cost of information ($\lambda_2 < \lambda_1$), the further away the prior expected values are from the posterior expected value of the new policy. This is also presented by the fact that the agent is learning even for the prior beliefs, where she was not acquiring information for λ_1. Therefore, in our example, when the cost of information

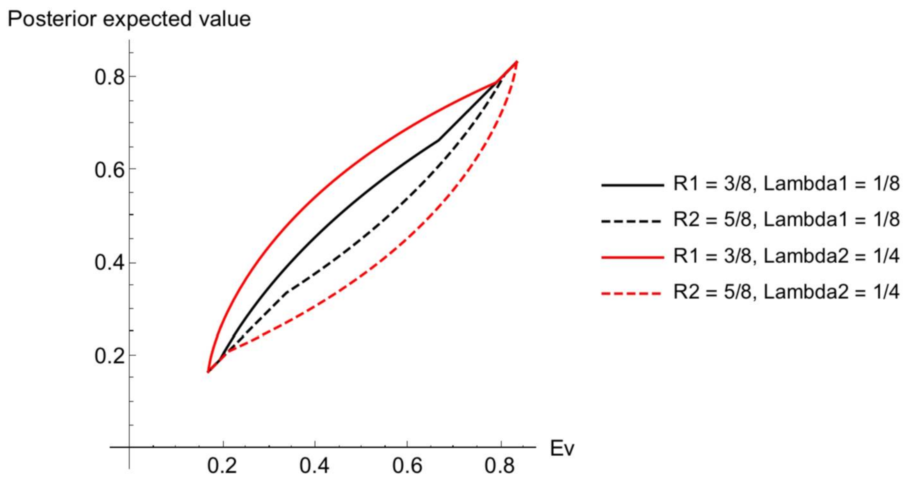

Figure 2.10: $\mathbb{E}_i[\mathbb{E}(v|i)|s^* = 2]$ as a function of $\mathbb{E}v$ for different levels of R and λ. The solid lines are the case with R_1 and dashed with R_2. Black corresponds to cases with λ_1 and red is used for λ_2.

is lower, the polarization of agents is more severe.

Another perspective on how the change in the mean of beliefs is influenced by the parameters is provided in Figure 2.11, which directly depicts $\Delta(s^* = 2)$ as a function of $\mathbb{E}v$ for $R_1 = 3/8$ and $\lambda_2 = 1/4$. The figure corresponds to the situation when the agent is updating to the right, towards v_3. The red region indicates the region where the agent is updating away from the realized value of the new policy v_2. An interesting insight is that since the maximal value of $\Delta(s^* = 2)$ is achieved for prior beliefs, which are close to the payoff associated with the true state $v_{s^*=2} = 1/2$, it suggests that someone who is updating towards the realized value of the new policy can move her belief from lower than v_2 to higher than v_2. Moreover, we observe that the more optimistic the agent is about the new policy, when she updates in the direction of v_3, the less she updates (see the decreasing part of Figure 2.11). This is not surprising in this example, because the g_2 is fixed and $\sum_{s=1}^{3} g_s = 1$.

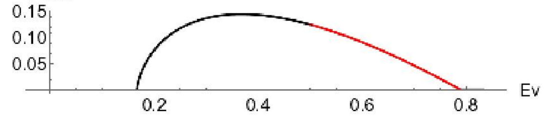

Figure 2.11: $\Delta(s^* = 2)$ as a function of $\mathbb{E}v$ for $R_1 = 3/8$ and $\lambda_2 = 1/4$. The red area depicts the region of updating in the opposite direction from the realized value.

This figure also provides an example of when two agents diverge in beliefs while they update in the same direction. We remind the reader of the illustrative figure of the situation in Figure 2.4, where we assume that the agents' valuations of the current policy are the same. Thus, since two agents differ only in their prior expectations about the new policy, it is sufficient to look at how a single agent's change in the mean of beliefs $\Delta(s^* = 2)$ depends on $\mathbb{E}v$. We are interested in finding two prior expected beliefs for which there is divergence of posterior beliefs. For that we need to find two points such that delta for the left point is lower than that for the right point. In our example the red part of the plot is a decreasing function. This means that, in our example, two agents updating in the same direction with the same valuation of the status quo might diverge in their opinions only when they are updating correctly. However, at the black part of the plot it is easy to find two points at which the agents diverge in their opinions.

2.6 Experimental design

Our theoretical results show that the agent-specific value of the status quo determines the information structure selected. As a consequence, two agents with different values of the status quo might become polarized in expectations. In particular, the rationally inattentive agent chooses to learn whether the outcome of the new policy is better or worse than the status quo, and not to learn the exact state-dependent outcome of the new policy. We have denoted this information strategy as state pooling behavior.

The main results of the model rely on several assumptions about a decision maker's preferences (risk neutrality), ability to estimate probabilities (by correctly updating beliefs), and motivation (information has a purely instrumental value). The experimental literature reports a large amount of evidence that casts doubts on human ability to perform these tasks as accurately as the theory suggests, and highlights that belief divergence could be mitigated or enhanced by human biases. We are interested in testing whether belief divergence in expectations, the main result of the model, can occur in a lab setting and whether behavioral components enhance or mitigate its magnitude.

In the following sections of the paper we investigate whether our normative model is also accurate in describing human behavior. We do so by running a lab experiment in which participants are allowed to collect information before making choices under uncertainty.

We collect actions and beliefs separately, and combine them to compute a cardinal indicator for beliefs divergence and to compare human behavior and theoretical predictions. Our design allows us to test the three most important predictions of the model: 1) the status quo influences the evaluation of and choice between information sources, 2) agents display state pooling behavior and disregard action-irrelevant information, and 3) beliefs become polarized in expectations.

2.6.1 Overview of the experimental design

The experiment comprises of four tasks and a final questionnaire. In the first and second tasks, subjects face a binary choice between the Opaque box (*risky action*), which contains a single "color ball", the value of which depends on the unknown color, and the Transparent box (*safe action*), containing a single ball whose value is known. The color ball is randomly drawn from a box containing three balls (*states*) with different colors (red, yellow, blue) with uniform probability of being selected. The two parts differ in the way we provide interim information about the color ball. In task 1 four possible advisors, (representing degenerate signal structures), are evaluated separately and subjects report their willingness to accept (Becker-DeGroot-Marscha method) to renounce to the signal. In each trial of task 2 only two advisors are displayed and the subject makes a binary choice between them. We ensure incentive compatibility by paying subjects for a single decision randomly selected from the entire experiment. Subjects never receive feedback about their decisions until the very end of the experiment. Each subject participates in all of the following tasks, in the order listed below. In tasks 3 and 4 we elicit unconditional and conditional beliefs for different advisors, assigned exogenously.

2.6.2 Task 1 - Colorblind advisor game

In each round of task 1 subjects i) choose an action contingent on the advisor and signal received (Figure 2.12, left) and then ii) indicate for each advisor the willingness to accept to renounce to its signal (Figure 2.12, right).

Subjects play 10 rounds with the same four advisors and different lottery return values. Three of the advisors in this game (named Red, Yellow, and Blue) are described as colorblind to all colors except the subject's own. Advisors are able to observe the color ball and report truthfully whether it matches her own name's color. For example, the Red

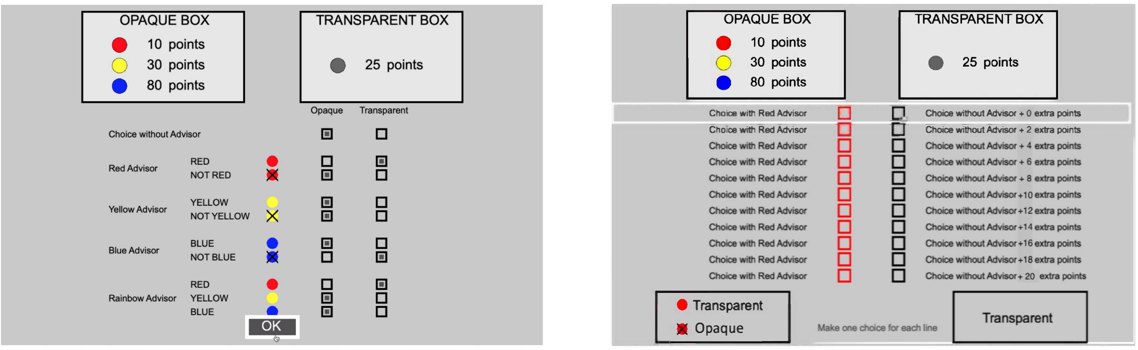

Task 1, Screen 1: Action choice — Task 2, Screen 2: WTA for each advisor

Figure 2.12: Task 1: Colorblind advisor game. Left: Subjects choose an action (box) contingent on the advisor and signal received. The possible values of each action are indicated on the top of the screen. Each state (ball color) is equally likely to occur. Right: Subjects indicate for each advisor the willingness to accept to renounce to its signal in a series of binary choices (BDM method). At most one switch is allowed. Action choices selected in the previous stage are reported on the bottom of the screen.

advisor returns a signal RED or NOT RED, which is easy to interpret. The fourth advisor is named Rainbow and reports every color accurately. This advisor provides a benchmark for the valuation of information without uncertainty. In each round, a subject chooses which box she would pick in each hypothetical advisor/answer scenario, as well as if she did not have access to any advisor (strategy method). Then, the subjects fill a multiple choice list for each of the four advisors, choosing between pairs of options: "Choice with the X Advisor" (X is replaced with the advisor's name) or "Choice without Advisor + w extra points", for w between 0 and 20 points, in 2 points intervals. The value w at which a subject i switches from preferring the former to the latter option reveals her valuation w_I^i, of the information structure I. Since each line counts as a separate decision, one of which might be randomly drawn for payment, truthful revelation is strictly optimal. We constrain subjects to have at most one switching point. If one round from this part is selected for the bonus payment, subjects receive the \$15 bonus with the percentage probability equal to the number of points that she collected in that round.

2.6.3 Task 2 - Imprecise advisor game

In each round of task 2 subjects i) choose one advisor between the two options available (Figure 2.13, left) and then ii) indicate the signal-contingent action for each signal (Figure 2.13, right).

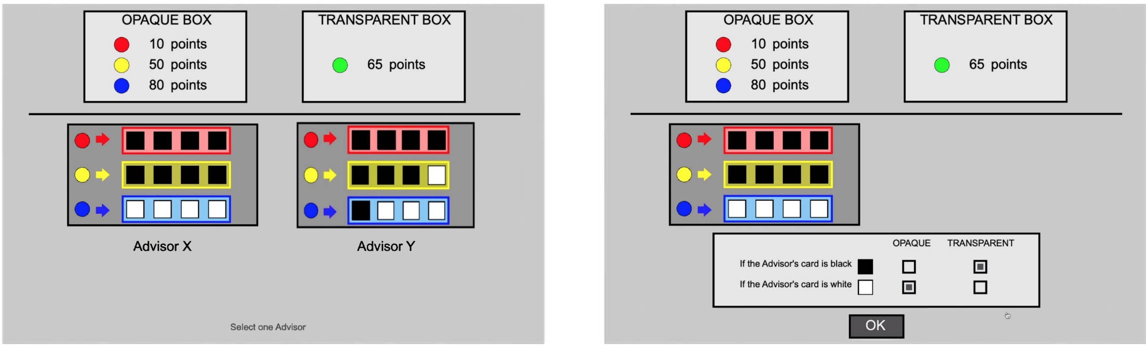

Task 1, Screen 1: Advisor choice | Task 2, Screen 2: Action choice

Figure 2.13: Task 2: Imprecise advisor game. Left: Subjects choose one signal structure (advisor) between the two options available. Each advisor is a triplet of state-contingent signal probabilities. Right: Subjects indicate the signal-contingent action for each signal (strategy method).

Subjects play 40 rounds with different pairs of advisors and values for the ball in the Transparent box. This ball can take two values (30 and 65 points). The values for the balls in the Opaque box are unchanged during the task (10, 50, and 80 points, uniform probability of being drawn). Each round comprises two parts. In the first part, subjects observe a pair of advisors and make a binary choice to select which advisor they want to consult. In the second part, only the selected advisor is consulted, a signal-contingent binary choice is implemented (as in task 1, part 1), and the participant chooses one box based on the signal received. Each advisor is defined as a triplet of decks of cards, corresponding to the triplet of conditional probabilities of providing a binary signal. Each deck is associated with one of the possible colors for the balls in the Opaque box, and contains a combination of four cards. Each card can be black or white; these colors are randomized and do not convey any intrinsic message. This means that the signal-contingent choice in the second part of the round requires the subjects to analyze every advisor separately, reducing the concern regarding inertia.

The 40 rounds are designed as a combination of 20 advisor pairs and two values for the ball in the Transparent box (safe option). The advisors are selected in order to examine preference over sources of information and formulate predictions about the effect of the safe option on information collection and posterior beliefs. In particular, 11 pairs of advisors out of 20 are designed such that a Bayesian agent would pick different advisors by changing the safe option. If one round from this part is selected for the bonus payment, subjects receive the $15 bonus with the percentage probability equal to the number of points that she collected in that round.

2.6.4 Task 3 - Card color prediction game

In each round of task 3 subjects indicate the likelihood of observing each signal for a given advisor (Figure 2.14, left). We elicit the subjects' signal probability beliefs for each of the 20 advisors using a single slider with sensibility to the unitary percentage level. Each round contains a single advisor from those used in task 2 and subjects are asked to report the probability of a black or white card being shown. We incentivize accurate and truthful reporting by using the quadratic loss scoring rule with two states (card colors). If one round from this part is selected for the bonus payment, the computer randomly determines the state and realized signal, and subjects receive the $15 bonus with the percentage probability determined by the quadratic loss scoring rule.

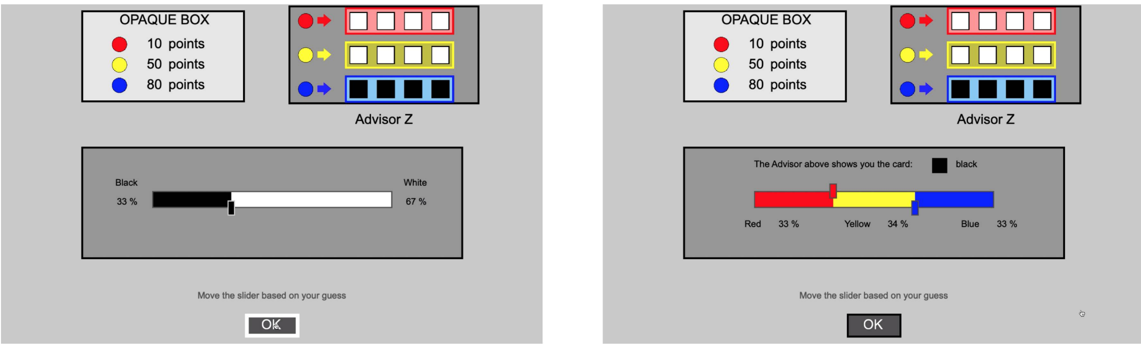

Task 3: Beliefs over signal likelihood

Task 4: Beliefs over state likelihood

Figure 2.14: Left: Task 3 (Card color prediction game). Subjects indicate the likelihood of observing each signal (card color) for the given advisor. Right: Task 4 (Ball color prediction game). Subjects indicate the likelihood of each state (ball color) given an advisor and signal. In both tasks subjects move the slider(s) and receive a number of points according to the quadratic loss scoring rule described in the instructions.

2.6.5 Task 4 - Ball color prediction game

In each round of task 4 subjects indicate the likelihood of each state given an advisor and signal (Figure 2.14, right). We elicit the subjects' posterior probability beliefs for each of the 20 advisors and for each possible signal realization, using a double slider with sensibility to the unitary percentage level. Each round contains a single advisor from those used in task 2 and one realized signal (black or white card). The subject is asked to report the probability of a red, yellow, or blue ball being in the Opaque box after observing the card color. We incentivize accurate and truthful reporting by using the quadratic loss scoring rule with three states (ball colors). If one round from this part is selected for the bonus payment, the computer randomly determines the state and realized signal, and subjects receive the $15 bonus with the percentage probability determined by the quadratic loss scoring rule.

2.6.6 Questionnaire

The final part of the experiment is a questionnaire designed to collect demographic variables (including field of study and familiarity with Bayes' rule), psychological measures (Life Orientation Test - Revisited, LOT-R hereafter), risk attitude (Holt-Laury risk elicitation method with multiple price list, Holt and Laury 2002, HL hereafter), cognitive ability (five questions from the Raven Progressive Matrices Test, Raven hereafter), as well as questions on the subjects' strategy in the first and second task.

2.6.7 Procedure

The experiment was run in CELSS (Columbia Experimental Laboratory of Social Sciences, Columbia University, New York, USA) between August and September 2019. The experiment was coded in MATLAB (Release 2018b) using Psychotoolbox 3 (Psychophysics Toolbox Version 3). Eighty-five volunteers were recruited using the platform ORSEE (Online Recruitment System for Economic Experiments) and were naive to the main purpose of the study. All subjects provided written, informed consent. The whole experiment took on average 85 minutes, including instructions and payment. On completion of the experiment, the subjects received payment in cash according to task performance. Each subject received a $10 show-up fee, and played for a bonus prize of $15. In addition, a subject could earn between $0.10 and $4 in the risk elicitation task and $0.50 for each question of the cognitive test, up to $5. The average payment was $25.

2.6.8 Subject understanding

Instructions were provided both on the computer screen, as slides that can be browsed by each subject at the desired pace, and as paper printout. The two versions of the instructions contained the same information verbatim. Subjects were required to answer correctly all the multiple-choice questions of the comprehension test to check understanding of the instructions before proceeding with every section of the experiment. The number of questions ranged from two to four for every section, and subjects received a one-minute timeout before having a new attempt. Subjects were initially informed about the payment structure, the no-deception policy of the laboratory, and that choices in one section of the experiment did not affect any other section, nor the questionnaire. A small number of subjects were recruited for each laboratory session (6 on average) in order to facilitate clarification of questions during the experiment.

2.7 Experimental investigation

This section aims to clarify the connection between the theoretical setup discussed in Section 2.3 and the experimental design introduced in Section 2.6. Readers who are familiar with the model can skip the following subsection and proceed directly to the hypothesis (Section 2.7.3).

2.7.1 Formal setup

We consider a setting with three possible states and two actions that generate state-contingent payoffs. The actions represent two policies: the current policy (the status quo), whose return R is known and independent of the state, and a new policy, whose return $v_s \in \mathbb{R}$ is uncertain. The state of the world $s \in S = \{r, y, b\}$ is represented by a color, associated with the deterministic return for the uncertain policy: r (red, low return), y (yellow, intermediate return), or b (blue, high return), with $0 < v_r < v_y < v_b < 100$ and $v_r < R < v_b$. An agent with a correct uniform prior belief $P(s) = \frac{1}{3}$, $\forall s$ observes an informative signal about the state and selects one of the two policies. The return $V \in \{\{v_s\}_S, R\}$ of own choice depends on the selected action and the realized states, and represents the probability of receiving a fixed prize k (\$15 in our laboratory experiment).

Information is valuable because it informs the subsequent binary choice between policies. We let $\sigma \in \{0, 1\}$ denote the realization of a stochastic signal that the subject may observe (note that in the case of degenerate probabilities the signal is deterministic). Since we have three states and two possible signal realizations, a signal structure is a triplet of state-dependent probabilities $P_I(\sigma = 1|s)$. We will refer to such a triplet I as an information source or advisor. Notice that even though the three states are equally likely, the two signals need not be equally likely.

The Bayesian agent represents a natural benchmark to consider the objective value of information in this environment. Let w_I denote the bonus that renders the agent indifferent between playing the game without additional signals[10] (but receiving additional w_I "tickets") and playing the game with the signal structure I. Since there are only two possible levels of payoff a subject can obtain (k if she wins the prize, or 0 if she does not), and intermediate rewards π are expressed in probability points, this measure of valuation is independent of the curvature of a subject's utility function for money. Subjective posterior beliefs $P_I(s|\sigma)$ and signal probability assessments $P_I(\sigma)$ are sufficient to obtain the valuation of information structure I for any agent who reduces compound lotteries into simple lotteries, has preferences that depend only on final outcomes,[11] and strictly prefers obtaining the prize $k > 0$ to not obtaining it, even if she is not Bayesian.

Given a binary choice environment $\{\{v_s\}_S, R\}$, and a signal structure I, the decision maker implements an updating rule $P(\sigma)$ (guess about the state of the world) and makes a binary decision. The valuation of the information structure I is given by a chosen lottery with value V and by the observed signal

$$w_I = EV[\underbrace{V|\text{signal observed}}_{V_1}] - EV[\underbrace{V|\text{no signal}}_{V_0}] \tag{2.6}$$

For notational reasons, we denote as V_0 the value of the action chosen without observing any signal and as V_1 the value of the action chosen after observing a signal realization. Similarly, we normalize the notation such that $\sigma = 0$ is the signal realization that leaves the chosen action unchanged, compared to the scenario without additional information. In contrast, $\sigma = 1$ is the choice-reverting signal.[12] Our design rules out any effect of risk preferences as the lotteries' returns are expressed in probability points. We can rewrite the valuation in equation 2.6 by taking into account that (at most) one signal realization

[10] Playing without any additional information is, from a theoretical perspective, equivalent to playing with a purely noisy signal. We prefer to refer to the former case for the sake of clarity.

[11] Note that a non-instrumental preference for information, e.g. preference towards certainty or preference towards negative information, would create a gap between the theoretical and observed valuation that cannot be explained by updating alone.

[12] We are implicitly assuming that such a signal exists. If both signal realizations are associated with the same final action, we can redefine the signal probabilities such that the choice-reverting signal $\sigma = 1$ occurs with probability 0 in every state.

is associated with a different choice with respect to the scenario without signal.

$$w_I = Pr(\text{signal} = 1) \cdot \Delta EV[V_{choice}|\text{signal} = 1] \tag{2.7}$$

where

$$\Delta EV[V_{choice}|\text{signal} = 1] = EV[V_1|\text{signal} = 1] - EV[V_0|\text{signal} = 1].$$

We can generalize the subjective valuation in order to include non-instrumental preference over information. A decision maker i has a subjective valuation w_I^i of the signal structure I that depends both on the instrumental value l_I and other characteristics of I, for example the type of "optimistic/pessimistic" information that it provides. We postpone further discussion about possible differences between Bayesian and subjective valuation of information to the results.

2.7.2 Bridging theory and experiment

Our experimental design allows us to estimate how agents evaluate an informative signal structure (advisor), and measure how the subjective evaluation depends on the properties of the signal structure, including instrumental value (expected improvement in the choice process) and non-instrumental properties (ease of interpretation).

The timing of the problem (as in task 2) can be summarized as follows:

1. The agent is informed of the prior $P(s) = \frac{1}{3}\ \forall s$ and the state-contingent returns $\{v_s\}_S,\ R$.
2. One state is realized, but the agent is unaware of it.
3. The agent is offered two sources of information (advisors) I_1 and I_2.
4. The agent chooses one advisor and discards the other.
5. The selected advisor observes the realized state (ball in the opaque box).
6. The selected advisor returns a binary signal, whose likelihood depends on the realized state.
7. The agent observes the realized signal.

8. The agent chooses one action (opaque or transparent box) and receives the payoff π.

9. The agent plays a lottery and receives the final prize k with probability $\frac{\pi}{100}$.

The problem presented in task 1 is similar up to a change in steps 3 and 4:

3'. The agent is offered one single source of information (advisor) I

4'. The agent indicates how much she is willing to accept to renounce to the advisor.

In the *Colorblind advisor game* (task 1), we elicit the probability w_I such that the agent is indifferent between making a choice after observing the realization of a known signal structure I, and choosing without additional signals but receiving additional w_I tickets to win the prize. In the *Imprecise advisor game* (task 2), we offer pairs of signal structures, and collect binary choices between advisors. If the valuation and choices differ from those a Bayesian expected utility maximizer would display, we would like to pinpoint the source of the deviation. For this reason, we add two control tasks to elicit a subjective signal of beliefs' realization (*Card color prediction game*, task 3) and subjective posterior beliefs (*Ball color prediction game*, task 4). We collect posteriors only after eliciting preference over advisors, so we do not nudge the subjects towards thinking about information valuation in a specific fashion.

2.7.3 Hypothesis

Our experiment allows us to test the main predictions of the model, as well as disentangle the possible factors that mitigate or enhance the results with respect to the behavior of an optimal decision maker. The first hypothesis refers to the crucial effect of the status quo on information acquisition.

Hypothesis 1. A change in the status quo (safe option) generates a reversal in the choice between advisors when such a reversal is optimal.

We tested this hypothesis by collecting choices over information structures under different values of the safe option. The optimal advisor choice would not be sufficient to generate belief polarization. The second hypothesis refers to the result that appears in the title of the paper.

Hypothesis 2. A change in the status quo (safe option) generates beliefs polarization in expectations.

We tested this hypothesis by collecting subjective beliefs after receiving a signal, in addition to the advisor choices. There are three main channels that may mitigate or enhance the results with respect to the behavior of an optimal decision maker: non-standard preferences for information structures, biased beliefs, and non-standard preferences over realized states. We can summarize the behavioral assumptions with these three hypothesis

Hypothesis 3. Subjects choose the information structure with the highest instrumental value.

Hypothesis 4. Subjects have an accurate estimate of probabilities and update own beliefs optimally after observing an informative signal.

Hypothesis 5. Subjects choose the action with the highest expected return after observing an informative signal.

Our setup allows us to test an additional hypothesis related to the general problem of information valuation by collecting detailed data about willingness to pay and binary choice between information structures. We extend traditional designs in a scenario with three states of the world and test the robustness of effects that are well-known in simpler settings. The behavioral counterparts for hypothesis 3 above can be summarized with the following alternative hypothesis.

Hypothesis 6. Subjects evaluate and choose information structures based on non-instrumental characteristics, including accuracy about the most desirable states (optimism) and ease of processing of the signals (certainty, state pooling).

2.8 Experimental results

This section contains the main results of the experimental investigation. We report aggregate choices between sources of information (Section 2.8.1) and provide evidence that (i) subjects do react to the value of the status quo as predicted by the theoretical model, meaning that they do switch between advisors providing the information about the best/worst state in accordance with their optimality. We also show that (ii) the probability of choosing an advisor increases with the instrumental value of the corresponding information structure. In addition to the instrumental value, we consider other characteristics of the information structures (Section 2.8.2), (iii) we replicate the well-known preference for certainty in our setting, and (iv) we document preference for state pooling information structures.

Possible deviations from optimality in the information collection stage may be reconciled with non-standard preferences and suboptimal action choice after observing the signal realization. We analyze subjects' actions (Section 2.8.3) and (v) we verify that actions are consistent with the optimal behavior of a risk-neutral agent. Biased perception or updating of probabilities are also possible causes of mistakes in the information collection stage. We compare the optimal beliefs about the state of the world of a Bayesian agent with the subjects' elicited beliefs (Section 2.8.4) and we report that (vi) subjects display conservatism in their beliefs, but the bias does not represent a major driver of deviation from optimality.

We combine actions and beliefs elicited in separate tasks (Section 2.8.5) and (vii) we observe belief polarization in our laboratory setting. Finally, we analyze the willingness to pay (WTP) for information structures in the first task (Section 2.8.6), where we observe that (viii) subjects display compression in their WTP and (ix) are willing to pay higher amounts for information about the most desirable state.

2.8.1 Optimal response to the change in the status quo

Figure 2.15 plots the aggregate probabilities at which subjects choose the best advisor, that is, the one that offers the information structure with higher instrumental value. The first bar displays the probability of choosing the best advisor in all trials where the best advisor exists (29 trials out of 40). For the other two bars we use 22 trials where it is optimal to switch between advisors when the value of the status quo changes. We can observe the high rate of optimal choices in all three cases. Importantly, the probabilities of choosing the best advisor for the status quo $R = 30$ and $R = 60$ are both significantly greater than 0.5, with values 0.66 and 0.70, respectively. The probability of choosing the optimal advisor when we also include trials with a dominant advisor, where the switch between the advisors does not occur, increases to 0.72. These high probabilities confirm that the subjects correctly recognize the optimal advisor and do show switching behavior

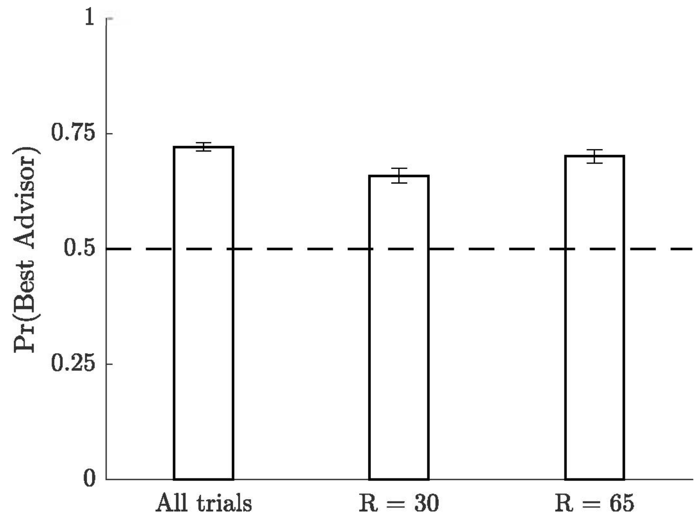

Figure 2.15: Probability of correct answers for all the participants. The first bar is for all trials where there is a dominant advisor (29 trials out of 40, 2465 observations). The remaining two bars indicate the pairs of trials with the same advisors and different status quo R, in which it is optimal to switch the choice of advisors when the status quo is changed (11 trials per bar, 935 observations).

between the advisors in line with the changes in the status quo.

Experimental Result 1. Subjects systematically react to the value of the status quo and choose the optimal advisor (information structure).

In order to take a more detailed look at the trials where the switch between the advisors is optimal, in Figure 2.16 we depict the probability of choosing the advisor (out of a pair of advisors labeled X and Y) as a function of the difference in the instrumental values of the two advisors presented. The probability of selecting Advisor X increases with the difference between the instrumental values of Advisor X and Advisor Y. All the trials except one lie in the first and the third quadrant, that is, whenever the instrumental value of Advisors X is lower than Y the probability is below or at most equal to 1/2. When the value of Advisor X is greater than Advisor Y the probability is higher than 1/2. Probabilities are increasing almost linearly and not stepwise near the zero, suggesting that subjects do not respond only to the sign of the difference in the instrumental values between the advisors, but to the actual value of the difference.

Experimental Result 2. The probability of choosing an advisor increases with the instrumental value of the corresponding information structure.

When we inspect Figure 2.16 closely, we notice that there are five trials (two of them overlap) that have a difference in instrumental values between advisors equal to -2.5 and

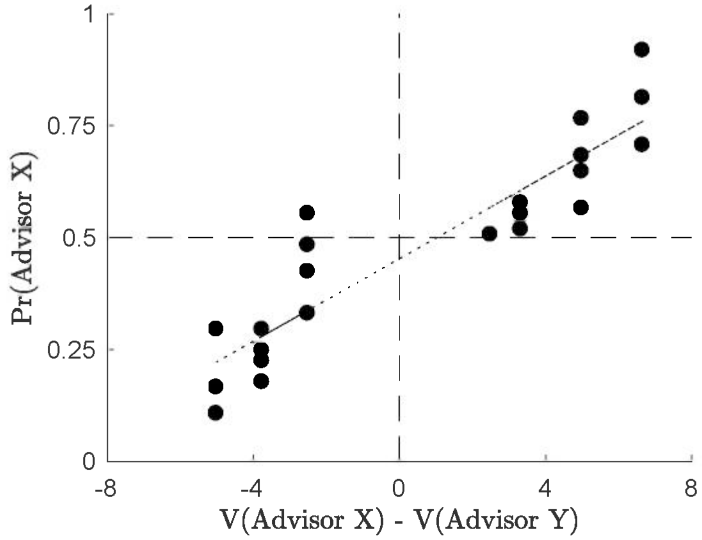

Figure 2.16: Probability of choosing Advisor X, given a pair of advisors X and Y, in those trials where a different advisor is optimal for different value of the status quo R (22 trials out of 40, 85 observations per trial). The probability of selecting Advisor X is increasing with the difference in the instrumental value of advisors and demonstrates that the accuracy of participants increases with the stakes.

their probabilities of selecting Advisor X are increasing with one trial having a probability higher than 1/2 and one very close to 1/2. How do these two trials differ from the others in that their probabilities of selecting Advisor X are very high even though it is optimal to choose Advisor Y? In both trials Advisor X offers a yes/no answer to the question whether the state is red. Such an advisor is offering the ideal state pooling information; however, it is more valuable to learn whether the blue state is happening because in both of these trials the status quo is equal to 65, so it is between payoffs from the yellow and blue states. This observation already provides us with an indication of the subjects' preference for state pooling, but in the following section we investigate this in greater detail.

2.8.2 State pooling and preference for certainty

The essential mechanism behind belief polarization of inattentive agents presented theoretically is that of state pooling behavior. Do people select the state pooling information? Before presenting the results we present the definition of the state pooling advisor (the concept of state pooling was introduced in the theoretical setting in Section 2.3.4).

Definition 5. An advisor with information structure I is a *state pooler* when it can provide a signal that generates a degenerate posterior about either state $s = r$ or $s = b$.

Additionally to the state pooling advisors we recognize another special category of sources of information, those that provide an answer to a question, “Is the state red(/yellow/blue)?" That is, the subject can learn with certainty if it is a particular state.

Definition 6. An advisor with information structure I provides certainty when there exists a state s such that every signal generates a degenerate posterior about the state being s.

Note that every advisor providing certainty about the best or the worst state, i.e., blue and red state, is a state pooling advisor. However, the opposite is not true. For instance, an advisor that gives a black signal with probability 1/2 when a state is blue and white signal otherwise is a state pooler, but does not provide certainty.

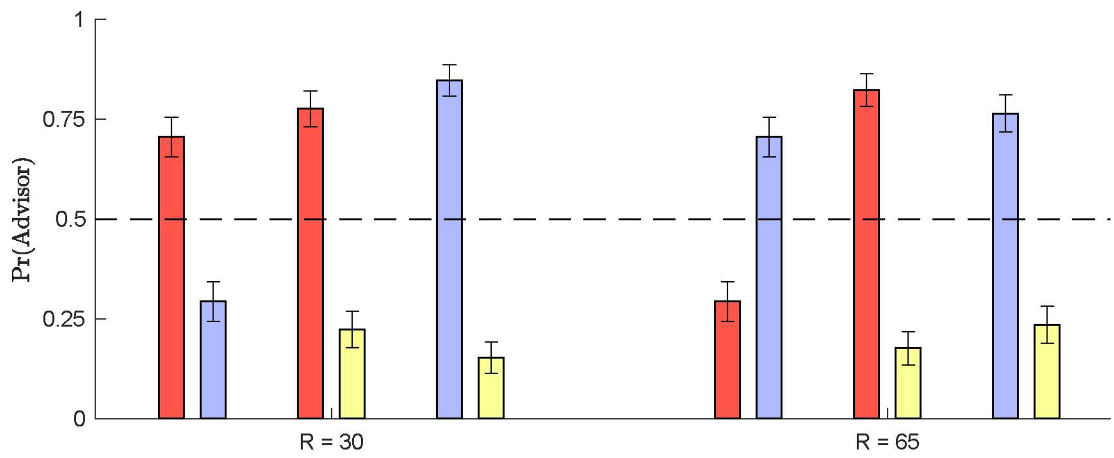

Figure 2.17: Comparison of advisors providing answer to the ideal state pooling question: “Is the state red/yellow/blue ?” for two different status quo values. The color of the bar shows which state the question is about. The figure demonstrates the state pooling behavior, and also that participants do switch between advisors when it is valuable to do so.

Figure 2.17 shows advisor choice in the trials in which both advisors provide certainty. We display separately the trials with different status quo values. When the subjects have to choose between advisors that provide certainty and are also state poolers, that is, between an advisor providing information whether the state is blue and an advisor providing information whether the state is red (first couple of bars for $R = 30$ and $R = 65$), they significantly select the former for high value of the status quo and latter when its value is low. This switch between advisors confirms our theoretically predicted state pooling effect. In particular, for a status quo value R the subject wants to learn whether the state-dependent payoff of the new policy is greater or lower than R. When subjects face a choice between a certainty state pooler and certainty advisors, they select the certainty state pooler with a probability close to 0.7.

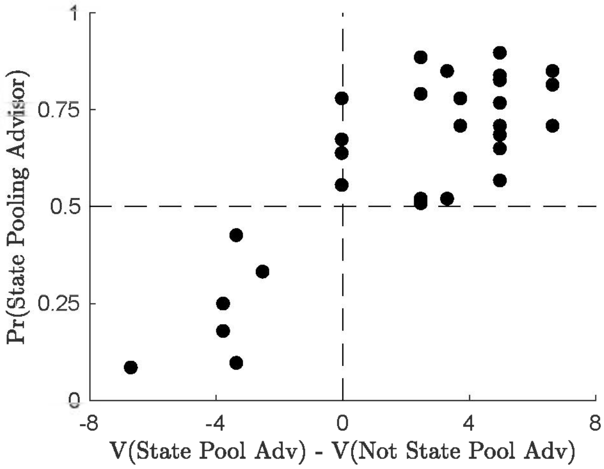

Figure 2.18: Probability of choosing the state pooling advisor for an appropriate state in trials where such a choise is possible. The participant strictly prefers to choose the state pooler when it is optimal to do so, as well as when both advisors are equally instrumentally valuable.

When we investigate how the probability of choosing the state pooling advisor depends on the instrumental value of such an advisor (see Figure 2.18), we notice that the probability is greater than 1/2 when it is optimal to select the state pooler, and below 1/2 otherwise. However, we can also notice that the probability of selecting the state pooler is increasing with the instrumental value. A particularly interesting observation is when the state pooling and non state pooling advisor both have the same instrumental value (0 on x-axis). In such a situation, subjects strictly prefer to select the state pooling advisor in comparison with a non state pooling advisor, even though it is not more informative.

How is the preference for the certainty advisor connected with the information valuation of such an advisor? Figure 2.19 plots the probability of choosing the certain advisor given the difference between the certain and uncertain advisor. Here we can see that the

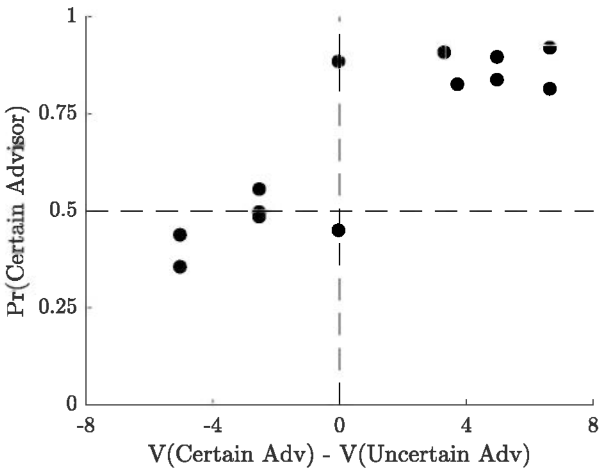

Figure 2.19: Probability of chosing a certain advisor in trials (14 trials, with 85 observations per trial) where the choice is between the certain and uncertain advisors. When the certain advisor is more valuable the probability of its being chosen is close to 1. When the certain advisor is less valuable, it is still chosen with a probability close to 1/2.

probability does not increase almost linearly with the informativeness as in the previous figures, but once it is optimal to choose the certain advisor it jumps to approximately 0.9. At the same time, also when it is optimal to choose an uncertain advisor, the probability of selecting the certain advisor is very close to 1/2.

This result highlights that subjects display a preference for state pooling and certainty advisors. In the case of a certainty advisor, we notice that the sign of the difference is enough to predict the correct choice. This result is driven by the ease of interpretation of the certainty advisor. We observe a similar pattern with the SP advisor, but the effect is mitigated by the value differences. Choices are accurate when advisors have significantly different values, whereas they are on average more imprecise when the difference is small. In order to understand these preferences and their interplay we estimate how the probability of selecting Advisor X depends on various components. Additionally to their instrumental value, we use a binary variable to indicate whether the advisor is the best in the pair, and three additional binary variables that capture whether an advisor gives certainty, state pooling, and the interaction between of the two.

The results from the regression are presented in Table 2.2. Not surprisingly, the instrumental value of the advisor w_I^{Bayes} has a significant effect on the probability of the advisor

	(1)	(2)	(3)	(4)
w_I^{Bayes}	0.210***	0.236***	0.2108***	0.253***
	(0.0346)	(0.0110)	(0.0113)	(0.0358)
Best Advisor	0.106			-0.186
	(0.143)			(0.149)
Certainty		0.1876***		0.596***
		(0.0708)		(0.154)
State Pooling			0.742***	1.041***
			(0.0757)	(0.111)
Certainty × SP			0.488***	-0.0601
			(0.0825)	(0.164)
Trials	All	All	All	All
Observations	3,400	3,400	3,400	3,400

Table 2.2: The significance levels concern the hypothesis that the coefficient is 0. Notation: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

being selected. However, both certainty and state pooling are also significant. From column (3) we can see that, when certainty is not included, the state pooling by itself is not able to capture the whole effect. When both certainty and state pooling are included, the interaction term is no longer significant. Another interesting observation is that the parameter for state pooling is twice as large as that for certainty. We can summarize the findings from this part in the following two results:

Experimental Result 3. Subjects significantly prefer state pooling advisors, even after controlling for the instrumental value of the available advisors.

Experimental Result 4. Subjects also significantly prefer advisors providing certainty, but the magnitude of this effect is smaller than the preference for state pooling advisors.

2.8.3 Action selection

Possible deviations from the optimal choice between alternative sources of information in the imprecise advisor game (task 2) can be rationalized by biases in preferences and beliefs. We collected separately action choices and subjective beliefs and we show here

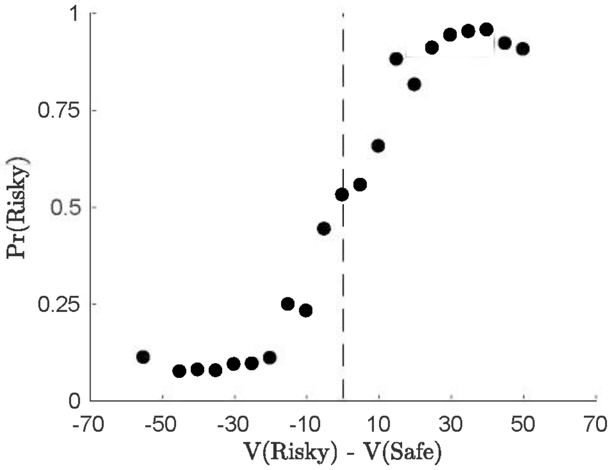

Figure 2.20: Observed probability of choosing the risky action in task 2. Trials were grouped based on the differences between the expected value from the risky option and a safe one, conditional on the realized signal. 3, 400 observations grouped in 21 bins with different sample sizes.

that participants' behavior is remarkably close to the theoretical predictions.

In task 2, given an advisor and a signal realization, participants were asked to choose one action, corresponding to the risky lottery (opaque urn) and safe option (transparent urn). A risk neutral agent that correctly calculates the expected value of the risky lottery should choose this option only if its value exceeds the status quo. Figure 2.20 shows the realized probability of selecting the risky option as a function of the difference in the EV between the actions. Trials are grouped based on the x-axis value for visualization purposes. The optimal agent would have a sharp jump in probability from 0 (when the value difference is negative) to 1. We observe in our data a smoother transition, suggesting that action probability is modulated by the cost of mistakes, similarly to our discussion in Figure 2.18 about the choice between advisors. Such a sigmoid curve is normally found in experiments involving choice under risk (Mosteller and Nogee 1951; Khaw, Li, and Woodford 2019). The indifference point appears close to the trials in which both actions have the same values, suggesting that the participants are, overall, close to risk neutrality. This outcome is consistent with the utilization of probability points to incentivize choices under uncertainty.

Experimental Result 5. Participants behave similarly to a risk neutral agent with a correct understanding of probabilities. They most likely choose the action with the highest expected value and their accuracy increases with the difference between the values of the two actions. Risk preferences do not represent a major driver of deviation from optimality in the choice between advisors.

2.8.4 Belief elicitation

A second concern is that participants are systematically biased in the computation of probabilities and this may influence the evaluation and therefore the likelihood of choice between advisors. We elicit subjective beliefs with a strategy-proof mechanism and verify that the participants hold, on average, very accurate beliefs. We are interested in eliciting both the likelihood of signal realization (given the advisor) and the posterior distribution (given the advisor and signal realization). We follow closely the design adopted by Charness, Oprea, and Yuksel (2018) and we adopt a Quadratic Scoring Rule to incentivize the accurate reporting of probabilities. This rule is incentive compatible and has been used extensively in decision making tasks with two or three possible states (Selten 1998; Costa-Gomes and Weizsäcker 2008).

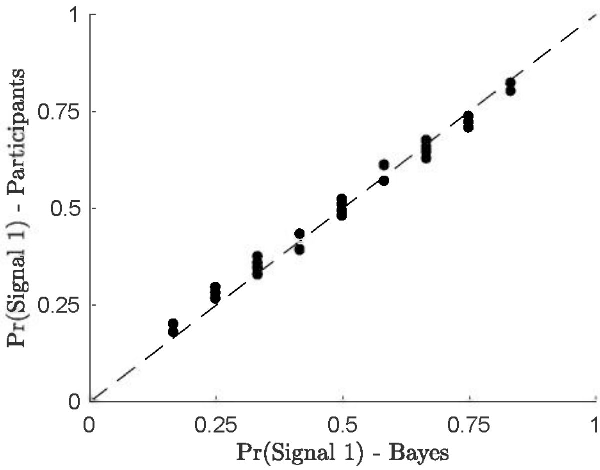

Figure 2.21: Estimated probability of receiving a signal realization in Task 3. The plot compares the average of the subjective estimates collected with the optimal estimates of a Bayesian decision maker. 1,700 observations across 20 trials (85 observations per point).

In both tasks we observe accurate probability estimates, close to the predictions of an optimal Bayesian agent. Figure 2.21 shows the subjective estimate of a signal realization (y-axis, averaged across participants) compared with the optimal estimates (x-axis).

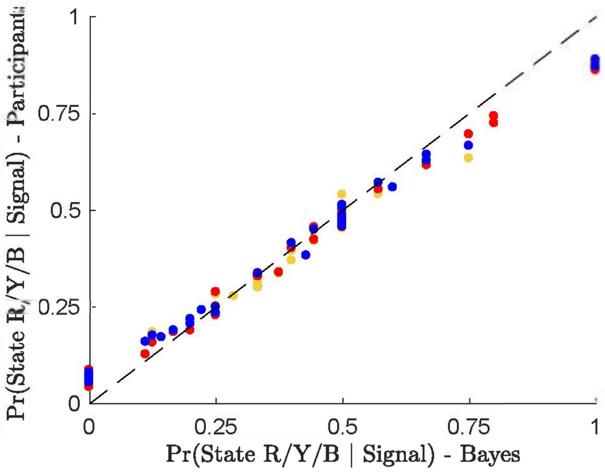

Figure 2.22: Estimated posterior probability of each state in task 4 conditional on the realized signal. Colors indicate which state was estimated (red, yellow, blue). The plot compares the average of the subjective estimates collected with the optimal estimates of a Bayesian decision maker. 20,400 observations across 40 trials (6 observations per trial, 85 observations per point in the plot).

Similarly, Figure 2.22 shows the subjective estimate of each of the three possible states in the posterior compared with the unbiased posterior, with different colors in the figure matching the state. In both plots, the 45 degree lines represent our theoretical benchmark and we can see that 1) participants are on average accurate in the estimate of probabilities, 2) we do not observe a systematic difference between estimates involving different states (i.e. we do not have evidence of motivated beliefs, Bénabou (2015)), and 3) both tasks show mild evidence of conservatism (central tendency of judgement), as is widely reported in experiments with subjective estimates (Hollingworth 1910; Anobile, Cicchini, and Burr 2012).

Experimental Result 6. Participants are on average accurate in the estimate of probabilities, both when they are asked to report signal realization beliefs and posterior probability beliefs. Biased beliefs do not represent a major driver of deviation from optimality in the choice between advisor.

2.8.5 Observed belief polarization

Our model predicts that a change in the status quo creates belief polarization because of the endogenous choice of information structures. In our experimental design this means that, given the true state, the same decision maker will have different beliefs (ex ante, before the signal realization) based on her status quo value. The experiment contains 11 pairs of trials that can be used to verify whether such polarization occurs. We combine the data collected for the binary advisor choice and signal-contingent choices (task 2) with the subjective beliefs about signal realization (task 3) and posterior distribution (task 4).

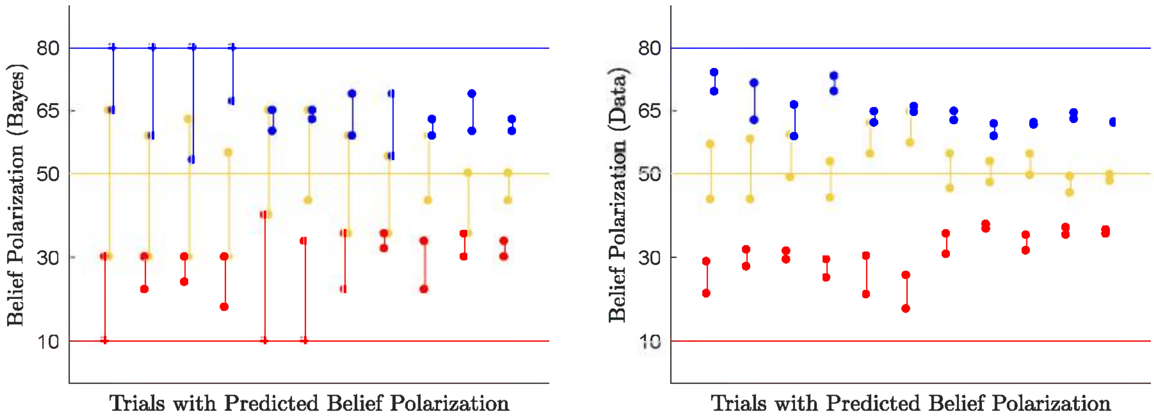

Figure 2.23: Beliefs divergence in the experiment: prediction and observations. For each of the 11 pairs of trials with predicted advisor switch (x-axis), we indicate the expected value of the risky action conditional on the state (color) and the chosen advisor (connected dots), but before the signal realization. The length of the connecting lines indicates the magnitude of belief polarization. Left: Predictions based on the optimal behavior of a Bayesian decision maker. Right: Observed polarization based on the advisor choice and elicited beliefs in tasks 2, 3, and 4. 1,870 observations across 11 pairs of trials (170 observations per pair of trials).

Figure 2.23 (left) shows the theoretical predictions for polarization of beliefs in each of the 11 pairs of trials in which agents should switch advisor. For each trial (column) and each state (color) we observe that the two status quos lead to different ex-ante expected values for the risky action. The two connected dots of the same color indicate the expected value, conditional on the state, for the two different agents, and the length of the connecting lines represents our measure of divergence. We want to compare the theoretical prediction (on the left) with the realized polarization that we observe in our data (on the right).

Experimental Result 7. Variations in the status quo value generate ex-ante belief divergence (before the signal realization, and after controlling for the true state) qualitatively analogous to the predicted ones, but with smaller magnitude.

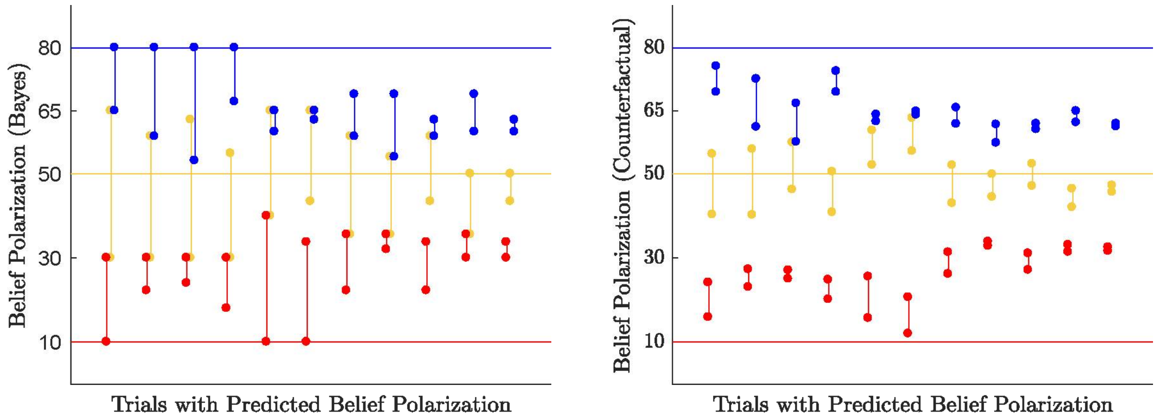

Figure 2.24: Beliefs divergence: prediction and counterfactual. Values are shown as in Figure 2.23. Left: Predictions based on the optimal behavior of a Bayesian decision maker. Right: Counterfactual analysis with optimal advisor choice and biased beliefs (estimated based on data collected in tasks 3 and 4).

We can separate the departure from the theoretical predictions into two components: advisor choice and beliefs. We discussed in Section 2.8.4 that subjective beliefs are on average remarkably close to the optimal ones. A counterfactual scenario in which agents choose the advisors as in the lab experiment but update beliefs optimally has a minor effect and increases divergence. Figure 2.23 shows the polarized beliefs in this alternative scenario. The fact that beliefs play a minor role in the departure from the theoretical predictions of belief divergence suggests that preference over sources of information is the main driver of this difference.

2.8.6 Compression effect in the willingness to pay

The Colorblind advisor game introduced in task 1 provides a different dataset that we can compare with the results from the other tasks. In each of the ten trials we collect signal-contingent actions (risky or safe options) as well as the subjective willingness to pay (WTP) in order to observe a certain signal structure. More precisely, we elicit the willingness to accept, expressed in probability points of winning the bonus, in exchange for the opportunity of playing the game without the advisor. For each of the four advisors in the game we elicit subjects' valuation of the advisor using multiple price lists, an incentive compatible implementation of the Becker-DeGroot-Marschak mechanism (BDM,

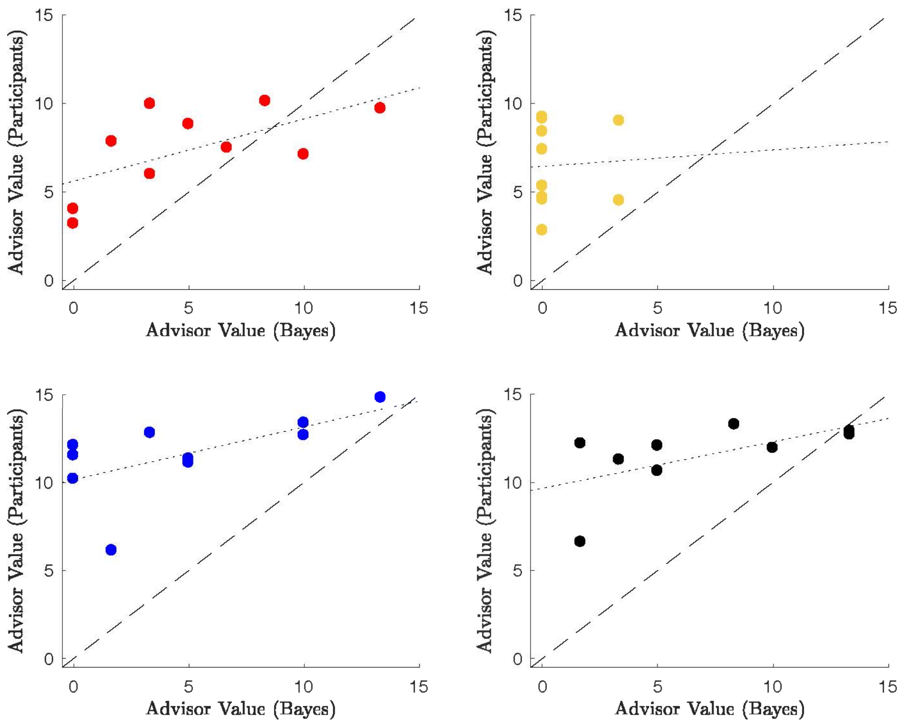

Figure 2.25: Comparison between the average subjective valuations of advisors across participants and the valuations for the optimal decision maker. The color used in each panel indicates the type of advisor: red, yellow, blue, and rainbow advisors (ordered from top-left to bottom-right). Optimal valuation (dashed lines) and linear regression estimates (dotted lines) are shown for comparison. $2,520$ observations across trials and advisor types (85 observations per point in each panel).

from Becker, DeGroot, and Marschak (1964)). The red, yellow, and blue advisor provide a binary message, whereas the rainbow advisor fully reveals the true state. Figure 22 shows, for each advisor type, the relation between subjective (averaged across participants) and theoretical advisor value (for an optimal decision maker). We notice that for all advisors the subjective evaluation tends to exceed the theoretical one (positive intercept) and there is a general positive relation between the two, with subjective values increasing with the theoretical ones, but not as much as the latter. This pattern is known as the compression effect and is well known in experiments with explicit elicitation of WTP for sources of information (e.g. Ambuehl and Li (2018)). We have several cases of advisors whose theoretical value is equal to zero (the case for most of the yellow advisors): the observation of a signal from them is not pivotal for the chosen action with respect to the decision without advisor, yet the subjects invest a significant amount of points to receive this piece of information.

The comparison of the plots of different advisors highlights a consistent pattern. The compression effect appears similarly in all four advisors, with a similar slope of the linear regression between observed and theoretical values. At the same time, intercepts are significantly different, with similar values for red and yellow advisors, and much higher levels for blue and rainbow advisors. This difference is aligned with preference for information structure biased in favor of the most desirable state (blue state).

This result is confirmed by running a simple OLS regression of the subjective advisor value using the theoretical value as regressor. Table 2.3 shows that the slope is positive but lower than one (compression effect) and the intercept is positive and significantly different from zero (analogous to the conservative probability estimates observed in tasks 3 and 4). When we allow the intercept to differ across advisors, we notice that they are not different between the red and yellow advisor, whereas the blue and rainbow advisors receive significantly higher WTPs. This result is consistent with those observed in environments with non-instrumental information (Masatlioglu, Orhun, and Raymond 2017) in which subjects display wishful thinking and desire to observe signals that are more accurate about the positive outcomes. The blue state represents the most desirable outcome in our setting, and is fully revealed by consulting either the blue or the rainbow advisors. The slope of the curves is not significantly different across advisors (not shown in the table) confirming that the effect does not arise from a different sensitivity to instrumental value. Instead, it provides evidence in favor of intrinsic (non-instrumental) preference for information structures, similarly to that described in Section 2.8.2 in favor of advisors providing certainty or state pooling in the posterior beliefs.

The result is qualitatively robust to the separate analysis of trials with high or low status quo. Columns 3 and 4 contain the regressions ran independently with the two parts of the dataset. Although the signs and significance of the estimates are unchanged, we observe different magnitudes. When the value of the status quo is higher than the intermediate state (column 3) the intercept is lower and the slope steeper. In this case the decision maker faces a safer choice problem and she reacts more to the incentive represented by the instrumental value of the advisor.

Experimental Result 8. Participants display a compression effect in their willingness to pay for information structures. They tend to overpay for advisors with low and even zero

Method: OLS, Dependent var: w_I^i

	(1)	(2)	(3)	(4)
Constant	7.45***	6.32***	4.06***	8.56***
	(0.207)	(0.287)	(0.370)	(0.421)
w_I^{Bayes}	0.413***	0.298***	0.494***	0.122**
	(0.0336)	(0.0378)	(0.0515)	(0.0586)
Red advisor		-0.424	-0.465	-0.0366
		(0.439)	(0.548)	(0.680)
Blue advisor		3.84***	3.84***	2.97***
		(0.434)	(0.615)	(0.604)
Rainbow advisor		3.12***	3.02***	2.96***
		(0.464)	(0.615)	(0.680)
Trials	All	All	$R > v_y$	$R < v_y$
Observations	2520	2520	1260	1260

Table 2.3: Aggregate valuations of information structures in task 1. Reported significance levels for w_I^{Bayes} concern the hypothesis that this coefficient is 0. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

instrumental value, and their subjective WTP increases with the theoretical values but with a slope smaller than one.

Experimental Result 9. Participants are willing to pay significantly higher amounts for advisors that provide evidence in favor of the most desirable state, as well as for those that fully reveal the true state.

2.9 Conclusion

Opinions about proposed policies and pertinent issues often become polarized. The literature provides several explanations of the phenomenon, including preference for information which confirms existing beliefs, imperfect memory, and interpretation of ambiguous evidence as confirming existing beliefs, among others. In this paper we explore a new source of belief polarization, which arises as a consequence of the state-pooling effect, if the information is costly to acquire.

We find that the key determinant of the direction of belief updating is the valuation of the status quo, as it directly affects the information acquisition strategy. In our interpretation, the agent partitions the states of the world into categories, instead of distinct states. This partition into categories is determined exactly by the valuation of the status quo. If the two agents have different valuations of the status quo, they might diverge in their opinions after information acquisition, but also in the interim moment before the realization of the private signal. Interestingly, the difference in their mean beliefs can become greater if the information becomes cheaper to acquire.

The large number of assumptions required by the model may cast doubts on whether belief divergence can be observed in human decision makers, as behavioral biases could mitigate or enhance the effect. We introduce an experiment in which the availability of advisors (information structures) is instrumental for the action choice. We are able to qualitatively replicate the model's prediction in our setting, and we observe that the magnitude of the polarization is lower than predicted. We explore the possible drivers of this difference and conclude that intrinsic (non-instrumental) preferences for information represent the leading factor. This result is consistent with well-known results in simpler settings that we are able to replicate (preference for certainty) and extend based on the additional conditions allowed by our setup (preference for state pooling).

Our paper sheds new light on the problem of opinion polarization in society that is currently taking place. It provides an explanation of why polarization can become more severe when information is cheaper to obtain. We believe that our results provide a useful starting point and encourage further exploration of this phenomenon in several directions. One possible extension of the model would include a larger action space with more than two options; this feature would allow the creation of several endogenous categories and provide a connection with the models of categorical thinking. Another interesting extension of the model is represented by the addition of a strategic voting layer on top of the model we presented. On the experimental side, the introduction of our design opens the path for further replications and tests of well-known paradigms and effects in more complex settings with three or more possible states. Possible extensions of the tasks we discussed would explore both instrumental and non-instrumental preferences for information in new scenarios. Finally, on the empirical side, we encourage future research testing the implications of our model on referendum data.

2.A Appendix A: Derivation of formula 2.5

The agent's posterior belief about the payoff of the new policy v given the fixed state s^* for option $i \in \{\text{status quo }, \text{new policy}\} = \{1, 2\}$ is

$$\mathbb{E}_i[\mathbb{E}(v|i)|s^*] = \mathcal{P}(i = 1|s^*)\mathbb{E}(v|i = 1) + \mathcal{P}(i = 2|s^*)\mathbb{E}(v|i = 2).$$

After substituting for the conditional probabilities $\mathcal{P}(i|s^*)$ $\forall i$ according to lemma 1 and applying the Bayes rule this can be rewritten as

$$\mathbb{E}_i[\mathbb{E}(v|i)|s^*] = \frac{\mathcal{P}(i=1)e^{\frac{v_{s^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \cdot \sum_{s=1}^{n} v_s g_s \frac{e^{\frac{v_s}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} +$$
$$+ \frac{(1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \cdot \sum_{s=1}^{n} v_s g_s \frac{e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}.$$

Lemma 1 shows that

$$\mathcal{P}(i = 1|s^*) = \frac{\mathcal{P}(i=1)e^{\frac{v_{s^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}.$$

Thus,

$$\mathbb{E}_i[\mathbb{E}(v|i)|s^*] = \sum_{s=1}^{n} v_s g_s \frac{\mathcal{P}(i=1|s^*)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1|s^*))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}((i=1))e^{\frac{R}{\lambda}}}.$$

2.B Appendix B: Proof of proposition 1

First we prove the following lemma that we further use for proving Proposition 1.

Lemma 2. *Relations* $\mathcal{P}(i = 1|s^*) \gtrless P(i = 1)$ *for* $0 < \mathcal{P}(i = 1) < 1$ *are equivalent to* $v_{s^*} \gtrless R$.

Proof. After substitution for the conditional probabilities, the conditions $\mathcal{P}(i = 1|s^*) \gtrless P(i = 1)$ can be rewritten as

$$\frac{\mathcal{P}(i=1)e^{\frac{v_{s*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \gtrless \mathcal{P}(i=1),$$

which are equivalent to

$$\left(\mathcal{P}(i=1) - \mathcal{P}^2(i=1)\right)\left(e^{\frac{v_{s*}}{\lambda}} - e^{\frac{R}{\lambda}}\right) \gtrless 0.$$

For $0 < \mathcal{P}(i = 1) < 1$ the term in the first parenthesis is always positive. Therefore, the left hand side of the inequality is positive when $v_{s^*} > R$ and negative for $v_{s^*} < R$. □

Now we can continue with the proof of Proposition 1.

Proof. In order to solve the agent's problem given by equations 2.1 - 2.4 we need to find $\mathcal{P}(i = 1)$ and $\mathcal{P}(i = 2)$ defined as $\mathcal{P}(i = 2) = 1 - \mathcal{P}(i = 1)$. These probabilities have to be internally consistent, i.e. $\mathcal{P}(i) = \sum_{s=1}^{n} \mathcal{P}(i|s)g_s$. After dividing both sides of these conditions by $P(i)$ we obtain the following conditions

$$1 = \sum_{s=1}^{n} \frac{e^{\frac{v_s}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + \mathcal{P}(i=2)e^{\frac{R}{\lambda}}} g_s, \quad \text{if } \mathcal{P}(i=1) > 0,$$

$$1 = \sum_{s=1}^{n} \frac{e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + \mathcal{P}(i=2)e^{\frac{R}{\lambda}}} g_s, \quad \text{if } \mathcal{P}(i=2) > 0.$$

The difference of these two equations is

$$\sum_{s=1}^{n} \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + \mathcal{P}(i=2)e^{\frac{R}{\lambda}}} g_s = 0.$$

For k for which holds that $v_k \leq R \leq v_{k+1}$ we can further write the above equation as

$$\frac{e^{\frac{v_k}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_k}{\lambda}} + \mathcal{P}(i=2)e^{\frac{R}{\lambda}}} v_k g_k = -\sum_{s \neq k} \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + \mathcal{P}(i=2)e^{\frac{R}{\lambda}}} v_k g_s. \tag{2.8}$$

We will use the last equation for determining the sign of $\Delta(s^*)$ that can be written as

$$\Delta(s^*) = \sum_{s=1}^{n} v_s g_s \frac{\mathcal{P}(i=1|s^*)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1|s^*))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} - \sum_{i=1}^{n} v_s g_s,$$

$$\Delta(s^*) = \sum_{s=1}^{n} v_s g_s \frac{\mathcal{P}(i=1|s^*)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1|s^*))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} - \sum_{i=1}^{n} v_s g_s \frac{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}},$$

$$\Delta(s^*) = \sum_{i=1}^{n} v_s g_s \frac{(\mathcal{P}(i=1|s^*) - \mathcal{P}(i=1))(e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}})}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}},$$

$$\Delta(s^*) = (\mathcal{P}(i=1|s^*) - \mathcal{P}(i=1)) \cdot \sum_{s=1}^{n} v_s g_s \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}.$$

Substituting the equation (2.8) into the sum in the last equation we obtain

$$\Delta(s^*) = (\mathcal{P}(i=1|s^*) - \mathcal{P}(i=1)) \left[\sum_{s \neq k} (v_s - v_k) g_s \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} \right].$$

The expression in the square brackets is positive, because for the above-defined k the sign of $(v_s - v_k)$ and the sign of $e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}$ are the same. Hence $\Delta(s^*)$ has the same sign as $(\mathcal{P}(i=1|s^*) - \mathcal{P}(i=1))$ that further, by Lemma 2, has the same sign as $(v_{s^*} - R)$.

□

2.C Appendix C: Proof of proposition 2

Proof. We are interested in the monotonicity of $\Delta(s^*)$ when the true state of the world s^* is changing. In appendix 2.B we derive that

$$\Delta(s^*) = (\mathcal{P}(i=1|s^*) - \mathcal{P}(i=1)) \left[\sum_{s\neq k} (v_s - v_k) g_s \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}\right].$$

Let us consider two states of the world s_1^* and s_2^*, such that $s_1^* > s_2^*$. Demonstrating that $\Delta(s_1^*) - \Delta(s_2^*) \geq 0$ would prove the monotonicity of $\Delta(s^*)$.

$$\begin{aligned}
\Delta(s_1^*) - \Delta(s_2^*) &= \left[\sum_{s\neq k} (v_s - v_k) g_s \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}\right] \cdot \\
&\cdot (\mathcal{P}(i=1|s_1^*) - \mathcal{P}(i=1) - \mathcal{P}(i=1|s_2^*) + \mathcal{P}(i=1)) = \\
&= \left[\sum_{s\neq k} (v_s - v_k) g_s \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}\right] \cdot \\
&\cdot (\mathcal{P}(i=1|s_1^*) - \mathcal{P}(i=1|s_2^*)) = \\
&= \left[\sum_{s\neq k} (v_s - v_k) g_s \frac{e^{\frac{v_s}{\lambda}} - e^{\frac{R}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_s}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}\right] \cdot \\
&\cdot \left(\frac{\mathcal{P}(i=1)e^{\frac{v_{s_1^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s_1^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} - \frac{\mathcal{P}(i=1)e^{\frac{v_{s_2^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s_2^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}}\right).
\end{aligned}$$

The term in the square brackets is positive, so the sign of $\Delta(s_1^*) - \Delta(s_2^*)$ is determined by the sign of the term in the round brackets.

Let us show that

$$\frac{\mathcal{P}(i=1)e^{\frac{v_{s_1^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s_1^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} - \frac{\mathcal{P}(i=1)e^{\frac{v_{s_2^*}}{\lambda}}}{\mathcal{P}(i=1)e^{\frac{v_{s_2^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}} > 0.$$

The last inequality is equivalent to

$$\mathcal{P}(i=1)e^{\frac{v_{s_1^*}}{\lambda}} \left(\mathcal{P}(i=1)e^{\frac{v_{s_2^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}}\right) -$$

$$- \mathcal{P}(i=1)e^{\frac{v_{s_2^*}}{\lambda}} \left(\mathcal{P}(i=1)e^{\frac{v_{s_1^*}}{\lambda}} + (1-\mathcal{P}(i=1))e^{\frac{R}{\lambda}} \right) > 0,$$

$$(1-\mathcal{P}(i=1))\, e^{\frac{v_{s_1^*}}{\lambda}} e^{\frac{R}{\lambda}} - (1-\mathcal{P}(i=1))\, e^{\frac{v_{s_2^*}}{\lambda}} e^{\frac{R}{\lambda}} > 0,$$

which, in turn, is equivalent to

$$e^{\frac{v_{s_1^*}}{\lambda}} > e^{\frac{v_{s_2^*}}{\lambda}}.$$

The last inequality holds, so $\Delta(s^*)$ is an increasing function of s^*.

□

Chapter 3

Estimating Models with Rationally Inattentive Agents

3.1 Introduction

A perennial question in macroeconomic theory is how firms set prices and more generally, the reason for the observed real effects of changes in monetary policy. A famous answer to this question is that people are not well enough informed about changes in market conditions, at least at the time these changes occur, to be able to immediately react in the way that would most fully serve their own interests. Woodford (2002) reconsidered this Phelps-Lucas hypothesis, according to which a temporary real effect of purely nominal disturbances results from imperfect information, and showed that nominal shocks have strong and persistent real effects due to imperfect common knowledge. Woodford (2002), however, assumed that firms pay little attention to the aggregate economy and that they set prices based on signals of the form "nominal aggregate demand plus i.i.d. noise". Rational inattention allows price-setting firms to decide what to pay attention to, subject to a constraint on information flow. Mackowiak and Wiederholt (2009) characterized conditions under which firms pay more attention to idiosyncratic conditions than to aggregate conditions. Maćkowiak, Matějka, and Wiederholt (2018) further showed that an agent with memory and limited attention wants to learn about the current optimal action and the best predictors of future optimal actions. Thus, an agent selects a signal of the AR(1) form only when it tracks a variable following the AR(1) process, which suggests that signal assumptions in Woodford (2002) might be too restrictive. Furthermore, even though it is known that there exist parameter values for which rational inattention models generate persistent real effects of nominal disturbances (see e.g., Mackowiak and Wiederholt (2015)), it is not known whether a full information method would select such parameter values.

In this paper, we estimate a DSGE model with rationally inattentive price-setting firms, based upon the estimated ICKM model presented in Melosi (2014). This allows us to directly compare three estimated models: an RIM, an ICKM and a Calvo pricing model. Importantly, such a comparison helps us to clarify the previously raised unknowns. That is, to see what parameter values would be selected for the RIM and how well it will match the data in contrast to other models. It also sheds light on the restrictiveness and implications of the ICKM assumed signal, structure versus the optimally selected signal given the information capacity constraint.

The RIM model shares with the Melosi (2014) model, which is a version of the ICKM model presented in Woodford (2002), that it has two exogenous state variables: the state of monetary policy and the state of technology. In contrast with previous studies, we do not assume any particular exogenously given signal form or independence of signals, but by modelling firms as rationally inattentive we allow firms to choose the optimal signals about the state variables optimally under the limited attention constraint. We show that the signal is a one-dimensional signal about the elements of the state vector even when optimal action may be driven by multiple shocks. Moreover, the selected noisy signal follows a more complex ARMA process as just a "true value plus i.i.d. noise", because the rational inattention firms want to learn about the current optimal action as well as about the best predictors of future optimal actions.

The substantive contribution of this study is in showing that the RI model matches the data better than the Calvo model, in particular by reproducing the persistence in the data more easily. Furthermore, the RI model reproduces the long-term mean-reversion better than the ICKM model, whereas the ICKM model seems to perform better in matching the short-run momentum of the hump-shaped response of output and inflation to monetary disturbances. This paper provides one important step in developing and estimating a medium-size DGSE model with RI agents, which might ultimately prove to be an alternative or complement to the New Keynesian models. Previous RI models were only calibrated because dynamic rational inattention models are difficult to solve. We build on the results presented in Maćkowiak, Matějka, and Wiederholt (2018), which allows us to overcome the key challenge to solve the model sufficiently fast so that estimation becomes feasible in a reasonable amount of time.

Our likelihood analysis makes another contribution. When we set the information flow in the RIM and ICKM models to be equal, the estimated parameter values differ mostly in the value of the posterior degree of the strategic complementarity. Specifically, the higher degree of strategic complementarity leads the RIM model to better capture the hump-shaped impulse response function. All the aforementioned results provide a strong empirical validation for the RIM, as well as indicating for what applications an RIM or other model would be more suitable.

Related Literature. We make a priority contribution to the growing body of macroeconomics studies modelling agents to be rationally inattentive. Sims (1998b) proposed the idea of rational inattention as an ultimate single information friction that could serve to explain the inertia in the macroeconomic data, instead of multiple sources of slow adjustment that were necessary for the models to match this inertia. Mackowiak and Wiederholt (2009) showed under what circumstances firms find it optimal to pay little attention to the aggregate economy. As a result, prices respond strongly and quickly to idiosyncratic shocks, but only weakly and slowly to nominal shocks. Nominal shocks have strong and persistent real effects. Mackowiak and Wiederholt (2015) developed a dynamic stochastic general equilibrium model with rational inattention households and decision-makers in firms and compared its prediction to the data. They showed that their model matches the empirical impulse responses to monetary policy shocks and aggregate technology shocks. These models were, however, at most calibrated. We build upon findings of Maćkowiak, Matějka, and Wiederholt (2018) that presented novel analytical results for solving dynamic rational inattention problems. The closest paper to ours is Melosi (2014) who estimated, using Bayesian methods, an ICKM model with two shocks: a disturbance to nominal aggregate demand and technology shock, and found that imperfect common knowledge is more successful than Calvo price stickiness model to account for the highly persistent effects of nominal shocks.

The importance of rational inattention models was further strengthened by empirical findings that there is a widespread dispersion in firms' beliefs about both past and future macroeconomic conditions, especially inflation, with average beliefs about recent and past inflation being much higher than those of professional forecasters (Coibon, Gorodnichenko, and Kumar 2015), and they also show that these patterns are in line with rational inattention theory. Kamdar (2019) using survey data, demonstrates that consumers' economic beliefs are driven by sentiment and thus, "optimistic" consumers, in contrast with recent U.S. experience, expecting an expansion also predict disinflation. Kamdar (2019) explains this observation by employing a model with rationally inattentive consumers who face fundamental uncertainty. Our findings are also related to the survey findings of Afrouzi (2020) that firms with fewer competitors pay less attention to monetary policy shocks and firms in oligopolies with fewer competitors have a lower degree of strategic complementarity. Rational inattention also proved to be important for policy analysis as policy changes incentives for allocation of attention (see e.g., Paciello and Wiederholt (2014),Mackowiak and Wiederholt (2015)).

Outline. The paper is organized as follows. Section 3.2 presents the common assumptions and features of the three models. In section 3.3 we introduce the rational inattention optimal signals into the model and present the solution method. Section 3.4 and section 3.5 show details of the imperfect common knowledge model and the Calvo pricing model, respectively. Section 3.6 presents the empirical findings and its discussion. Section 3.7 concludes.

3.2 The Model

In this section, closely following the model introduced in Melosi (2014), we present the main building blocks of the economy that are shared among the various models introduced later. In particular, first in section 3.2.1 we present a shared assumptions. Sections 3.2.2-3.2.5 introduce the agents in the economy: representative household, financial intermediary, monetary authority and price-setting firms. Section 3.2.6 describes the log-linearization of the model and the law of motions for price level and output.

3.2.1 Assumptions

The economy consists of the representative households, a financial intermediary, a monetary authority, and a continuum $(0,1)$ of monopolistically competitive firms. The model has two exogenous state variables: the stock of money and the state of technology.

We assume that all information is publicly available to every agent in the economy. Firms cannot attend perfectly to all available information. How we model information acquisition varies for each consider model considered. The share assumption for the information acquisition is that all other agents except the firms perfectly observe the past and current realizations of all the model variables.

Timing is as follows. At the beginning of period t the household inherits the entire money stock of the economy M_{t-1}. All the shocks and signals are realized. After observing current-period shocks, the representative household decides how much money D_t to deposit at the financial intermediary that yields interest at a rate $R_t - 1$. The financial intermediary lends to firms, at a fixed fee τ, funds that were collected from households' deposits and from the monetary authority. After receiving signals, firms set their prices and hire labor to which they then pay a nominal hourly wage W_t for H_t hours worked. Households face a cash-in-advance (CIA) constraint. That is, they have to pay for all consumption with the accumulated cash balance, which after receiving wages increased to $M_{t-1} - D_t + W_t H_t$. Afterwards firms pay back their loans $L_{i,t}$, and households receive their deposits plus dividends and interest.

3.2.2 The representative household

The representative household derives utility from consuming the consumption good C_t and disutility from hours worked H_t, and solves

$$\max_{C_t,H_t,D_t\geq 0} \mathbb{E}_t \sum_{s=0}^{\infty} \left[\ln C_{t+s} - \alpha \frac{H_{t+1}^{1+\gamma}}{1+\gamma} \right], \tag{3.1}$$

where β is the discount factor, α is a parameter that affects the marginal utility of leisure, and $\gamma > 0$ is the inverse of Frisch labor elasticity. C_t is given by the Dixit-Stiglitz aggregator

$$C_t = \left(\int_0^1 C_{i,t}^{\frac{\nu-1}{\nu}} di \right)^{\frac{\nu}{\nu-1}}, \tag{3.2}$$

where $C_{i,t}$ is the consumption of differentiated good i in period t. The parameter ν is the elasticity of substitution between consumption goods and is assumed that $\nu > 1$. Households are subject to the cash-in-advance constraint, due to which they hold money upfront to finance their consumption,

$$P_t C_t \leq M_{t-1} - D_t + W_t H_t \tag{3.3}$$

where P_t denotes the price level, which is given by

$$P_t = \left(\int (P_{i,t})^{1-\nu} di \right)^{\frac{1}{1-\nu}} \tag{3.4}$$

and the law of motion of households' cash is

$$M_t = (M_{t-1} + W_t H_t - D_t - P_t C_t) + R_t D_t + \Pi_t + \Pi_t^b. \tag{3.5}$$

3.2.3 The financial intermediary

The financial intermediary solves every period t, following the static problem:

$$\max_{\{L_t, D_t\}} (1 - R_t) D_t + X_t + \tau \cdot \vartheta_t \cdot \mathbb{I}\{L_t > 0\} \tag{3.6}$$

such that

$$L_t \leq X_t + D_t, \tag{3.7}$$

where $X_t = M_t - M_{t-1}$ is the monetary injection, L_t is the aggregate amount of liquidity supplied to firms. $\mathbb{I}\{L_t > 0\}$ is an indicator function that equals one, if liquidity supplied to firms is positive, and the fraction of firms that borrow is indicated by ϑ_t . We consider a fixed fee τ that is paid to the financial intermediary for its services, instead of an equilibrium interest rate on loans. This is in order to keep the model as close as possible to the previously estimated model by Melosi (2014) and also as it significantly simplifies solution of the model.

3.2.4 The monetary authority

The monetary authority sets the money stock M_t according to an empirical monetary policy rule without feedback:

$$\Delta \ln M_t = (1 - \rho_m)\mu_M + \rho_m \Delta \ln M_{t-1} + \sigma_m \epsilon_{m,t}, \quad \epsilon_{m,t} \sim \mathcal{N}(0,1) \tag{3.8}$$

where Δ stands for the first-difference operator and μ_M is a parameter that represents the long-run average growth rate of money. $\rho_m \in [0,1)$ is a degree of smoothness in conducting monetary policy. The monetary shock $\epsilon_{m,t}$ captures unexpected changes in the growth rate of the money stock in every period t. Market clearing for the monetary market requires that:

$$\ln M_t = \ln Y_t + \ln P_t. \tag{3.9}$$

3.2.5 Firms

Firm i's expected profit in period t (as valued by households) conditional on its information set at time t, $\mathcal{I}_{i,t} = \mathcal{I}_{i,0} \cup \{S_{i,1}, \ldots, S_{i,t}\}$, where $\mathcal{I}_{i,0}$ denotes the initial information set. We assume that firms are endowed with an infinite sequence of signals at time 0. $S_{i,t}$ denotes the signal vector received in period t is

$$\mathbb{E}[\beta^t Q_t (P_{i,t} Y_{i,t} - W_t N_{i,t} - \tau \mathbb{I}\{L_{i,t}\} | \mathcal{I}_{i,t})] \tag{3.10}$$

where $\beta^t Q_t$ is the time 0 value of one unit of the consumption good in period t to the representative household, $Y_{i,t}$ is the amount of goods i produced by firm i at time t, and $N_{i,t}$ is the labor input demanded by firm i at time t. The production function is given by

$$Y_{i,t} = A_{i,t} N_{i,t}^{\phi} \tag{3.11}$$

where $\phi \in (0,1)$ is the return-to-scale parameter and $A_{i,t}$ is the firm's level of technology. Based on Lorenzoni (2009), the firm's level of technology is modeled as

$$\ln A_{i,t} = \ln A_t + \sigma_\eta \eta_{i,t}, \quad \eta_{i,t} \sim \mathcal{N}(0,1) \tag{3.12}$$

and the aggregate level of technology, $\ln A_t$, follows a random walk with drift

$$\ln A_t = \mu_A + \ln A_{t-1} + \sigma_a \epsilon_{a,t}, \quad \epsilon_{a,t} \sim \mathcal{N}(0,1) \tag{3.13}$$

We assume that the technology shocks, $\epsilon_{a,t}$ and $\eta_{i,t}$, are orthogonal to monetary shocks $\epsilon_{m,t}$, at all leads and lags. Firms borrow liquidity $L_{i,t}$ at the fixed cost τ from the financial intermediary to pay their nominal labor costs:

$$L_{i,t} = W_t N_{i,t} \tag{3.14}$$

In every period t, firm i sets the price of good i and commits to supply any quantity of the good demanded at that price so that $Y_{i,t} = C_{i,t}$. Each firm i faces the following demand function

$$C_{i,t} = \left(\frac{P_{i,t}}{P_t} \right)^{-\nu} C_t. \tag{3.15}$$

3.2.6 Price-setting equation and law of motions

In order to simplify the signal-extraction issues, we assume that firms use the log-linearized model, rather than the original nonlinear model, when addresing their signal-extraction problem. Since the exogenous processes, i.e. the stock of money $\ln M_t$ and aggregate technology $\ln A_t$ have a unit root, all endogenous variables but labor are also nonstationary. Thus, we first need to detrend the variables before log-linearizing the models. We define the stationary endogenous variables as follows

$$y_t \equiv Y_t/A_t;\ y_{i,t} \equiv Y_{i,t}/A_t;\ p_t \equiv P_t A_t/M_t \text{ and } p_{i,t} \equiv P_{i,t}/P_t$$

We also define $m_t \equiv \ln M_t - \mu_M t$ and $a_t \equiv \ln A_t - \mu_A t$. We can obtain the following log-linearized price-setting equation for firm i at time t and law of motions for price level and output, respectively

$$\ln P_{i,t} = (1-\lambda)\mathbb{E}[\ln P_t|\mathcal{I}_{i,t}] + \lambda\mathbb{E}[\ln M_t|\mathcal{I}_{i,t}] + \lambda\mathbb{E}[\ln A_t|\mathcal{I}_{i,t}] - \frac{\lambda}{\gamma+1} + \lambda\mathbb{E}[\eta_{i,t}|\mathcal{I}_{i,t}] - \lambda\ln\bar{y} \tag{3.16}$$

$$\ln P_t = \left[\sum_{j=0}^{\infty}(1-\lambda)^j\lambda\left(m_{t|t}^{(j+1)} - a_{t|t}^{(j+1)}\right)\right] - \lambda\ln\bar{y} + \mu_M\cdot t - \mu_A\cdot t \tag{3.17}$$

$$\ln Y_t = \left[m_t - \sum_{j=0}^{\infty}(1-\lambda)^j\lambda m_{t|t}^{(j+1)}\right] + \sum_{j=0}^{\infty}(1-\lambda)^j\lambda a_{t|t}^{(j+1)} + \lambda\ln\bar{y} + \mu_A\cdot t, \tag{3.18}$$

Derivation of equations 3.16-3.18 is presented in Appendix 3.A, $\bar{y}$ denotes the steady-state value of the detrended real output, y_t. The parameter $(1-\lambda)$, defined as $(1-\lambda) \equiv 1-(\gamma+1)\phi^{-1}/[\nu(\phi^{-1}-1)+1]$, denotes the degree of strategic complementarity. Because the received signals are private to price setters, they are uncertain about what other price setters believe about the realized shocks. When $1-\lambda > 0$, firms react to their beliefs about other firms' price-setting decisions $\mathbb{E}[\ln P_t|\mathcal{I}_{i,t}]$. The variables $m_{t|t}^{(j)}$ and $a_{t|t}^{(j)}$ are the average expectations of order j about the state of monetary policy, m_t, and the state of technology, a_t, respectively. The average j-th order expectations about the state of monetary policy are defined as $m_{t|t}^{(j)} \equiv \int m_{t|t}^{(j)}(i)di$, where $m_{t|t}^{(j)}(i) \equiv \mathbb{E}\left[m_{t|t}^{(j-1)}|\mathcal{I}_{i,t}\right]$. Average expectations about technology are analogously defined.

Previous sections introduced the definition of all the economic agents and specifications that are shared among all the models considered. The main feature that differentiates the models is the information acquisition technology and a restriction on signal's availability.

In the following section, we introduce the main contribution that considers firms to be rationally inattentive. Sections 3.4 and 3.5 present the imperfect common knowledge model and Calvo pricing, respectively.

3.3 Rational inattention model

In the rational inattention version of the model, firms decide how to allocate their limited attention and choose the signal structure in order to maximize the expected profit subject to the information flow constraint. The economy is driven by two shocks to the (log of) the money stock m_t and (log of) firm specific technology $a_{i,t}$. We assume that m_t follows an ARMA (p_m, q_m) process and that $a_{i,t}$ follows an ARMA (p_a, q_a) process, with p_m, q_m, p_a, q_a finite. Therefore, firm i chooses the signals for tracking these shocks subject to the information constraint. Specifically, the firm chooses the number of signals K, the content of the signals matrices: A_1, A_2, B_1, B_2 and the variance-covariance matrix of noise in the signals Σ_ψ. In period zero, the decision-maker in firm i solves

$$\max_{K,A_1,A_2,B_1,B_2,\Sigma_\psi} \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^t \hat{\pi}_t(\hat{p}_{i,t}, \hat{p}_t, \hat{y}_t, \hat{q}_t)\right], \tag{3.19}$$

subject to:

$$\ln P_{i,t} = \mathbb{E}\left[(1-\lambda)\ln P_t + \lambda m_t - \lambda a_t - \frac{\lambda}{\gamma+1}\eta_{i,t} | \mathcal{I}_{i,t}\right] - \lambda\left[(\mu_A - \mu_M)t + \ln \bar{y}\right], \tag{3.20}$$

$$\mathcal{I}_{it} = \mathcal{I}_{i,0} \cup \{S_{i,1}^K \dots S_{i,t}^K\}$$

$$S_{i,t}^K = A_1 \begin{pmatrix} m_t \\ \vdots \\ m_{t-M+1} \end{pmatrix} + A_2 \begin{pmatrix} a_{i,t} \\ \vdots \\ a_{i,t-M+1} \end{pmatrix} + B_1 \begin{pmatrix} \epsilon_{m,t} \\ \vdots \\ \epsilon_{m,t-N+1} \end{pmatrix} + B_2 \begin{pmatrix} \epsilon_{a_i,t} \\ \vdots \\ \epsilon_{a_i,t-N+1} \end{pmatrix} + \psi_{it}^K \tag{3.21}$$

$$\lim_{T\to\infty} \frac{1}{T} I\left(\bar{m}_0, \bar{a}_{i,0}, m_1, a_{i,1}, \ldots, m_T, a_{i,T}; S_{i,1}^K, \ldots, S_{i,T}^K\right) \leq \kappa \tag{3.22}$$

where the vectors $\bar{m}_0$ and $\bar{a}_{i,0}$ in the information flow constraint denotes the vector of initial conditions for the processes for the optimal actions [1]. $\tilde{\pi}_t(\cdot)$ is the log-quadratic approximation of $Q_t\pi_t$; π_t is the profit function. $\hat{p}_{i,t} \equiv ln(P_{i,t}/P_t)$ is the profit maximizing price; and $\hat{q}_t$ is the log deviations of $q_t \equiv M_tQ_t$ from its value at the deterministic steady state. ψ_{it}^K follows a Gaussian vector white noise process with variance-covariance matrix Σ_ψ. $\mathcal{I}_0$ denotes the initial information set and S_t^K denotes the signal vector received in period t, where $K \geq 1$ denotes the dimension of the signal vector. Thus, a firm's information set in any period $t \geq 1$ includes the initial information and all signals received up to time t. When firms decide how to allocate their limited attention, they are aware that their choices will affect their optimal price-setting policy as given by (3.20) in any subsequent periods.

Using results from Maćkowiak, Matějka, and Wiederholt (2018) we can rewrite the problem (3.19)-(3.22). In section 3.3.1 we explore how to alternatively quantify the information flow constraint and in section 3.3.2 we rewrite the objective of the maximization problem.

3.3.1 Quantifying information flow

The idea that the firm has a limited amount of attention is captured by the information flow constraint (3.22). Rational inattentiveness refers to the fact that the firm selects signals optimally given these constraints. In quantifying the information flow we follow

[1]That is $\bar{m}_0 = (m_{1-p_1}, \ldots, m_0, \epsilon_{m,1-q_1}, \ldots, \epsilon_{m,0})$ and $\bar{a}_{i,0} = (a_{i,1-p_2}, \ldots, a_{i,0}, \epsilon_{a_i,1-q_2}, \ldots, \epsilon_{a_i,0})$ as we have assumed that m_t follows the ARMA (p_m, q_m) and $a_{i,t}$ follows the ARMA (p_a, q_a), respectively

Sims (2003a). Thus the information flow constraint is quantified by the reduction in uncertainty, where uncertainty is measured by entropy. The entropy of a random vector $X^T = (X_1, \ldots, X_T)$ as $H(X^T)$ and the conditional entropy of X^T given knowledge of a random vector $S^T = (S_1, \ldots, S_T)$ as $H(X^T|S^T)$. We define the mutual information between two random vectors X^T and S^T as

$$I(X^T; S^T) = H(X^T) - H(X^T|S^T)$$

Dividing both sides by T and taking the limit $T \to \infty$ we obtain the previously stated constraint (3.22)

$$\lim_{T\to\infty} \frac{1}{T} I\left(X^T; S^T\right) = \lim_{T\to\infty} \frac{1}{T} H(X^T) - \lim_{T\to\infty} \frac{1}{T} H(X^T|S^T) \tag{3.23}$$

the first term on the right-hand side measures how total uncertainty about X^T grows per unit of time. The second term on the right-hand side measures how total uncertainty about X^T grows per unit of time given knowledge of S^T, and the difference between the two terms measures the information flow to the firm.

Maćkowiak, Matějka, and Wiederholt (2018) showed that the information flow constraint (3.22) is equivalent to a constraint on the difference between prior uncertainty and posterior uncertainty at a given point in time. This result, adopted for this paper setup with two shocks, can be stated as follows:

Lemma 3. *Let $\mathcal{S}^{K,t} = \{S_1^K, \ldots, S_{-t}^K\}$ denote the set of signals received up to and including time t. The information flow constraint (3.22) is equivalent to*

$$\lim_{T\to\infty} \left[H\left(\varphi_T|\mathcal{S}^{K,T-1}\right) - H\left(\varphi_T|\mathcal{S}^{K,T}\right)\right] \tag{3.24}$$

where the vector φ_t can be any vector with the following two properties: (i) m_t^M, a_t^M, $\epsilon_{m,t}^N$, $\epsilon_{a,t}^N$ can be computed from φ_t and (ii) φ_t contains no redundant elements.

Proof. Direct consequence of Lemma 1 in Maćkowiak, Matějka, and Wiederholt (2018). □

3.3.2 The objective

In every period, firm i sets the price of good i to maximize the present discounted value of profits. Since the firm can reset the price next period, this is equivalent to setting the price to maximize current profit. After a log-quadratic approximation to the profit function, the loss in profit in the case of a deviation of the actual price, $\hat{p}_{it}$, from the profit-maximizing price, $\hat{p}^*_{it}$, is proportional to $(\hat{p}^*_{it} - \hat{p}_{it})^2$ (for details see Appendix 3.B). Therefore, a firm's optimal price given any information set $\mathcal{I}_{it}$ is $\hat{p}_{it} = \mathbb{E}[\hat{p}^*_{it}|\mathcal{I}_{it}]$. Substituting this equation into the quadratic loss function yields the mean square error $\mathbb{E}\left[\hat{p}^*_{it} - \mathbb{E}[\hat{p}^*_{it}|\mathcal{I}_{it}]\right]$.

Applying the assumption that the agent receives a long sequence of signals in period zero such that conditional second moments are independent of time. Thus, conditional second moments can be computed using the steady-state Kalman filter. This assumption about the information set $\mathcal{I}_{i0}$ after the agent has made the information choice in period zero has three implications. First, it does not matter which period $t \geq 1$ one is referring to in the agent's loss function. Second, replacing the agent's loss function by the loss function

$$\mathbb{E}\left[\sum_{t=1}^{\infty} \beta^t (\hat{p}^*_{it} - \mathbb{E}[\hat{p}^*_{it}|\mathcal{I}_{it}])^2\right] = \frac{\beta}{1-\beta}\mathbb{E}\left[(\hat{p}^*_{it} - \mathbb{E}[\hat{p}^*_{it}|\mathcal{I}_{it}])^2\right], \; \beta \in (0,1) \tag{3.25}$$

is a monotone transformation of the objective and thus does not affect the solution to the dynamic rational inattention problem. Third, the conditional expectation $\mathbb{E}[\hat{p}^*_{it}|\mathcal{I}_{it}]$ is a time-invariant function of the signals.

Finally, the rational inattention problem (3.19)-(3.22) can be analogously rewritten as

$$\min_{K,A,B,\Sigma_\psi} \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^t (\hat{p}^*_{it} - \mathbb{E}[\hat{p}^*_{it}|\mathcal{I}_{it}])^2\right] \tag{3.26}$$

where $\beta \in (0,1)$ is a parameter, subject to the information flow constraint (3.22)

$$\mathcal{I}_{it} = \mathcal{I}_{i0} \cup \{S^K_{i1}, \ldots, S^K_{it}\} \tag{3.27}$$

and

$$S_{it}^K = A \begin{pmatrix} \hat{p}_{it}^* \\ \vdots \\ \hat{p}_{i,t-M+1}^* \end{pmatrix} + B \begin{pmatrix} \epsilon_t \\ \vdots \\ \epsilon_{t-N+1} \end{pmatrix} + \psi_{it}^K \tag{3.28}$$

where ψ_{it}^K follows a Gaussian vector white noise process with variance-covariance matrix Σ_ψ and the optimal price is driven by the shock to the firm's level of technology a_{it} and by the shock to the stock of money m_t. The decision-maker minimizes the expected discounted sum of profit losses due to suboptimal pricing. She understands that in every period $t \geq 1$ she will set the price equal to the conditional expectation of the profit-maximizing price and she will remember all past signals.

3.3.3 Rational inattention Kalman filter formulation

Having introduced the rational inattention problem (3.19)-(3.22) and showed that it can be reformulated as (3.26)-(3.28), we briefly describe the solution method for the reformulated problem. We build directly on a so-called rational inattention Kalman filter that was introduced in Maćkowiak, Matějka, and Wiederholt (2018). We first restrict the dimensionality of the optimal signal and then we specify that an optimal signal includes variables only about the state variables and its innovations.

Lemma 4. *Firm i can attain the optimum with a one-dimensional signal, i.e. $K = 1$.*

Proof. Based on Proposition 2 from Maćkowiak, Matějka, and Wiederholt (2018), which imposes a restriction that the optimum in an economy with two shocks can be obtained with $K \leq 2$. Further restriction is possible through application of Proposition 3 and Proposition 9 from Miao, Wu, and Young (2019). □

Lemma 5. *Assume that state variables m_t and a_{it} follow an ARMA(p_m, q_m) process and ARMA (p_a, q_a) process, respectively. Then, any optimal signal vector is on linear combinations of the elements of $\xi_{m,t}$ and $\xi_{a,t}$ only, where $\xi_{m,t}$ and $\xi_{a,t}$ are defined as follows*

$$
\xi_{m,t} = \begin{cases} (m_t, \ldots, m_{t-(p_m-1)})'; & \textit{if } p_m \geq 1 \textit{ and } q_m = 0 \\ (m_t, \ldots, m_{t-(p_m-1)}, \epsilon_t, \ldots, \epsilon_{t-(q_m-1)})'; & \textit{if } p_m \geq 1 \textit{ and } q_m \geq 1 \\ m_t; & \textit{if } p_m = 0 \textit{ and } q_m = 0 \\ (m_t, \epsilon_t, \ldots, \epsilon_{t-(q_m-1)})'; & \textit{if } p_m = 0 \textit{ and } q_m \geq 1 \end{cases} \tag{3.29}
$$

$$
\xi_{a,t} = \begin{cases} (a_{it}, \ldots, a_{it-(p_a-1)})'; & \textit{if } p_a \geq 1 \textit{ and } q_a = 0 \\ (a_{it}, \ldots, a_{it-(p_a-1)}, \epsilon_t, \ldots, \epsilon_{t-(q_m-1)})'; & \textit{if } p_a \geq 1 \textit{ and } q_a \geq 1 \\ a_{it}; & \textit{if } p_a = 0 \textit{ and } q_a = 0 \\ (a_{it}, \epsilon_t, \ldots, \epsilon_{it-(q_a-1)})'; & \textit{if } p_a = 0 \textit{ and } q_a \geq 1 \end{cases} \tag{3.30}
$$

Proof. Direct generalization of the Proposition 1 from Maćkowiak, Matějka, and Wiederholt (2018). □

Applying Lemma 4 and Lemma 5 the search for the optimal signal can be restricted to one-dimensional signals that have the following state-space representation

$$
\xi_{t+1} = F\xi_t + v_{t+1} \tag{3.31}
$$

$$
S_t = h'\xi_t + \psi_t \tag{3.32}
$$

where $\xi_t = (\xi_{m,t}, \xi_{a,t})'$ and the noise ψ_t follows a Gaussian white noise process with variance $\sigma_\psi^2 > 0$. The next step is to find the vector of signal weights h and the variance of noise σ_ψ^2 that minimizes the loss function (3.26) subject to the information flow constraint (3.22). Maćkowiak, Matějka, and Wiederholt (2018) show that this problem can be solved using the Kalman filter with an information cost constraint. Specifically, they show the problem of finding the vector of signal weights reduces to:

$$\min_{h\in\mathbb{R}^{\max\{1,p\}+q},\sigma_\psi^2>0} (1\ 0\ \ldots\ 0)\ \Sigma_0 \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \tag{3.33}$$

where the conditional variance-covariance matrices of the state vector, Σ_1 and Σ_0, are given by

$$\Sigma_1 = F\Sigma_0 F' + Q \tag{3.34}$$

$$\Sigma_0 = \Sigma_1 - \frac{(1-2^{-2\kappa})}{h'\Sigma_1 h}\Sigma_1 hh'\Sigma_1 \tag{3.35}$$

where Q denotes the variance-covariance matrix of v_{t+1}. $\Sigma_1 \equiv \lim_{t\to\infty}\Sigma_{t|t-1}$ and $\Sigma_0 \equiv \lim_{t\to\infty}\Sigma_{t|t}$, where $\Sigma_{t|t-1}$ denote the conditional variance-covariance matrix ξ_t given $\mathcal{I}_{t-1}$. Note that we have used the fact that in the case of the one-dimensional signal (3.32) the information flow constraint (3.24) can be rewritten as $\frac{1}{2}\log_2\left(\frac{h'\Sigma_1 h}{\sigma_\psi^2}+1\right) \leq \kappa$ and it is always binding, in order to obtain the presented formulation of a problem.

Solving the problem 3.33-3.35 yields the vector of optimal signal weights h. The problem 3.33 - 3.35 can be solved by a Kalman filter. Since multiplying the signal 3.31 by a non-zero constant does not change the matrices Σ_1 and Σ_0, it is helpful to normalize either an element of h or σ_ψ^2 before solving the problem 3.33-3.35.

3.3.4 Optimal signal selection

The final part to be specified for the rational inattention model is the formulation of the optimal action that firm i is tracking, that is to specify the profit-maximizing price $\hat{p}_{it}^*$ as a function of the state variables. In order to manifest implications of the rational inattention on the optimal signal selections and price-setting behavior of firms and consequential real effects of nominal shocks we investigate two cases with a different degree of strategic complementarity.

No strategic complementarity case

In the case when there is no strategic complementarity in price setting ($\lambda = 1$), the profit-maximizing price has the following form (see Appendix 3.C for derivation)

$$\hat{p}_{it}^{*} = m_t^{(1)}(i) - a_t^{(1)}(i) - \sum_{j=0}^{\infty}\left(m_{t|t}^{j+2} - a_{t|t}^{j+2}\right) - \mathbb{E}_{i,t}\frac{1}{\gamma+1}\eta_{i,t} \tag{3.36}$$

The profit-maximizing price is then dependent on the two state variables and the higher-order beliefs. In this case we can directly use the rational inattention Kalman filter to find the optimal signal. This case helps us to build an intuition that firms aim is to learn about the difference between the state variables and not necessarily about their exact values. Such indexation is one of the main reasons behind the sluggish behavior observed in the RIM's impulse response functions. Next we consider the case with a positive degree of the strategic complementarity, which slightly complicates the solution method.

Strategic complementarity case

In the case with strategic complementarity in the price-setting between the firms , i.e. $0 < \lambda < 1$, the profit-maximizing price depends also on an endogenous variable, the price level. Formally, the profit-maximizing price is (see Appendix 3.C for derivation)

$$\hat{p}_{it}^{*} = (1-\lambda)\left[\mathbb{E}_{it}\ln P_t - \sum_{j=0}^{\infty}(1-\lambda)^j\lambda\left(m_{t|t}^{j+2} - a_{t|t}^{j+2}\right) + \ln\bar{y} - \mu_M t + \mu_A t\right] - \mathbb{E}_{i,t}\frac{1}{\gamma+1}\eta_{i,t} \tag{3.37}$$

Such dependence complicates the solution method. We apply a guess and verify method. The solution algorithm is described schematically below.

ALGORITHM: **Solution of the model with the endogenous variable**
Step 1: We guess that the profit-maximizing price follows an ARMA(p,q) process
Step 2: Given the guess, we apply the analytical results from section 3.3.3 to establish the form of an optimal signal and we compute the optimal signal weights and the implied actions.

Step 3: We then calculate the price level and the actual law of motion for the profit-maximizing price.

Step 4: If the actual law of motion for the profit-maximizing price differs from our guess, we update the guess and go to **Step 1**; otherwise a fixed point is reached.

The rational inattention model shares with the ICKM the transition equation specified in the next section. This concludes the model. The key challenge for estimation of the rational inattention macroeconomic models is to solve the model sufficiently fast. We implement the solution method introduced in MATLAB using the Artelys Knitro[2] solver. The computing time needed for solving the model depends on various variables, especially the degree of strategic complementarity in price setting. The code was run on at 2,4 GHz Intel Core i5, 16GB DDR3, Inter Iris 1536 MB and the running time ranged from 14 seconds to 350 seconds. For comparison, with implementation of the usual solvopt solver the lower bound achieved was 400 seconds.

3.4 Imperfect common knowledge model

Following Melosi (2014), in the imperfect common knowledge model (ICKM), it is assumed that firms observe one idiosyncratic noisy signal about each exogenous state variable and face strategic complementarities in price setting with no cost of price adjustment. Thus, a signal of firm i is defined as:

$$S_{i,t}^2 = \begin{pmatrix} m_t \\ a_{i,t} \end{pmatrix} + \begin{pmatrix} \tilde{\sigma}_m & 0 \\ 0 & \tilde{\sigma}_a \end{pmatrix} \mathbf{e}_{i,t} \tag{3.38}$$

where $S_{i,t}^2 \equiv [S_{m,i,t}, S_{a,i,t}]'$; $m_t \equiv \ln M_t - \mu_M t$; $a_{i,t} \equiv \ln A_{i,t} - \mu_A t$; $\mathbf{e}_{i,t} \equiv [e_{m,i,t}, e_{a,i,t}]'$ and $\mathbf{e} \overset{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_2)$. Note that m_t and $a_{i,t}$ represent two exogenous state variables of the model and the signal noises $e_{m,i,t}$ and $e_{a,i,t}$ are assumed to be independent and identically distributed across firms and time. In every period t, firms observe the history of their signals $S_{i,t}$ and choose their optimal price $P_{i,t}$ so as to maximize the objective function (3.10) subject to equations (3.11)-(3.15) and (3.38), the transition equation for the stochastic

[2]For a summary of the algorithms implemented in the Knitro solver see Byrd, Nocedal, and Waltz (2006)

discount factor Q_t, and that for the nominal wage W_t.

We solve the ICKM model by guessing and verifying the laws of motion for the vector of higher-order beliefs. As shown by Woodford (2002), under some conditions, the equilibrium dynamics of model variables can be expressed as a function of a weighted average of the (average) higher-order expectations about the exogenous state variables. This approach has several advantages. First, there is no need to trunctate the state vector to solve the model. Second, the dimensionality of the state vector often becomes very small, making the task of solving the model both quite fast and fairly accurate. Nevertheless, this method of solution is not applicable to every imperfect common knowledge model. Sufficient conditions for this solution method to be applicable are that the nominal output follows an exogenous process and no past and forward-looking endogenous variables enter the price-setting problem.[3] The former condition is satisfied in the model presented because of the equilibrium condition (3.9) and the exogeneity of the monetary rule (3.8). The latter condition is satisfied by abstracting from capital accumulation and assuming that firms do not pay interest on loans.

As mentioned earlier, following Woodford (2002), we solve the ICKM by guessing and verifying the law of motion for a finite number of weighted averages of higher-order expectations; that is, $\mathbf{F}_t \equiv \sum_{j=1}^{\infty}(1-\lambda)^{j-1}\lambda \mathbf{X}_t^{(j)}$, where $\mathbf{X}_t^{(j)}$ denotes the vector of average j-th order expectations about the exogenous state variables $\mathbf{X}_t \equiv [m_t, m_{t-1}, a_t]'$. The transition equations of the ICKM can be shown to be:

$$\bar{\mathbf{X}}_t = \bar{\mathbf{B}}\bar{\mathbf{X}}_{t-1} + \bar{\mathbf{b}}\mathbf{u}_t, \tag{3.39}$$

where we denote $\bar{\mathbf{X}}_t \equiv [\mathbf{X}_t'; \mathbf{F}_t']'$; $\bar{B} \equiv \begin{pmatrix} \mathbf{B}_{3x3} & \mathbf{0}_{3x3} \\ \mathbf{G}_{3x3} & \mathbf{H}_{3x3} \end{pmatrix}$; $\bar{\mathbf{b}} = [\mathbf{b}'; \mathbf{d}']$; $\mathbf{u}_t = [\epsilon_{m,t}, \epsilon_{a,t}]'$ with $\mathbf{u}_t \overset{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \Sigma_u)$, for all t and $\Sigma_u = \begin{pmatrix} \sigma_m^2 & 0 \\ 0 & \sigma_a^2 \end{pmatrix}$; the matrices $\mathbf{B}$ and $\mathbf{b}$ are given by the

[3] Paciello and Wiederholt (2014) show that this solution method also applies when one introduces dependence of the money supply on real shocks in a specific way and the utility function is restricted to being logarithmic.

exogenous processes (3.8) and (3.12); and $\mathbf{G}$, $\mathbf{H}$ and $\mathbf{d}$ are matrices that are yet to be determined. The evolution of the log-deviations of the (stationary) real output y_t and price level p_t from their steady-state value can be expressed as linear functions of the state vector $\bar{\mathbf{X}}_t$. Thus, we obtain an equilibrium of the ICKM by characterizing these three matrices: $\mathbf{G}$, $\mathbf{H}$ and $\mathbf{d}$. Further, these three matrices are a function of the model parameters: $\rho_m, \sigma_m, \sigma_a, \tilde{\sigma}_m, \tilde{\sigma}_a, \lambda$ and the Kalman-gain matrix associated with firms' signal extraction problem. For the full description of the ICKM solution see Appendix 3.D.

3.5 Calvo pricing model

In the Calvo model firms and all the other agents (i.e., households, the financial intermediary and the monetary authority) perfectly observe the past and current realizations of the model variables. The key assumption is that with exogenously given probability each firm i can reoptimize the price in a period t. This probability is independent of the state of the model and of the last time the firm reoptimized the price. Specifically, only a fraction $(1 - \theta_p)$ of firms reoptimize their prices, while the remaining θ_p fraction adjusts them to a geometric weighted average of the steady-state rate of inflation π^* and of the last period's inflation rate π_{t-1} with weights $(1-\omega)$ and ω, respectively. Thus, the degree of price indexation is captured by a parameter ω. After detrending the nonstanionary variables and log-linearizing the model around the deterministic steady state one can obtain the standard New Keynesian Philips curve (for its derivation see e.g. Section 2.2 in Woodford (2003)):

$$\hat{\pi}_t = \frac{\omega}{1+\omega\beta}\hat{\pi}_{t-1} + \frac{\kappa_{pc}}{1+\omega\beta}\hat{y}_t + \frac{\beta}{1+\omega\beta}\mathbb{E}_t\hat{\pi}_{t+1} \tag{3.40}$$

where $\kappa_{pc} = (1-\theta_p)(1-\theta_p\beta)\lambda/\theta_p$ with $\lambda \equiv (\gamma+1)\phi^{-1}/[\nu(\phi^{-1}-1)+1]$.

3.6 Empirical Analysis

In this section, we present the Bayesian estimation of all models presented. Section 3.6.1 introduces the dataset used, followed by the VAR analysis in section 3.6.2. Sets of identifiable parameters in various models and their prior and posterior distributions are presented in 3.6.3. Finally, section 3.6.4 displays a comparison of impulse response functions, as well as of the marginal data densities of various models, to show the goodness of fit for different models.

3.6.1 Data

For estimation we adopt the same dataset as in Melosi (2014) in order to allow for a direct comparison with the results previously obtained about the ICKM and Calvo model. We use quarterly data that cover the third quarter of 1954 through the fourth quarter of 2011. We fit two time series of the U.S. GDP deflator and the U.S. per capita real GDP, obtained by dividing the nominal GDP by the civilian non-institutional population aged 16 years and older and deflating using the chained-price GDP deflator. The data on the U.S. GDP deflator and on the U.S. per capita real GDP are denoted as $\{P_t, t = 1, 2, \ldots T\}$ and $\{Y_t, t = 1, 2, \ldots, T\}$, respectively.The measurement equation for real GDP per capita and the GDP price deflator in the log-linearized RIM and ICKM model are the same and given by equations (3.18) and (3.17); and for the Calvo model they are standard and hence omitted. The models are estimated using unfiltered data.

3.6.2 VAR analysis

We fit a VAR with four lags to the data set. We follow Sims and Zha (1998) to specify the prior distribution for the VAR parameters. We obtain 100 000 posterior draws through the Gibbs sampler. We then compute the IRFs of output and inflation to monetary shocks in the VAR and compare these IRFs with those implied by the other three DSGE models. In this comparison, the IRFs implied by the VAR are used as the benchmark. In section 3.6.4, we show that this comparison is astute, from a Bayesian perspective, because the VAR attains a larger posterior probability than that of the RIM, ICKM and the Calvo model (see e.g., Schorfheide (2000)).

3.6.3 Prior and posterior distributions

After the log-linearization, the set of identifiable parameters in the RIM, ICKM and Calvo models, respectively, are

$$\theta_{RIM} = (\rho_m, \sigma_m, \sigma_a, \mu_M, \mu_A, \lambda) \tag{3.41}$$

$$\theta_{ICKM} = (\rho_m, \sigma_m, \sigma_a, \tilde{\sigma}_m^2, \tilde{\sigma}_a^2, \mu_M, \mu_A, \lambda) \tag{3.42}$$

$$\theta_{Calvo} = (\rho_m, \sigma_m, \sigma_a, \mu_M, \mu_A, \beta, \kappa_{pc}, \omega) \tag{3.43}$$

First, we discuss how the observables in the RIM and ICKM differ from each other. We can see directly that the sets of observables for these models do not include the discount factor β and α that dropped out during the log-linearization. The technology parameter ϕ, the inverse of Frisch labor supply elasticity γ and the demand elasticity ν are not separately identified. We also do not estimate the standard deviation of the firm-specific technology σ_η, because we do not use any firm-level data for estimation. The main difference between the RIM and ICKM sets of observables lies in the fact that the ICKM has an exogeneously fixed signal form of two independent signals for each state variable and corresponding noises $\tilde{\sigma}_m^2$ and $\tilde{\sigma}_a^2$. The RIM on the other hand uses the optimal signals given the information flow constraint, so for each selection of parameter specified in equation 3.41 it might select different a ARMA representation of the optimal signal. We specify the selected optimal signal in section 3.6.4 when discussing the IRFs' comparison.

In setting up the prior statistics of all the parameters we follow the approach of Melosi (2014); thus, we just briefly summarize it below. The prior for autoregressive parameter ρ_m and the standard deviation of the monetary shock σ_m is determined by estimating it using the presample observation from the first quarter of 1949 through the second quarter of 1954. This is possible thanks to the market clearing condition for the monetary market. The prior mean for the standard deviation of the aggregate technology is centered at 0.007 as is common in the real business cycle literature (see e.g., Kydland and Prescott (1982)). As we will see, the crucial parameter for the analyses presented is the strategic complementarity parameter $(1-\lambda)$. In the ICKM, if the technology parameter ϕ is set to 0.65 (Cooley and Prescott 1995) and the Frisch labor-supply elasticity γ is set to 0.5 (Fuentes-Albero et al. 2009), then the prior median of $(1-\lambda)$ is associated with the net markup $(\nu-1)^{-1}$ in the range from 5 to 23 percent.

The parameters for the RIM and ICKM differ mainly in the signal structure. For the ICKM the priors for standard deviations of signal noises $\tilde{\sigma}_m^2$, $\tilde{\sigma}_a^2$; and the standard deviation of the signal noise in the RIM σ_ψ^2 are set up such that it is ensured that signals are quite informative about the business cycle frequency variations of aggregate variables. To investigate direct comparison of the ICKM model with the RIM and thus how restric-

	RIM			ICKM			Calvo			Prior		
Name	Median	5%	95%	Median	5%	95%	Median	5%	95%	Median	5%	95%
λ	0.19	0.07	0.35	0.31	0.12	0.50	-	-	-	0.41	0.21	0.60
ρ_m	0.43	0.34	0.51	0.40	0.30	0.49	0.30	0.22	0.38	0.5	0.17	0.82
$100\sigma_m$	0.88	0.79	0.93	0.89	0.82	0.95	0.90	0.83	0.97	2.00	0.43	12.81
$100\sigma_a$	0.85	0.74	1	0.87	0.71	1.03	0.90	0.83	0.97	0.70	0.51	0.87
$100\mu_M$	1.25	1.10	1.43	1.25	1.09	1.42	1.25	1.11	1.39	0.00	-41.00	41.00
$100\mu_A$	0.43	0.24	0.59	0.43	0.33	0.53	0.42	0.10	0.73	0.00	-41.00	41.00
$100\tilde{\sigma}_m$	-	-	-	10.02	5.48	14.35	-	-	-	5.01	2.12	7.91
$100\tilde{\sigma}_a$	-	-	-	1.35	0.69	1.98	-	-	-	1.06	0.24	1.87
$100\kappa_{pc}$	-	-	-	-	-	-	1.09	1.08	1.09	12.00	0.00	22.00
ω	-	-	-	-	-	-	0.02	0.00	0.05	0.50	0.08	1.00

Table 3.1: Posterior and prior statistics for the parameters of the RIM, ICKM and Calvo model. Results are obtained by the Metropolis-Hastings algorithm with one million posterior draws.

tions on the signal structure matter, we assume that the information flow is equal in both models and thus we do not estimate and identify κ in RIM. In the Calvo model we can identify the slope of the New Keynesian Philips curve κ_{pc} and we set its prior belief such that it ranges from 0.00 to 0.22 (see e.g., Schorfheide (2008)). For indexation parameter ω we set a broad prior. The deterministic discount factor β and the Calvo parameter θ_p cannot be indentified separately.

We evaluate the posterior distributions numerically through the randomwalk Metropolis-Hastings algorithm, by generating one million draws. The prior and posterior statistics are shown in table 3.1. The majority of the parameters that are shared among the models have quite similar posterior medians. The main parameter that differs between the RIM and the ICKM is the degree of strategic complementarity $(1 - \lambda)$. We have assumed the prior median of strategic complementarity to be 0.59 and Bayesian inference significantly increased the posterior median both for the RIM and ICKM. Moreover, the estimated degree of strategic complementarity for the RIM is undoubtedly bigger than in the ICKM, i.e. 0.81 in comparison with 0.69, respectively. The estimated number is very close to a strategic complementarity of 0.82 obtained from a micro-level survey from New Zealand and reported in Afrouzi (2020). The RIM estimated degree of strategic complementarity is also closer than other models to the usual value of 0.9 used for calibration of the U.S. economy (see e.g., Mankiw and Reis (2002),Woodford (2003)).

3.6.4 Results comparison

In this section we present a Marginal data density comparison of the models and the comparison based on the impulse response functions. The former is necessary for us to establish that the VAR attains larger posterior probability than the other models. Otherwise, it would not be meaningful to use the VAR as a benchmark model for the IRFs' comparison. The latter provides us with information on how different models capture the hump-shaped behavior of the IRFs.

Marginal data density comparison. The posterior probability of model M_i, where $i \in M := \{\text{RIM, ICKM, Calvo, VAR}\}$ is given by

$$\mathcal{P}_{t,M_i} = \frac{\mathcal{P}_{0,M_i} \cdot P(Y|M_i)}{\sum_{i \in M} \mathcal{P}_{0,M_i} \cdot P(Y|M_i)}$$

where $\mathcal{P}_{0,M_i}$ is the prior probability of the model M_i and are assumed to be the same across models, $\mathcal{P}_{0,M_i} = 1/3$. The dataset is denoted as Y, the marginal data density(MDD) of a model M_i with a vector of parameters Θ_i is defined as $P(Y|M_i) = \int \mathcal{L}(\Theta_i|Y, M_i)P(\Theta_i|M_i)d\Theta_i$, where $P(\Theta_i|M_i)$ is the prior distribution, $\mathcal{L}$ is the likelihood function[4]. Table 3.2 shows the log marginal data densities for all the models. The VAR proves to have the largest posterior probability and the RIM dominates the ICKM in the MDD comparison.

VAR	RIM	ICKM	Calvo
1920.05	1802.71	1737.88	1706.02

Table 3.2: Log of the marginal data densities for all the models.

Impulse response function comparison. Figures 3.1 and 3.2 depict the IRFs of real GDP and inflation to an unexpected monetary shock. To obtain these IRFs we use the full-information method. We apply the restriction that the monetary policy has no long-run real effects (see e.g., Blanchard and Quah (1989)) for identification of the monetary

[4] Chib's method is used for the VAR model. For all the other models we apply Geweke (1999) harmonic mean estimator.

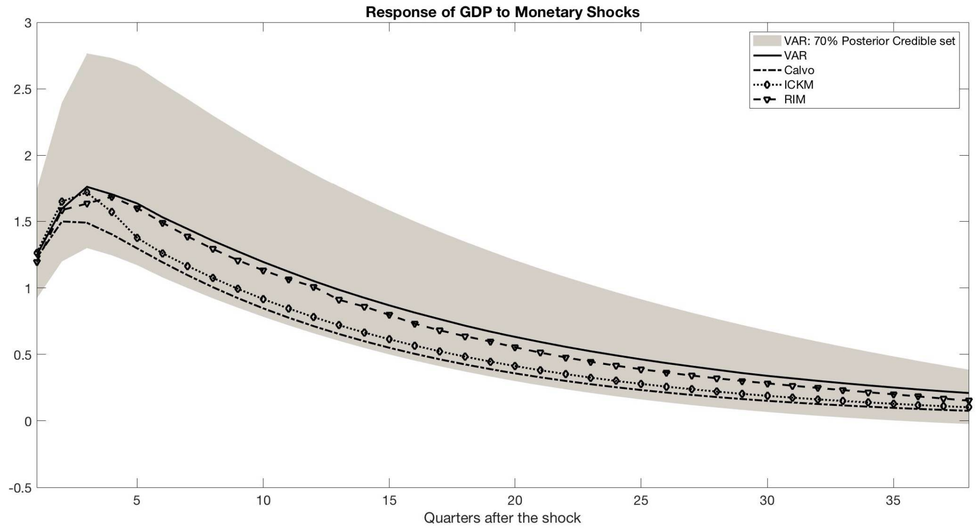

Figure 3.1: IRF of real GDP to a monetary shock for the RIM, ICKM, Calvo and VAR.

shock in the case of VAR. The benchmark VAR-IRFs present a hump-shaped behavior, with short run sluggish momentum (with maximum achieved 3 quarters after the monetary shock for a real GDP reaction) with a long-term persistence. Both the RIM and ICKM prove to be capturing this hump-shaped behavior much more closely than the Calvo model. The RIM mainly captures the persistence much more closely than other models. However, the ICKM outperforms the RIM in timely and more appropriate mimicking of the short-run momentum. These differences are caused mainly by i) optimal signal selection given the information flow constraint and ii) the different estimated value of the strategic complementarity.

AR	MA	$100\sigma_\psi^2$
3.0122	-0.1541	7.0437
-3.2928	0.3646	
1.5345	-0.2816	
-0.2538	0.0810	

Table 3.3: Parameters for the selected signal in RIM

For the estimated parameter values in the RIM the selected signal has an ARMA(4,4) form and the parameters are presented in table 3.3. From the form of the signal it is clear that it differs significantly from the assumed AR(1) form in the ICKM models following

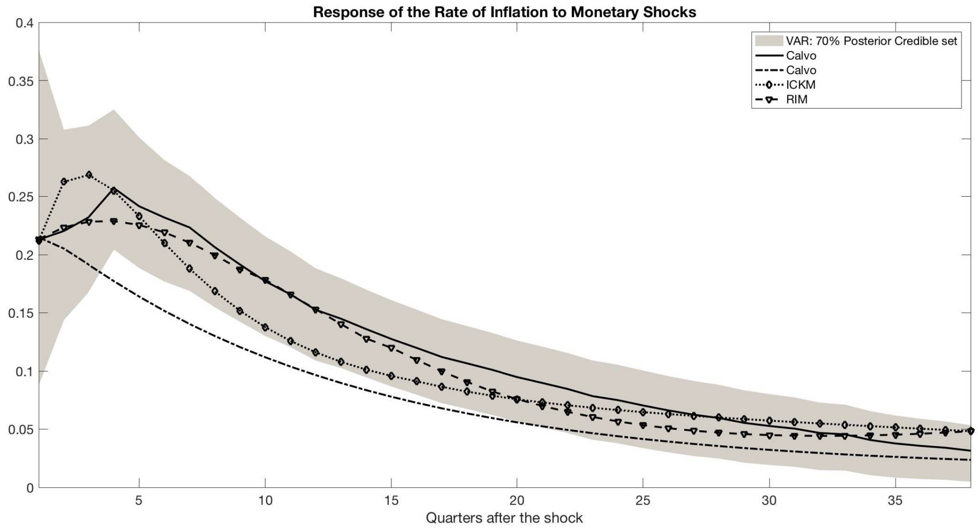

Figure 3.2: IRF of inflation to a monetary shock for the RIM, ICKM, Calvo and VAR.

the tradition of Woodford (2002). The rationally inattentive firms are selecting signals that at the same time provide them with information about the current optimal action as well as the future optimal action. Also, firms selected one signal that is a construct index of the state variables instead of obtaining two independent noisy signals for each state variable. Such signal selection occasions that sluggish behavior is not only caused by imperfect information due to obtaining noisy signals, but also because from the selected signal form it is more complicated to understand what exactly is causing a variation.

The larger estimated degree of strategic complementarity for the rational inattention model has a significant impact on the persistence as it magnifies the mechanism of shock propagation. This can be observed from equations (3.17) and (3.18) where the persistence of response to a monetary shock depends on the weighted averages of higher-order expectations about the state variables $\sum_{j=0}^{\infty}(1-\lambda)^j\lambda\left(m_{t|t}^{(j+1)} - a_{t|t}^{(j+1)}\right)$ and $\sum_{j=0}^{\infty}(1-\lambda)^j\lambda m_{t|t}^{(j+1)}$, respectively. Here we can notice the effect of the strategic complementarity as the weights associated with the average expectations. The larger the degree of the strategic complementarity the more it matters for price setting decisions and more sluggishly the higher order expectations adjust after shocks.

3.7 Concluding Remarks

In this paper, we conduct full-information Bayesian inference and comparison of the Rational inattention model, Imperfect common knowledge model and Calvo pricing model. To our knowledge, this is the first paper that directly compares the performance of the estimated RIM model with full flexibility in the signal structure and the ICKM model. We show that the RIM captures data better than the Calvo model. In comparison with the ICKM, the RIM reproduces the persistence more closely; however, in the capturing of the aftershock momentum, the ICKM seems to be more successful. The difference between the models' performance is caused mainly by the fact that the optimal signal is an index about the state variables with more lags, as is assumed in the ICKM, and by the fact that we estimated a large degree of the strategic complementarity in the RIM.

To reveal the role of rational inattention, we build upon and keep as close as possible to the model presented in Melosi (2014) that introduced the estimation of the ICKM model. However, the given model is still stylized in several dimensions to provide tractability. This allowed us to conduct the next important step and estimate a version of the model presented in Melosi (2014) with inattentive firms. Previous rational inattention models without signal structure restrictions were only calibrated because dynamic rational inattention models are difficult to solve. In this paper, we are building on the results presented in Maćkowiak, Matějka, and Wiederholt (2018) which allows us to overcome the key challenge to solve the model sufficiently fast so that estimation becomes feasible in a reasonable amount of time.

There are still several unsolved challenges. As was advocated by Mackowiak and Wiederholt (2015), it is crucial to estimate a model that would have not only inattentive firms, but also households. That, however, presents an additional computational challenge. Because our aim was to investigate the direct comparison with the ICKM, we have assumed in the RIM the same information flow as in the ICKM. This assumption can be relaxed and direct estimation of the cost of information in a macroeconomic setting might be very insightful.

3.A Appendix A: Derivation of the log-linearized firm's price-setting equation

The first-order necessary condition (w.r.t. $P_{i,t}$) of the price-setting in the RIM and ICKM is:

$$\mathbb{E}_{i,t}\left[\beta Q_t\left(Y_{i,t}-\nu P_{i,t}\left(\frac{P_{i,t}}{P_t}\right)^{-\nu-1}\frac{Y_t}{P_t}+\nu\phi^{-1}\frac{W_t}{A_{i,t}}\left(\frac{Y_{i,t}}{A_{i,t}}\right)^{\phi^{-1}-1}\left(\frac{P_{i,t}}{P_t}\right)^{-\nu-1}\frac{Y_t}{P_t}\right)\right]=0$$

Using $C_t=\left(\int_0^1 C_{i,t}^{\frac{\nu-1}{\nu}}di\right)^{\frac{\nu}{\nu-1}}$ we can write

$$\mathbb{E}_{i,t}\left[\beta Q_t\left(\left(\frac{P_{i,t}}{P_t}\right)^{-\nu}Y_t-\nu P_{it}\left(\frac{P_{i,t}}{P_t}\right)^{-\nu-1}\frac{Y_t}{P_t}\right)\right]+$$
$$+\mathbb{E}_{i,t}\left[\beta Q_t\nu\phi^{-1}\frac{W_t}{A_{i,t}}\left(\frac{1}{A_{i,t}}\left(\frac{P_{i,t}}{P_t}\right)^{-\nu}Y_t\right)^{\phi^{-1}-1}\left(\frac{P_{i,t}}{P_t}\right)^{-\nu-1}\frac{Y_t}{P_t}\right]=0$$

From the representative household's problem we know that labor supply is $W_t/P_t=\alpha Y_t H_t^{\gamma}$ after substituting this result we obtain

$$\mathbb{E}_{i,t}\left[\beta Q_t(1-\nu)\left(\frac{P_{i,t}}{P_t}\right)^{-\nu}Y_t\right]+$$
$$+\mathbb{E}_{i,t}\left[\beta Q_t\nu\phi^{-1}\frac{\alpha Y_t H_t^{\gamma}}{A_t e^{\eta_{i,t}}}\left(\frac{1}{A_{i,t}}\left(\frac{P_{i,t}}{P_t}\right)^{-\nu}Y_t\right)^{\phi^{-1}-1}\left(\frac{P_{i,t}}{P_t}\right)^{-\nu-1}Y_t\right]=0$$

Define the stationary variables

$$y_t\equiv\frac{Y_t}{A_t},\ y_{i,t}\equiv\frac{Y_{i,t}}{A_t},\ p_{i,t}=\frac{P_{i,t}}{P_t},\ h_{i,t}=H_{i,t}$$

Afterwards we can rewrite it to

$$\mathbb{E}_{i,t}\left[\beta q_t p_{i,t}^{-\nu} y_t \left((1-\nu) + \nu\phi^{-1}\frac{\alpha y_t h_t^{\gamma}}{e^{\eta_{i,t}}}\left(\frac{1}{e^{\eta_{i,t}}}(p_{i,t})^{-\nu} y_t\right)^{\phi^{-1}-1}(p_{i,t})^{-1}\right)\right]$$

We take the log-linear approximation of the equation above around the deterministic symmetric steady-state. We do not need to take care of what is outside the round brackets because the expression within them is zero at the deterministic symmetric steady-state.

The price-setting condition can be approximated as:

$$\mathbb{E}_{i,t}\left[\gamma\hat{h}_t - [\nu(\phi^{-1}-1)+1]\hat{p}_{i,t} + \phi^{-1}(\hat{y}_t - \eta_{i,t})\right] = 0$$

From the production function we know that $\hat{h}_{i,t} = \phi^{-1}(\hat{y}_{i,t} - \eta_{i,t})$ and using that $\hat{y}_t = \int \hat{y}_{i,t} di$ we get that $\hat{h}_t = \phi^{-1}\hat{y}_t$. By substituting those results into the equation above, we obtain:

$$\mathbb{E}_{i,t}\left[\gamma\phi^{-1}\hat{y}_t + \phi^{-1}(\hat{y}_t - \eta_{i,t}) - [\nu(\phi^{-1}-1)+1]\hat{p}_{i,t}\right] = 0$$

then

$$[\nu(\phi^{-1}-1)+1]\mathbb{E}_{i,t}\hat{p}_{i,t} = \mathbb{E}_{i,t}[\gamma\phi^{-1}\hat{y}_t + \phi^{-1}(\hat{y}_t - \eta_{i,t})]$$

$$\mathbb{E}_{i,t}\hat{p}_{i,t} = \frac{\mathbb{E}_{i,t}[(\gamma+1)\phi^{-1}\hat{y}_t - \phi^{-1}\eta_{i,t}]}{[\nu(\phi^{-1}-1)+1]}$$

$$\mathbb{E}_{i,t}\hat{p}_{i,t} = \frac{(\gamma+1)\phi^{-1}}{\nu(\phi^{-1}-1)+1}\mathbb{E}_{i,t}\hat{y}_t - \frac{\phi^{-1}}{\nu(\phi^{-1}-1)+1}\mathbb{E}_{i,t}\eta_{i,t}$$

If we define $\lambda \equiv (\gamma+1)\phi^{-1}/[\nu(\phi^{-1}-1)+1]$ then

$$\mathbb{E}_{i,t}\hat{p}_{i,t} = \lambda\mathbb{E}_{i,t}\hat{y}_t - \frac{\lambda}{\gamma+1}\mathbb{E}_{i,t}\eta_{i,t}$$

this is

$$\mathbb{E}_{i,t}\left[\ln P_{i,t} - \ln P_t\right] = \lambda\mathbb{E}_{i,t}\left[\ln Y_t - \ln A_t - \ln\bar{y}\right] - \frac{\lambda}{\gamma+1}\mathbb{E}_{i,t}\eta_{i,t}$$

equivalently

$$\ln P_{i,t} = \mathbb{E}_{i,t}[\lambda\ln Y_t + \ln P_t - \lambda\ln A_t] - \frac{\lambda}{\gamma+1}\mathbb{E}_{i,t}\eta_{i,t} - \lambda\ln\bar{y}$$

From eq. 3.9 we have

$$\ln P_t + \ln Y_t = \ln M_t \implies \ln Y_t = \ln M_t - \ln P_t$$

after plugging it in

$$\ln P_{i,t} = \mathbb{E}_{i,t}\left[(1-\lambda)\ln P_t + \lambda\ln M_t - \lambda\ln A_t - \frac{\lambda}{\gamma+1}\eta_{i,t}\right] - \lambda\ln\bar{y}$$

In the main text we define $m_t \equiv \ln M_t - \mu_M t$ and $a_t \equiv \ln A_t - \mu_A t$ thus we obtain

$$\ln P^*_{i,t} = \mathbb{E}\left[(1-\lambda)\ln P_t + \lambda m_t - \lambda a_t - \frac{\lambda}{\gamma+1}\eta_{i,t}|z_i^t\right] - \lambda[(\mu_A - \mu_M)t + \ln\bar{y}], \qquad (3.44)$$

Further, after log-linearization of equation 3.4 around the deterministic steady-state, we obtain $\hat{p}_t = \int \hat{p}_{i,t}di$. Thus, when we integrate equation 3.44 across firms we get

$$\ln P_t = (1-\lambda)\ln P^{(j+1)}_{t|t} + \lambda\ln\ln M^{(j+1)}_{t|t} - \lambda\ln\ln A^{(j+1)}_{t|t} - \lambda\ln\bar{y}$$

for $j \in \{1, 2, \dots\}$. After repeatedly substituting these results into the average-price equation and substitution, we can obtain

$$\ln P_t = \left[\sum_{j=0}^{\infty}(1-\lambda)^j\lambda\left(m_{t|t}^{(j+1)} - a_{t|t}^{(j+1)}\right)\right] - \lambda\ln\bar{y} + \mu_M\cdot t - \mu_A\cdot t. \tag{3.45}$$

Using $\ln M_t = \ln Y_t + \ln P_t$ we receive

$$\ln Y_t = \left[m_t - \sum_{j=0}^{\infty}(1-\lambda)^j\lambda m_{t|t}^{(j+1)}\right] + \sum_{j=0}^{\infty}(1-\lambda)^j\lambda a_{t|t}^{(j+1)} + \lambda\ln\bar{y} + \mu_A\cdot t \tag{3.46}$$

Equations 3.45 and 3.46 are the law of motion for price level and output, respectively.

3.B Appendix B: Analogous objective of rational inattention model

The objective function of the model with rationally inattentive firms is

$$\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^t\hat{\pi}_t(\hat{p}_{i,t},\hat{p}_t,\hat{y}_t,\hat{q}_t)\right]$$

where $\hat{\pi}_t$ is the log-quadratic approximation of $Q_t\pi_t$. We take Q_t - the stochastic discount factor as exogenously given for firms. Let's define the profit-maximizing price that solves the log-quadratic price-setting problem under perfect information as $\hat{p}^*_{i,t}$. We can show that

$$\hat{\pi}_t\left(\hat{p}^*_{i,t},\hat{p}_t,\hat{y}_t,\hat{q}_t\right)-\hat{\pi}_t\left(\hat{p}_{i,t},\hat{p}_t,\hat{y}_t,\hat{q}_t\right)\propto\left(\hat{p}^*_{i,t}-\hat{p}_{i,t}\right)^2$$

is a function of structural parameters up to a constant

$$\hat{\pi}_t\left(\hat{p}^*_{i,t},\hat{p}_t,\hat{y}_t,\hat{q}_t\right)-\hat{\pi}_t\left(\hat{p}_{i,t},\hat{p}_t,\hat{y}_t,\hat{q}_t\right)\propto\left(\ln\left(P^*_{i,t}\right)-\ln\left(P_{i,t}\right)\right)^2$$

The first expression on the left side of the equation is not affected by the rational inattention problem as the profit-maximizing price is obtained for complete information. Therefore we can rewrite the objective as

$$\begin{aligned}\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^t\hat{\pi}_t(\hat{p}_{i,t},\hat{p}_t,\hat{y}_t,\hat{q}_t)\right]&\propto-\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^t\left[\hat{\pi}_t\left(\hat{p}^*_{i,t},\hat{p}_t,\hat{y}_t,\hat{q}_t\right)-\hat{\pi}_t\left(\hat{p}_{i,t},\hat{p}_t,\hat{y}_t,\hat{q}_t\right)\right]\right]\\&\propto-\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^t\left(\hat{p}^*_{i,t}-\hat{p}_{i,t}\right)^2\right]\end{aligned}\quad(3.47)$$

Because $\hat{p}^*_{i,t}$ and $\hat{p}_{i,t}$ are stationary processes, they do not depent on t.

$$\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^t\hat{\pi}_t(\hat{p}_{i,t},\hat{p}_t,\hat{y}_t,\hat{q}_t)\right]\propto-\mathbb{E}\left[\beta^t\left(\hat{p}^*_{i,t}-\hat{p}_{i,t}\right)^2\right]$$

The log-quadratic problem $\hat{p}_{i,t}$ is defined as $\ln(P_{i,t}/P_t)$. Hence

$$\hat{p}^*_{i,t} = \ln\left(P^*_{i,t}\right) - \ln\left(P_t\right)$$
$$\hat{p}_{i,t} = \ln\left(P_{i,t}\right) - \ln\left(P_t\right)$$

Then

$$\mathbb{E}\left[\sum_{t=1}^{\infty} \beta^t \hat{\pi}_t(\hat{p}_{i,t}, \hat{p}_t, \hat{y}_t, \hat{q}_t)\right] \propto -\mathbb{E}\left[\beta^t \left(\ln\left(P^*_{i,t}\right) - \ln\left(P_{i,t}\right)\right)^2\right]$$

3.C Appendix C: Profit-maximizing price for the rational inattention model

We want to derive $\hat{p}^*_{it}$ for two cases: i) with zero degree of the strategic complementarity and ii) with the degree of the strategic complementarity between 0 and 1. Note that the $\hat{p}^*_{it} = \ln P^*_{it} - \ln P_t$. Recall that $m_t = \ln M_t - \mu_M t$ and $a_t = \ln A_t - \mu_A t$ and then from the equation (3.16) the former ($\ln P^*_{it}$) can be shown that

$$\ln P_{it} = \sum_{j=0}^{\infty}(1-\lambda)^{j+1}\lambda\left(m^{j+2}_{t|t}(i) - a^{j+2}_{t|t}(i)\right) - \ln \bar{y} + \mu_M t - \mu_A t + \lambda m^{(1)}_t(i) - \lambda a^{(1)}_t(i) - \frac{\lambda}{\gamma+1}\mathbb{E}_{it}\eta_{it}$$

the latter ($\ln P_t$) is given by the equation (3.17). The equations (3.36) and (3.37) are then directly obtain as a difference of $\ln P^*_{it} - \ln P_t$ first for the case when $\lambda = 1$ and second for the case when $0 < \lambda < 1$.

3.D Appendix D: Solving the ICKM

This solution method was directly stated in the Appendix B of Melosi (2014) and it is generalization of the solution method presented in Woodford (2002) for two state variables. We introduce this solution method here in order to provide stand-alone feature of this paper and to highlight how the signal acquisition assumptions influence the solution of the model.

In order to find an equilibrium for the ICKM we have to characterize the equilibrium law of motion for the economy's aggregate states. The transition equations are

$$\hat{y}_t = -\hat{p}_t \tag{3.48}$$

$$\hat{p}_t = \mathbf{r}'\bar{\mathbf{X}}_\mathbf{t} \tag{3.49}$$

$$\bar{\mathbf{X}}_\mathbf{t} = \bar{\mathbf{B}}\bar{\mathbf{X}}_\mathbf{t-1} + \bar{\mathbf{b}}\mathbf{u}_\mathbf{t} \tag{3.50}$$

where $\hat{y}_t$ and $\hat{p}_t$ denote the log-linear deviations of the stationary output ($y_t = \frac{Y_t}{A_t}$) and price ($p_t = \frac{P_t A_t}{M_t}$), from their deterministic steady-state. Further, $\bar{\mathbf{X}}_t \equiv [\mathbf{X}_t'; \mathbf{F}_t']'$; $r = (-1, 0, 1, 1, 0, -1)'$;

$\bar{\mathbf{B}} \equiv \begin{pmatrix} \mathbf{B}_{3x3} & \mathbf{0}_{3x3} \\ \mathbf{G}_{3x3} & \mathbf{H}_{3x3} \end{pmatrix}$; $\bar{\mathbf{b}} = [\mathbf{b}'; \mathbf{d}']$; $\mathbf{B} \equiv \begin{pmatrix} 1+\rho_m & -\rho_m & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$; $\mathbf{b} \equiv \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{pmatrix}$; $\mathbf{u}_t =$ $[\epsilon_{m,t}, \epsilon_{a,t}]'$ with $\mathbf{u}_t \overset{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \Sigma_u)$, for all t and $\Sigma_u = \begin{pmatrix} \sigma_m^2 & 0 \\ 0 & \sigma_a^2 \end{pmatrix}$. The vector $\mathbf{F}_t$ is defined as

$$\mathbf{F}_t \equiv \sum_{j=1}^{\infty} (1-\lambda)^{j-1} \lambda \mathbf{X}_t^{(j)} \tag{3.51}$$

where

$$\mathbf{X}_t \equiv (m_t, m_{t-1}, a_t)' \tag{3.52}$$

matrices $\mathbf{G}$, $\mathbf{H}$ and $\mathbf{d}$ are yet to be determined in order to find an equilibrium. The transition equation (3.48) is obtained after detrending and log-linearizing of $P_tY_t = M_t$ the market clearing conditions in the money market. The transition equation (3.49) can be obtained from the law of motion for the price level (3.17) combining with $\hat{p}_t = \ln P_t + \ln A_t - \ln M_t - \ln \bar{p}$ and $\ln \bar{p} + \ln \bar{y} = 0$.

Signal structure for the ICKM can be written as

$$S_{it} = \mathbf{D}\bar{\mathbf{X}}_{\mathbf{t}} + e_{it} \tag{3.53}$$

where $\mathbf{D} \equiv (\mathbf{D_1}; \mathbf{0_{2x3}})$, $\mathbf{D_1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$, $\mathbf{e}_{it} \equiv (\epsilon_{m,i,t}, \epsilon_{a,i,i})'$, $\mathbf{e}_{it} \overset{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_e)$, for all t, and i, $\mathbf{\Sigma_e} = \begin{pmatrix} \tilde{\sigma}_m^2 & 0 \\ 0 & \tilde{\sigma}_a^2 \end{pmatrix}$.

We follow Woodford (2002) and use the method of undetermined coefficients to identify unknown matrices $\mathbf{G}$, $\mathbf{H}$ and $\mathbf{d}$. Specifically, we want to solve the fixed point problem that given the conjectured law of motion (3.50), optimal firms' behavior must exactly aggregate to the conjectured law of motion (3.50).

After we have introduced the transition equations and the signal structure, it follows (see, e.g., Chow (1975),Harvey (1989)) that the firm i's optimal estimate of the state vector evolves according to a Kalman filter equation

$$\bar{\mathbf{X}}_{\mathbf{t|t}}(i) = \bar{\mathbf{X}}_{\mathbf{t|t-1}}(i) + \mathbf{k}[S_{it} - \mathbf{D}\bar{\mathbf{X}}_{\mathbf{t|t-1}}(i)]$$

where $\mathbf{k}$ is $6x2$ Kalman gain matrix that has to be determined. After plugging the one step ahead forecast of the state vector into the Kalman filter equation, integrating over firms and applying that $\int \mathbf{e}_{it} di = \mathbf{0}$ and substituting for $\bar{\mathbf{X}}_{\mathbf{t}}$ we obtain the average estimates of the current state vector

$$\bar{\mathbf{X}}_{t|t} = [\mathbf{I} - \mathbf{kD}]\bar{\mathbf{B}}\mathbf{X_{t-1|t-1}} + \mathbf{kD}[\bar{\mathbf{B}}\mathbf{X}_{t-1} + \bar{\mathbf{b}}\mathbf{u_t}]. \tag{3.54}$$

Let us define the $6x3$ vector φ such that $\varphi = (\lambda \cdot \mathbf{I}_3; (1-\lambda)\cdot \mathbf{I}_3)'$. Then after some manipulation $\varphi'\bar{\mathbf{X}}_t^{(1)} = \mathbf{F}_t$. when we substitute there equation (3.54) we obtain

$$\mathbf{F}_t = [\varphi' - \tilde{\mathbf{k}}\mathbf{D}]\bar{\mathbf{B}}\mathbf{X_{t-1|t-1}} + \tilde{\mathbf{k}}\mathbf{D}[\bar{\mathbf{B}}\mathbf{X_{t-1}} + \bar{\mathbf{b}}\mathbf{u}_t] \tag{3.55}$$

where $\tilde{\mathbf{k}} \equiv \varphi'\mathbf{k}$. Following three equations can be derived

$$\varphi'\bar{\mathbf{B}} = (\lambda\mathbf{B} + (1-\lambda\mathbf{G}); (1-\lambda)\mathbf{H}) \tag{3.56}$$

$$\mathbf{D}\bar{\mathbf{B}} = \left(\mathbf{B}^\dagger; 0_{2x3}\right) \tag{3.57}$$

where $\mathbf{B}^\dagger = (\mathbf{B}_1'\mathbf{B}_3')$ and $\mathbf{B}_j$ is the j-th row of the matrix $\mathbf{B}$.

$$\mathbf{D}\bar{\mathbf{b}} = \mathbf{D_1 b} \tag{3.58}$$

Using equations (3.56) - (3.58), one can rewrite equation (3.55) as

$$\mathbf{F}_t = \left[\lambda\mathbf{B} + (1-\lambda)\mathbf{G} - \tilde{\mathbf{k}}\mathbf{B}^\dagger\right]\mathbf{X}_{t-1|t-1} + (1-\lambda)\mathbf{H}\mathbf{F}_{t-1|t-1} + \tilde{\mathbf{k}}\mathbf{B}^\dagger\mathbf{X_{t-1}} + \tilde{\mathbf{k}}\mathbf{u_t} \tag{3.59}$$

From the equation for φ and by lagging by one period we obtain $(1-\lambda)\cdot\mathbf{F}_{t-1|t-1} = \mathbf{F}_{t-1} - \lambda\mathbf{X}_{t-1|t-1}$ what after plugging into the equation (3.59) and rearranging we get

$$\mathbf{F}_t = \left[\lambda\mathbf{B} + (1-\lambda)\mathbf{G} - \tilde{\mathbf{k}}\mathbf{B}^\dagger - \lambda\mathbf{H}\right]\mathbf{X}_{t-1|t-1} + \mathbf{H}\cdot\mathbf{F}_{t-1} + \tilde{\mathbf{k}}\mathbf{B}^\dagger\mathbf{X_{t-1}} + \tilde{\mathbf{k}}\mathbf{u_t} \tag{3.60}$$

After comparison with (3.50), we can identify the matrices

$$\mathbf{G} = \tilde{\mathbf{k}}\mathbf{B}^{\dagger} \tag{3.61}$$

$$\mathbf{d} = \tilde{\mathbf{k}} \tag{3.62}$$

$$\left[\lambda\mathbf{B} + (1-\lambda)\mathbf{G} - \tilde{\mathbf{k}}\mathbf{B}^{\dagger} - \lambda\mathbf{H}\right] \stackrel{!}{=} 0 \tag{3.63}$$

After substituting (3.61) into (3.63), we can identify matrix **H** as

$$\mathbf{H} \stackrel{!}{=} \mathbf{B} - \tilde{\mathbf{k}}\mathbf{B}^{\dagger}. \tag{3.64}$$

The steady-state matrix of Kalman gain is given by

$$\mathbf{k} = \mathbf{P}\mathbf{D}'[\mathbf{D}\mathbf{P}\mathbf{D}' + \mathbf{\Sigma_e}]^{-1} \tag{3.65}$$

and the matrix **P** solves the algebraic Riccati equation

$$\mathbf{P} = \bar{\mathbf{B}}\left[\mathbf{P} - \mathbf{P}\mathbf{D}'[\mathbf{D}\mathbf{P}\mathbf{D}' + \mathbf{\Sigma_e}]^{-1}\mathbf{D}\mathbf{P}\right]\bar{\mathbf{B}}' + \bar{\mathbf{b}}\mathbf{\Sigma_u}\bar{\mathbf{b}}' \tag{3.66}$$

ALGORITHM:

Numerical loop to find out the fixed-point and determining P

Step 1: Given parameter values and a guess of the Kalman-gain $\mathbf{k^0}$, use equations (3.61) - (3.63) to characterize matrices **G**, **H** and **d**.

Step 2: Solve Riccati equation (3.66) and obtain **P** and a new Kalman-gain matrix **k** from (3.65).

Step 3: If the **k** and $\mathbf{k}^0$ are sufficiently close, the fixed point is found and stop. Otherwise set new $\mathbf{k}^0 = \mathbf{k}$ and repeat from **Step 1**.

Step 4: When fixed-point Kalman gain is found, the steady-state system of ICKM is fully characterized.

Bibliography

Afrouzi, H. 2020. "Strategic inattention, inflation dynamics, and the non-neutrality of money."

Ambuehl, S., and S. Li. 2018. "Belief updating and the demand for information." *Games and Economic Behavior* 109, no. 21-39.

Andersen, S., A. Di Girolamo, G.W. Harrison, and M.I. Lau. 2014. "Risk and Time Preferences of Entrepreneurs: Evidence from a Danish Field Experiment." *Theory and Decision* 77 (3): 341–357.

Andersen, S., G.W. Harrison, M.I. Lau, and E.E. Rutström. 2008. Chapter Risk Aversion in Game Shows of *Risk Aversion in Experiments*, edited by J. C. Cox and G. W. Harrison, Volume 12. Bingley, UK: Emerald, Research in Experimental Economics.

Anobile, G., G. M. Cicchini, and D. C Burr. 2012. "Linear Mapping of Numbers onto Space Requires Attention." *Cognition* 122 (3): 454–459.

Bandiera, O., A. Prat, L. Guiso, and R. Sadun. 2011. "Matching Firms, Managers and Incentives." *NBER Working Paper No. 16691.*

Barsky, R., F. Juster, M. Kimball, and M. Shapiro. 1997. "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Survey." *Quarterly Journal of Economics* 112 (2): 537–579.

Beaudry, P., and T. Willems. 2018. "On the macroeconomic consequences of over-optimism." *National Bureau of Economic Research*, vol. No. w24685.

Becker, G. M., M. H. DeGroot, and J. Marschak. 1964. "Measuring utility by a single-response sequential method." *Behavioral science* 9 (3): 226–232.

Bénabou, R. 2015. "The economics of motivated beliefs." *Revue d'économie politique* 125 (5): 665–685.

Blackwell, D., and L. Dubins. 1962. "Merging of Opinions with Increasing Information." *The Annals of Mathematical Statistics* 33 (3): 882–886.

Blanchard, O. J., and D. Quah. 1989. "The dynamic effects of aggregate demand and aggregate supply." *The American Economic Review* 79 (4): 655–673.

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

Bonatti, A., and J. Hörner. 2017. "Career Concerns with Exponential Learning." *Theoretical Economics* 12 (1): 425–475.

Byrd, R.H., J. Nocedal, and R.A. Waltz. 2006. Chapter KNITRO: An integrated package for nonlinear optimization of *Large-Scale Nonlinear Optimization*, edited by G. di Pillo and M. Roma, 35–59. Springer.

Cantillon, R. 1775. *Essai sur la nature du commerce en general.* London: Gyles.

Caplin, Andrew, and Mark Dean. 2015. "Revealed preference, rational inattention, and costly information acquisition." *American Economic Review* 105 (7): 2183–2203.

Chancelier, J.P., M. De Lara, and A. De Palma. 2009. "Risk Aversion in Expected Intertemporal Discounted Utilities Bandit Problems." *Theory and Decision* 67 (4): 433–440.

Charness, G., R. Oprea, and S. Yuksel. 2018. "How do people choose between biased information sources? Evidence from a laboratory experiment." *Working Paper.*

Chow, G.C. 1975. *Analysis and control of dynamic economic systems.* New York : Wiley.

Ciani, E., B. Etheridge, A. Delavande, and M. Francesconi. 2019. "Policy Uncertainty and Information Flows: Evidence from Pension Reform Expectations." *IZA Discussion Paper Series*, no. 12604 (September).

Coibon, O., Y. Gorodnichenko, and S. Kumar. 2015. "How Do Firms Form Their Expectations? New Survey Evidence." *(No. w21092). National Bureau of Economic Research.*

Cooley, T. F., and E. C. Prescott. 1995. Chapter Economic growth and business cycles of *Frontiers of business cycle research, 1*, edited by Thomas F. Cooley. Princeton:Princeton University Press.

Costa-Gomes, M. A., and G. Weizsäcker. 2008. "Stated beliefs and play in normal-form games." *The Review of Economic Studies* 75 (3): 729–762.

Cover, Thomas, and Joy Thomas. 2012. *Elements of information theory.* John Wiley & Sons.

Cramer, J., J. Hartog, N. Jonker, and C. Van Praag. 2002. "Low Risk Aversion Encourages the Choice for Entrepreneurship: An Empirical Test of a Truism." *Journal of Economic Behavior and Organization* 48 (1): 29–36.

De Roos, N., and Y. Sarafidis. 2009. "Decision making under risk in Deal or No Deal." *Journal of Applied Econometrics* 25 (6): 987–1027.

Dixit, A.K., and J.W. Weibull. 2007. "Political polarization." *Proceedings of the National Academy of Sciences* 104 (18): 7351–7356.

Fuentes-Albero, C., M. Kryshko, J. V. Ríos-Rull, R. Santaeulàlia-Llopis, and F. Schorfheide. 2009. "Methods versus substance: Measuring the effects of technology shocks on hours." *NBER Working Paper 15375.*

Gentzkow, M., J.M. Shapiro, and M. Taddy. 2016. "Measuring Polarization in High-Dimensional Data: Method and Application to Congressional Speech." NBER Working Paper Series.

Gerber, A., and D. Green. 1999. "Misperceptions about perceptual bias." *Annual Review of Political Science* 2 (1): 189–210.

Geweke, J. 1999. "Using simulation methods for Bayesian econometric models: inference, development, and communication." *Econometric reviews* 18 (1): 1–73.

Hamilton, B.H. 2000. "Does Entrepreneurship Pay? An Empirical Analysis of the Returns to Self-Employment." *Journal of Political Economy* 108 (3): 604–631.

Harvey, A.C. 1989. *Forecasting, structural time series models, and the Kalman filter.* Cambridge University Press.

Heidhues, P., S. Rady, and P. Strack. 2015. "Strategic Experimentation with Private Payoffs." *Journal of Economic Theory* 159:531–551.

Herranz, N., S. Krasa, and A.P. Villamil. 2015. "Entrepreneurs, Risk Aversion, and Dynamic Firms." *Journal of Political Economy* 123 (5): 1133–1176.

Hollingworth, H. L. 1910. "The Central Tendency of Judgment." *The Journal of Philosophy, Psychology and Scientific Methods* 7 (17): 461–469.

Kamdar, R. 2019. "The Inattentive Consumer: Sentiment and Expectations." *manuscript.*

Keller, G., and S. Rady. 2010. "Strategic Experimentation with Poisson Bandits." *Theoretical Economics* 5 (2): 275–311.

Keller, G., S. Rady, and M. Cripps. 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica* 73 (1): 39–68.

Kerr, W.R., R. Nanda, and M. Rhodes-Kropf. 2014. "Entrepreneurship as Experimentation." *Journal of Economic Perspectives* 28 (3): 25–48.

Khaw, M.W., Z. Li, and M. Woodford. 2019. "Cognitive Imprecision and Small-Stakes Risk Aversion." *NBER Working Paper*, vol. no. 24978.

Kihlstrom, R., and J.J. Laffont. 1979. "A General Equilibrium Entrepreneurial Theory of Firm Formation Based on Risk Aversion." *Journal of Political Economy* 87 (4): 719–749.

Knight, F. 1921. *Risk, Uncertainty and Profit.* Boston: Hart, Schaffner & Marx.

Kydland, F. E., and E. C. Prescott. 1982. "Time to build and aggregate fluctuations." *Econometrica: Journal of the Econometric Society*, pp. 1345–1370.

Lilienfeld, S.O. *et al.* 2012. "Fearless Dominance and the U.S. Presidency: Implications of Psychopathic Personality Traits for Successful and Unsuccessful Political Leadership." *Journal of Personality and Social Psychology* 103 (3): 489–505.

Lorenzoni, G. 2009. "A theory of demand shocks." *American Economic Review* 99 (5): 2050–84.

Mackowiak, B., F. Matějka, and M. Wiederholt. 2018. "Survey: Rational inattention, a disciplined behavioral model." *CEPR Discussion Papers (No. 13243).*

Maćkowiak, B., F. Matějka, and M. Wiederholt. 2018. "Dynamic rational inattention: Analytical results." *Journal of Economic Theory* 176:650–692.

Mackowiak, B., and M. Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99 (3): 769–803.

———. 2015. "Business Cycle Dynamics Under Rational Inattention." *The Review of Economic Studies* 82 (4): 1502–1532.

Mankiw, G.N., and R. Reis. 2002. "Sticky Information versus Sticky Prices: A Pro- posal to Replace the New Keynesian Phillips Curve." *The Quarterly Journal of Economics* 117 (4): 1295–1328.

Manso, G. 2015. "Experimentation and the Returns to Entrepreneurship." *Mimeo, Haas School of Business, UC Berkeley.*

Masatlioglu, Y., A. Y. Orhun, and C. Raymond. 2017. "Intrinsic information preferences and skewness."

Matějka, Filip, and Alisdair McKay. 2015. "Rational inattention to discrete choices: A new foundation for the multinomial logit model." *American Economic Review* 105 (1): 272–298.

McCarty, N., K.T. Poole, and H. Rosenthal. 2008. *Polarized America: The dance of ideology and unequal riches.* MIT Press Books.

Melosi, L. 2014. "Estimating models with dispersed information." *American Economic Journal: Macroeconomics* 6 (1): 1–31.

Miao, J., J. Wu, and E. Young. 2019. "Multivariate Rational Inattention." *(No. WP2019-07). Boston University-Department of Economics.*

Moskowitz, T., and A. Vissing-Jorgensen. 2002. "The Returns to Entrepreneurial Investment: A Private Equity Premium Puzzle?" *American Economic Review* 92 (4): 745–778.

Mosteller, F., and P. Nogee. 1951. "An Experimental Measurement of Utility." *Journal of Political Economy* 59:371–404.

Nimark, K., and S. Sundaresan. 2019. "Inattention and Belief Polarization." *Journal of Economic Theory* 180:203–228. Meeting Papers, Society for Economic Dynamics.

Paciello, L., and M. Wiederholt. 2014. "Exogenous information, endogenous information, and optimal monetary policy." *Review of Economic Studies* 81 (1): 356–388.

Poole, K.T., and H. Rosenthal. 1984. "The polarization of American politics." *The Journal of Politics* 46 (4): 1061–1079.

Post, T., M.J. van den Assem, G. Baltussen, and R.H. Thaler. 2008. "Deal or No Deal? Decision Making under Risk in a Large-Payoff Game Show." *American Economic Review* 98 (1): 38–71.

Pratt, John W. 1964. "Risk Aversion in the Small and the Large." *Econometrica* 32 (1-2): 122–136.

Roberts, K., and M.L. Weitzman. 1981. "Funding Criteria for Research, Development, and Exploration Projects." *Econometrica* 49 (5): 1261–1288.

Rothschild, M. 1974a. "Searching for the Lowest Price When the Distribution of Prices Is Unknown." *Journal of Political Economy* 82 (4): 689–711.

———. 1974b. "A Two-armed Bandit Theory of Market Pricing." *Journal of Economic Theory* 9 (2): 185–202.

Savage, L. J. 1954. *The foundations of statistics.* Dover Books.

Schiller, B., and P. Crewson. 1997. "Entrepreneurial Origins: A Longitudinal Inquiry." *Economic Inquiry* 35 (3): 523–531.

Schorfheide, F. 2000. "Loss function-based evaluation of DSGE models." *Journal of Applied Econometrics* 15 (6): 645–670.

———. 2008. "DSGE model-based estimation of the New Keynesian Phillips curve." *FRB Richmond Economic Quarterly* 94 (4): 397–433.

Selten, R. 1998. "Axiomatic characterization of the quadratic scoring rule." *Experimental Economics* 1 (1): 43–61.

Sims, C. 1998a. Chapter Stickiness of *Carnegie-Rochester Conference Series on Public Policy*, Volume 49, 317–356. Elsevier.

Sims, C.A. 1998b. "Stickeness." *Carnegie-Rochester Conference Series on Public Policy* 49:317–356.

———. 2003a. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665–690.

Sims, C.A., and T. Zha. 1998. "Bayesian methods for dynamic multivariate models." *International Economic Review*, no. 949-968.

Sims, Christopher. 2003b. "Implications of rational inattention." *Journal of Monetary Economics* 50 (3): 665–690.

Smith, M. 2019. Most voters have only become more sure about their EU referendum vote.

Steiner, J., C. Stewart, and F. Matějka. 2017. "Rational Inattention Dynamics: Inertia and Delay in Decision-Making." *Econometrica* 85 (2): 521–553.

Sutton, R.S., and A.G. Barto. 1998. *Reinforcement Learning: An Introduction.* Cambridge, MA: MIT Press.

Weitzman, M.L. 1979. "Optimal Search for the Best Alternative." *Econometrica* 47 (3): 641–654.

Woodford, M. 2003. *Interest and Prices.* Princeton:Princeton University Press.

Woodford, M. 2002. Chapter Imperfect Common Knowledge and the Effects of Monetary Policy of *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, edited by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, 25–58. Princeton University Press, Princeton.