UTS Business School –
Department of Economics
Dr Chau Chak Wing Building
14 – 28 Ultimo Road, Ultimo NSW 2007

T: +61 2 9514 7713
economics@uts.edu.au

Department of Economics
Working Paper Series

‘A Theory of Intuition and Contemplation’

Benjamin Balzer [1]
Benjamin Young [2]

[1 & 2] University of Technology Sydney

No. 01 / 2020 **May 2020 (updated)**

A Theory of Intuition and Contemplation*

Benjamin Balzer† & Benjamin Young‡

May 4, 2020

Abstract

We introduce a model of intuition and contemplation in decision problems under uncertainty. Intuition is a garbling of true information and contemplation is the ability to recover the true informational content of signals. We define natural orders on the quality of intuition and on contemplative ability. In *any* non-strategic decision problem, the agent's utility increases as either the quality of her intuition or her contemplative ability improves. We derive versions of Blackwell's Informativeness Theorem for our intuitive agent and apply the model to the canonical Bayesian persuasion problem.

*We thank Alex Jakobsen, Kentaro Tomoeda, Toru Suzuki, Antonio Rosato, Jun Zhang, Emil Temnyalov, Isa Hafalir, and seminar participants at the University of Calgary, UTS, and UNSW for insightful feedback.

†University of Technology Sydney (benjamin.balzer@uts.edu.au)
‡University of Technology Sydney (benjamin.young@uts.edu.au)

1 Introduction

The problem of an individual trying to draw inference from observations is the backbone of many economic problems. Researchers often make the simplifying assumption that individuals are able to extract the true Bayesian content of information for free. However, extracting information and drawing inference in the real world is complicated and individuals may not be Bayesian in every context. Instead, they may first rely on their intuition to proxy the meaning of revealed information and then further contemplate in order to discover its true content. In this paper, we provide a formal model of such a cognitive process which is portable to any setting where an individual faces an information extraction problem.

Consider the following motivating example. An investor needs to manage her portfolio. There are three states of the economy: *Good*, *Fair*, and *Bad*. Each state determines a set of optimal trades. The investor, however, does not know the state and, instead, relies on news to inform her decision making. Suppose she comes across a newspaper article which she knows will reveal the true state if carefully read. However, fully reading and interpreting the article is costly. Instead, the investor can quickly skim the article to intuit whether the state of the economy is either *Good* or in the set {*Fair*, *Bad*}. Suppose the true state is *Bad*. Then, the first impression of the investor from her quick skim of the article is to rule out state *Good*. She must now decide whether to adjust her portfolio on the basis of this first impression or to fully read the article to resolve the remaining uncertainty. The following diagram conceptualizes this situation.

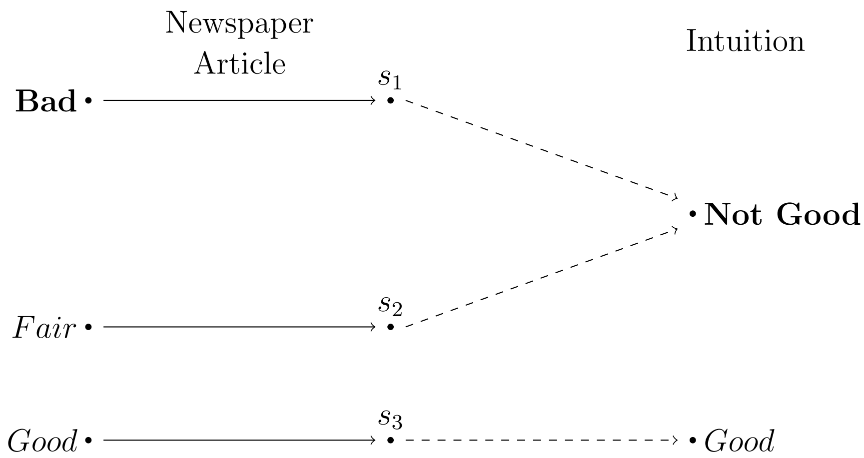

It is natural to model this newspaper article as an *experiment* that is fully-revealing. However, instead of observing its outcome (i.e s_1 since the state is *Bad*), the investor observes for free the noisier signal *Not Good* (i.e. either s_1 or s_2 has realized). We think of *intuition* as the Bayesian update associated with this first impression. With intuition *Not Good*, however, the investor still faces uncertainty as to whether the true signal is s_1 or s_2 (i.e. the state is *Bad* or *Fair*). We model *contemplation* as a costly process that allows the investor to resolve this residual uncertainty.

The details of the model are as follows. An agent faces an experiment which is a stochastic map from payoff-relevant states of the world into a signal space. The agent, however, does not initially process the signal generated by the experiment. Instead, she processes information on a *mental signal space*, which is reached via some garbling of the experiment. Her intuition is the Bayesian update of the generated mental signal given the experiment and the garbling. We allow the *intuition-generating process* to be any garbling. As such, it is always noisier than true information and permits stochastic intuition. We model contemplation as the agent's ability to invert her intuition back to the true signal. Specifically, for each mental signal, the agent chooses an intensity of contemplation which determines the probability of successfully identifying the meaning of the true signal realization. Thus, contemplation serves as a bridge between intuitive decision-making and its Bayesian counterpart. The agent's contemplation decision is governed by a trade-off between the perceived benefits from contemplation (i.e. being able to use all information that has been revealed) and exogenous cognitive costs.

We define natural orders on the agent's quality of intuition and contemplative ability. Our order on the quality of intuition implies that the agent's intuition improves as it becomes less noisy. We show that the agent ascribes a higher value to information in *any* decision problem as her intuition improves. We say that the agent's contemplative ability improves if her cognitive costs uniformly decrease. In any decision problem, the agent is better off as her contemplative ability improves. Moreover, we show that, in the prominent case of convex cognitive costs, this order is also necessary to ensure the agent is better off in any decision problem.

We characterize the set of experiments that are unambiguously welfare-ranked for our intuitive agent. First, if nothing is known about the agent's intuition-generating process, we can only conclude that the agent prefers some information to no information. Second, we introduce a condition on the agent's intuition-generating process, *intuitive sufficiency*, under which Blackwell's equivalence between garblings and the value of information (Blackwell (1953)) is restored. Intuitive sufficiency essentially guarantees that, when experiments become less noisy in the Blackwell sense, the agent's intuition preserves this reduction in noise.

Our model fits with recent empirical evidence. In the context of noisy information, Ambuehl and Li (2018) find that individuals (i) differ consistently in how their beliefs respond to information, (ii) underreact in their value of information to increased informativeness of the experiment, and (iii) value experiments that reveal states of the world with certainty disproportionately highly. We can capture these findings within our model by assuming that (i) different individuals have different intuition-generating processes, (ii) the intuition-generating process does not preserve the increased informativeness of experiments in one-to-one proportion, and (iii) individual's have perfect intuition at signals that induce certainty, but relatively poor intuition otherwise. In another experimental study, Enke et al. (2020) find that, for tasks that do not require high-level problem solving skills, increasing stakes can improve an individual's task performance through reduced reliance on intuition. Our model captures this finding as the agent will contemplate more (i.e. depend on her intuition less) as the stakes increase.

We embed our intuitive agent in a strategic setting; the canonical Bayesian persuasion problem introduced by Kamenica and Gentzkow (2011). In this problem, a sender, whose preferences are misaligned with those of a receiver, designs an experiment to persuade the receiver to take sender-preferred actions. By focusing on a particular intuition-generating process and cognitive cost function, we characterize the sender's optimal experiment. We show that the informativeness of this experiment may increase as either the receiver's quality of intuition or contemplative ability worsens. This is because poor intuition and low contemplative ability both serve as a source of *commitment power* for the receiver to take undesirable actions from the sender's perspective when using her intuition. This incentivizes the sender to increase the amount of information he reveals. We show, in turn, when this increased informativeness increases the receiver's equilibrium utility. Consequently, becoming more cognitively limited, which hurts the agent in non-strategic settings, can make her better off when information is designed.

The paper proceeds as follows. In Section 2, we discuss the related literature. Section 3 introduces and discusses the model, and provides a simple example for illustrative purposes. We present the general implications of the model in Section 4. In Section 5 we apply the model to the canonical Bayesian persuasion problem. Section 6 concludes.

2 Related Literature

Our model of intuition and contemplation is related to the literature on costly information processing, most closely to theories of rational inattention (Sims (2003), Caplin and Dean (2015), De Oliveira et al. (2017)).[1] We provide a detailed comparison of our model and rational inattention in Section 3.2.3, where we explain that neither model nests the other. Specifically, while our model of contemplation is nested in the rational inattention framework, our model of intuition is a novel concept. This allows us to study the interesting interaction between intuition and contemplation which is absent in rational inattention models. Outside the paradigm of rational inattention, Kominers, Mu, and Peysakhovich (2018) provide a model in which an individual can

[1]Those cited constitute a selection of conceptual contributions to the field of rational inattention. Rational inattention has been applied extensively across of myriad of economic fields, including macroeconomics, finance, and behavioral economics. See Maćkowiak, Matejka, and Wiederholt (2018) for a detailed survey.

make decisions using her prior for free but must pay a fixed cost to access the true Bayesian content of information. Instead, the information that our agent receives for free is more general than the prior (i.e. her intuition) and we allow for more general contemplation processes.

Our paper is related to the model of 'local thinking' provided in Gennaioli and Shleifer (2010). In their model, an agent evaluates new information using past experience to form representative comparisons. Such an agent can exhibit heuristical behavior that is difficult to explain within the Bayesian paradigm. In our model of intuition, past experience may indirectly affect the agent's evaluation of information by shaping her intuition-generating process. However, given her imperfect intuition, she processes information 'as-if' she were Bayesian. This allows us to provide stark predictions for how the value of information changes as the model's primitives vary; a task which is difficult to achieve within the local-thinking framework. Moreover, we also consider the interaction between intuition and contemplation, which allows the agent to correctly evaluate information.

We also relate to models of coarse belief updating (Mullainathan (2002), Jakobsen (2020)). In these papers, the agent's interpretation of a signal may also not coincide with its Bayesian update. However, these models assume that the Bayesian update of a given signal still fully determines the agent's resulting belief update at this signal. In contrast, our agent's intuition-generating process ensures that the intuition associated with a realized signal depends on both the Bayesian update of that signal *and* counterfactual signals that did not realize. In particular, our agent may hold different intuitive beliefs for signals with the same informational content. Moreover, we also explore the role of contemplation for our intuitive agent.

In our model of information extraction, the agent makes contemplation decisions by trading off a perceived value of contemplation against exogenous cognitive costs. Hence, our paper relates to contributions in the literature that explore a similar trade-off, albeit in completely different contexts. In the level-k framework, Alaoui and Penta (2016) model an individual who can endogenously transition between different depths of reasoning as a function of the perceived benefits and the exogenous costs of doing so. Ergin and Sarver (2010) provide a decision-theoretic framework where costly contemplation allows for state-contingent planning by revealing information about the individual's menu-dependent preferences.

The main application of our theory is Bayesian persuasion, as introduced by Kamenica and Gentzkow (2011). Other papers have also explored how deviations from the paradigm of Bayesianism impact optimal persuasion. de Clippel and Zhang (2019) allow for the receiver's updating rule to not coincide with Bayes' rule. Beauchêne, Li, and Li (2019) consider a model with a receiver who is ambiguity averse with maxmin expected utility, where a sender can utilize ambiguous experiments. These papers are mainly interested in characterizing the sender's value of persuasion. In contrast, our paper is focused on the welfare of the intuitive receiver who is being persuaded. Matyskova and Montes (2018) are also interested in receiver welfare in a model where the receiver has access to costly information acquisition. However, their receiver only has the ability to learn *more* information than what is provided by the experiment. In contrast, our receiver can, *at most*, learn the information provided by the experiment. As a complement to our finding that cognitive limitations can increase the receiver's welfare, they show that it can be harmful to the receiver to have access to additional information.

Models of cognitively limited agents have been applied elsewhere in the literature. Ravid (2019) provides a model of bargaining with a rationally inattentive buyer. He shows that being rationally inattentive can benefit the buyer as it leads the seller to increase the quality of his products. Similarly, we show that our intuitive receiver can be better off in a persuasion context as bounded intuition induces higher quality of information. Martin (2017) considers a rationally inattentive buyer in a strategic price-setting game. He shows that the equilibrium informativeness of prices may actually increase in the buyer's cost of attention. We establish a similar result: increasing contemplation costs can lead to more informative experiments.

3 The Model

3.1 Formal Details

3.1.1 Preliminaries

Stochastic Maps: For a finite set X, let $\Delta(X)$ denote the set of probability distributions over X. A *stochastic map* between two finite sets X and Y is a function $\alpha : X \to \Delta(Y)$. We write $\alpha(y|x)$ to denote the probability of y given x under stochastic map α. For two stochastic maps $\alpha : X \to \Delta(Y)$ and $\beta : Y \to \Delta(Z)$, define the *composite* of α and β, denoted $\beta \circ \alpha$, as

$$\beta \circ \alpha(z|x) = \sum_{y \in Y} \beta(z|y)\alpha(y|x)$$

for all $z \in Z$ and $x \in X$. Take two stochastic maps $\alpha : X \to \Delta(Y)$ and $\beta : X \to \Delta(Z)$. Then, β is a *garbling* of α if there exists a stochastic map, $g : Y \to \Delta(Z)$ such that $\beta = g \circ \alpha$.

Information Primitives: There is a finite state space, $\Omega = \{\omega_1, \ldots, \omega_n\}$, $n \geq 2$, with associated full-support prior $\mu \in \Delta(\Omega)$. An *experiment* is the combination of a finite signal space, $S = \{s_1, \ldots, s_k\}$, and a stochastic map $\sigma : \Omega \to \Delta(S)$, where $\sigma(s|\omega)$ denotes the conditional probability of signal $s \in S$ given $\omega \in \Omega$. In a slight abuse of notation we write σ to denote experiment (S, σ). Let $\mathcal{E}$ denote the set of all experiments.

Given experiment $\sigma \in \mathcal{E}$ and signal realization $s \in S$, define the Bayesian operator

$$B(s, \sigma) \equiv (B(\omega_1|s, \sigma), \ldots, B(\omega_n|s, \sigma)),$$

where

$$B(\omega|s, \sigma) \equiv \frac{\sigma(s|\omega)\mu(\omega)}{\sum\limits_{\omega' \in \Omega} \sigma(s|\omega')\mu(\omega')}.$$

That is, $B(s, \sigma) \in \Delta(\Omega)$ is the Bayesian update given signal realization $s \in S$ in experiment σ.

Agent Preferences: An agent evaluates the informational content of experiments in order to make decisions. Formally, the agent chooses action $a \in A$ where A is a compact action space. The agent has a state-dependent utility function, $u : A \times \Omega \to \mathbb{R}$, where $u(a, \omega)$ denotes the utility of taking action $a \in A$ in state $\omega \in \Omega$. We assume $u(\cdot, \omega)$ is continuous for all $\omega \in \Omega$. Let $p \in \Delta(\Omega)$ denote an arbitrary belief over the state space. The agent chooses an action $a \in A$ that maximizes her expected utility which is given by

$$U(a|p) \equiv \sum_{\omega \in \Omega} u(a, \omega) p(\omega). \tag{1}$$

Let $a^*(p)$ denote the set of optimal actions at belief p.[2]

3.1.2 Intuition and Contemplation

Fix an experiment $\sigma : \Omega \to \Delta(S)$ and suppose signal $s \in S$ realizes. A Bayesian agent observes s, computes Bayesian update $B(s, \sigma)$, and chooses an action from the set $a^*(B(s, \sigma))$. However, an intuitive agent may not process the meaning of the true signal realization from experiment σ for free. Figure 1 provides a general overview of how we model an intuitive and contemplative agent.

We now discuss each component of this cognitive process in detail.

Intuition: Fix an experiment σ defined on signal space S and suppose signal $s \in S$ realizes. We want to model a situation in which the agent does not necessarily understand the true informational content of s, $B(s, \sigma)$, but does have an *intuition* for its meaning. We capture this idea through the modeling assumption that the agent does not observe the true signal realization $s \in S$ but rather observes a noisier signal, $\tilde{s}$, in a mental signal space, $\tilde{S}$. She then uses the informational content of $\tilde{s}$ to approximate the true information. We call this approximation her intuition.

[2]Since $u(\cdot, \omega)$ is continuous for all $\omega \in \Omega$, $a^*(p) \neq \emptyset$. Generically, $a^*(p)$ will be a singleton. However, it may contain multiple elements in situations in which p makes the agent indifferent between multiple actions. Depending on the application, one may need to define a selection rule to break indifferences.

The agent observes the experiment σ and signal $s \in S$ realizes

$\Downarrow$

The agent does not directly process s but rather observes signal $\tilde{s}$ in mental signal space $\tilde{S}$

$\Downarrow$

The agent forms an intuition using the informational content of $\tilde{s}$ as a proxy for true information

$\Downarrow$

The agent chooses intensity of contemplation, ρ, which determines the likelihood of discovering the true signal realization s

$\Downarrow$

If contemplation is successful, the agent takes a Bayesian optimal action. Else, she bases her decision on her intuition.

Figure 1: The Cognitive Process

Formally, for any experiment σ there exists a stochastic mapping $\Gamma_\sigma : S \to \Delta(\tilde{S})$ which (probabilistically) associates true signals in S with signals in some mental signal space $\tilde{S}$. This stochastic mapping depends on experiment σ for two reasons: (i) σ determines the true signal space, S; and (ii) we allow for situations in which the informational content of σ can affect the agent's intuition.[3] Note that $\Gamma_\sigma \circ \sigma$ is itself an experiment on $\tilde{S}$; that is, $\Gamma_\sigma \circ \sigma \in \mathcal{E}$. We refer to $\Gamma_\sigma \circ \sigma$ as the agent's *intuitive experiment*. This intuitive experiment is simply a garbling of the true experiment through Γ_σ. We term Γ_σ the agent's *intuition-generating process at* σ and let $\Gamma \equiv \{\Gamma_\sigma\}_{\sigma \in \mathcal{E}}$ denote the agent's *intuition-generating process*.

Given σ and true signal realization s, process Γ_σ generates a mental signal,

[3] For example, the way σ is *framed* may impact how the agent processes this information (Tversky and Kahneman (1981), Entman (1993)) or the aspects of this information that are salient (Price, Tewksbury, and Powers (1997)).

$\tilde{s}$. We model the agent's *intuition* at true signal s to be the informational content of the realized mental signal $\tilde{s}$ given intuitive experiment, $\Gamma_\sigma \circ \sigma$; that is, her intuition is $B(\tilde{s}, \Gamma_\sigma \circ \sigma)$. Note that since Γ_σ is a stochastic map, it holds that (i) the agent's intuition at true signal s may be generated stochastically; and (ii) her intuitive experiment is, overall, noisier than the true experiment σ.

If the agent relies on her intuition, she chooses an action in $a^*(B(\tilde{s}, \Gamma_\sigma \circ \sigma))$. We depict the two extremes cases of intuition-generating processes in Figure 2: *perfect intuition* (i.e. the agent intuits the true informational content) and *prior-based intuition* (the agent intuits no informational content).

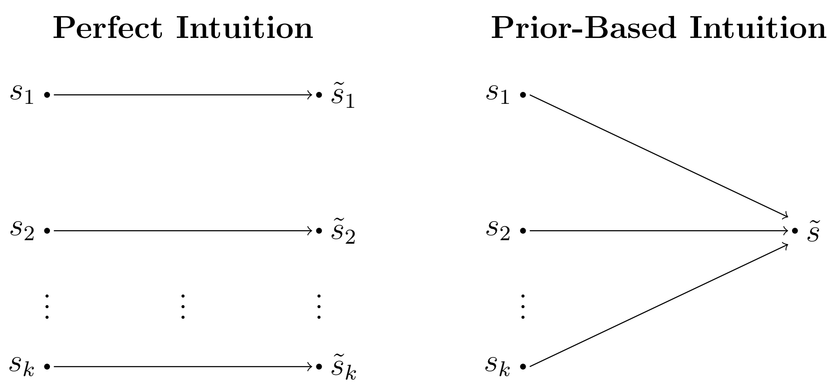

Figure 2: The Extremes of Intuition.

Contemplation: For given mental signal $\tilde{s}$, the agent perceives uncertainty as to which $s \in S$ generated it. We model contemplation as the ability to resolve this uncertainty. We assume the agent has access to a continuous contemplation technology. Formally, the agent chooses $\rho \in [0, 1]$, where ρ is the probability that she successfully inverts mental signal $\tilde{s}$ to true signal s. Instead, with $1 - \rho$ probability, this inversion fails and the agent uses her intuition to make a decision. Contemplation is costly: success probability ρ comes at cognitive cost (in utils) $C(\rho) \geq 0$. We assume that $C(0) = 0$, $C(\cdot)$ is weakly increasing, and left-continuous.[4]

Given σ and Γ_σ, we can quantify the agent's perceived uncertainty over true signals at mental signal $\tilde{s}$. In particular, the conditional probability of s given mental signal $\tilde{s}$ is given by

$$Pr(s|\tilde{s}, \sigma, \Gamma_\sigma) = \frac{\Gamma_\sigma(\tilde{s}|s) \sum_{\omega \in \Omega} \sigma(s|\omega)\mu(\omega)}{\sum_{s' \in S} \Gamma_\sigma(\tilde{s}|s') \sum_{\omega \in \Omega} \sigma(s'|\omega)\mu(\omega)}.$$

We now describe the agent's contemplation problem. Let $a_s \in a^*(B(s, \sigma))$ denote a Bayesian optimal action at true signal $s \in S$. Let $a_{\tilde{s}} \in a^*(B(\tilde{s}, \Gamma_\sigma \circ \sigma))$ denote an intuitive optimal action at mental signal $\tilde{s}$. Then, the agent chooses contemplation intensity ρ to solve

$$\max_{\rho \in [0,1]} \quad \rho \sum_{s \in S} Pr(s|\tilde{s}, \sigma, \Gamma_\sigma) U(a_s|B(s, \sigma)) + (1-\rho) U(a_{\tilde{s}}|B(\tilde{s}, \Gamma_\sigma \circ \sigma)) - C(\rho). \quad (2)$$

According to (2), the agent trades off the *perceived benefits from contemplation* against the *cost of contemplation*.[5] Contemplation is valuable because it allows the agent to attune her action to the true information that has been revealed, which may be different from her intuition.

Notice that the benefit from contemplation is not correct at the interim stage. In particular, while some true signal s has realized (which generated some mental signal $\tilde{s}$), the agent perceives that there is uncertainty as to which true

[4] These two assumptions together imply that $C(\cdot)$ is lower semi-continuous, which ensures there is a solution to the agent's optimal contemplation problem.

[5] See Alaoui and Penta (2018) for an axiomatization of the cost-benefit approach to cognitive problems.

signal has realized; that is, she may attach positive probability to signals that have not realized. If contemplation is successful, the agent *will observe the true signal realization with probability one*. However, her perception of what she *might* observe is correct from the ex-ante perspective. This is precisely because we assume the agent *fully understands* how her intuition is generated.

3.2 Discussion of Model

3.2.1 Intuition

We start by emphasizing two important ingredients to our model of intuition. First, the agent *understands the true experiment* but does not necessarily process its true informational content. Instead, she holds an intuition which is a garbling of this true information and, as such, is stochastically noisier. Second, the agent *understands her intuition-generating process*. These two assumptions together imply that the agent's perception of the benefits from contemplation are correct from an ex-ante perspective, which is important for tractability and the generality of our results presented in Section 4. This means we do not capture situations in which the agent's intuition is systematically biased on average, as her intuition always averages out to the prior. As such, our model of intuition is only a minor deviation from the paradigm of Bayesianism.

3.2.2 Contemplation

The contemplation processes we consider are a restricted class of experiments that reveal information about true signal realizations: our contemplation strategies characterized by $\rho \in [0, 1]$. Similar to the interpretation in Tirole (2009), one can think of this contemplation process as a form of sampling. Using the example from the introduction, the choice of ρ could be interpreted as a random sampling of paragraphs from the article that the investor carefully reads. If the true information she requires is in a sampled paragraph, she learns it; otherwise, she learns nothing. Cognitive costs associated with contemplation can arise from a number of sources. For example, one could equally think of these as opportunity costs or a physical cost in the form of disutility from mental strain.

For these sampling-based contemplation processes, the agent will choose actions using either her intuition or the true information (i.e. as a Bayesian). Thus, these processes are the most straightforward bridge between intuitive decision-making and its Bayesian counterpart. We impose this restriction to starkly highlight the difference between intuitive decision-making and its Bayesian counterpart. Moreover, this simplifying assumption greatly improves tractability in applications while still providing insightful conclusions.

3.2.3 Relation to Rational Inattention

Our model is most closely related to static models of rational inattention. There are, however, two important differences. First, in rational inattention models, the agent chooses her attention strategy using the prior (i.e. acquires information about the true experiment at the ex-ante stage). In contrast, our agent receives first some information for free (i.e. her intuition) and then, conditional on this first impression, makes a contemplation decision (i.e. acquires information about the true experiment at an interim stage). This implies that our intuitive agent can correlate her contemplation strategy more effectively with the true signal than a rationally inattentive agent can correlate her attention strategy. In this sense, our model of intuition is a generalization of static rational inattention.

Second, rational inattention models generally allow the agent to choose any experiment that reveals information about true signals as an attention strategy. In contrast, our intuitive agent is restricted to a specific class of experiments as described in the previous section. In this sense, rational inattention models allow for more general attention strategies than our contemplation technology.

These two points together imply that, while clearly related, neither model nests the other. However, if we were to allow for more general contemplation strategies (i.e. any experiment that reveals information about the true signal realization), then our model of intuition and contemplation would *strictly generalize* the static rational inattention framework. Specifically, rational inattention would be the special case in which our intuitive agent has prior-based intuition as in Figure 2.

3.3 An Illustrative Example

In what follows, we define a particular intuition-generating process and specification of cognitive costs and apply them to a particular decision problem in order to illustrate the mechanics of the model.

Intuition and Contemplation: Take any experiment $\sigma : \Omega \to \Delta(S)$. We describe an intuition-generating process that is a mixture of perfect intuition and prior-based intuition (see Figure 2). In particular, the mental signal space is given by $\tilde{S} \equiv S \cup \{\tilde{s}_{\emptyset}\}$. Define the intuition-generating process at σ to be $\Gamma^{\phi}_{\sigma}(s|s) = \phi \in [0,1]$ and $\Gamma^{\phi}_{\sigma}(\tilde{s}_{\emptyset}|s) = 1 - \phi$ for all $s \in S$. Figure 3 illustrates this process.

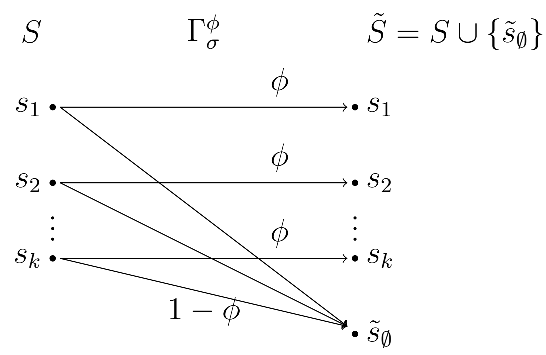

Figure 3: A ϕ-mixture of perfect intuition and prior-based intuition.

Intuition-generating processes within this class are parameterized by a single variable, $\phi \in [0, 1]$. Given true signal s, the agent either learns its content for free (with probability ϕ) or entertains mental signal $\tilde{s}_\emptyset$ (with probability $1-\phi$). At mental signal $\tilde{s}_\emptyset$, the agent's intuition is equal to the prior, μ, *regardless of the experiment.* In this sense, this class of intuition-generating processes is a ϕ-mixture of perfect intuition and prior-based intuition.

We assume that cognitive costs are given by $C(\rho; \kappa) = \kappa\rho^2/2$, where the parameter $\kappa \geq 0$ determines the marginal cost of contemplation.

The Decision Problem: There are two states, $\Omega = \{\omega_1, \omega_2\}$, with a uniform prior. The action space coincides with the state space, $A = \Omega$, and the agent wants to choose an action to match the state. Formally, utility is given by

$$u(a, \omega) = \begin{cases} 1 & \text{if } a = \omega \\ 0 & \text{if } a \neq \omega. \end{cases} \tag{3}$$

The experiment σ is fully informative; that is, there exists a binary signal space $S = \{s_1, s_2\}$ such that $\sigma(s_i|\omega_i) = 1$ for $i = 1, 2$. Clearly, $B(s_1, \sigma) = (1, 0)$ and $B(s_2, \sigma) = (0, 1)$ so that if the agent were Bayesian she would take the action ω_i when signal s_i realizes, $i = 1, 2$.

Our agent, however, does not always process information on S. Instead, her intuitive experiment is given by $\Gamma_\sigma^\phi \circ \sigma(s_i|\omega_i) = \phi$ and $\Gamma_\sigma^\phi \circ \sigma(\tilde{s}_\emptyset|\omega_i) = 1 - \phi$ for $i = 1, 2$. Hence, the agent's intuitive beliefs at each mental signal are

$$B(s_1, \Gamma_\sigma^\phi \circ \sigma) = (1, 0); \quad B(s_2, \Gamma_\sigma^\phi \circ \sigma) = (0, 1); \quad B(\tilde{s}_\emptyset, \Gamma_\sigma^\phi \circ \sigma) = \mu = (1/2, 1/2).$$

In particular, for a given true signal s, the agent's intuition is $B(s, \sigma)$ with probability ϕ and is $(1/2, 1/2)$ with probability $1 - \phi$. Therefore, she takes a Bayesian optimal action at true signal s with probability ϕ. However, at $\tilde{s}_\emptyset$ she is unable to correctly correlate her action with the true signal realization. As such, her action choice at $\tilde{s}_\emptyset$ may not be Bayesian optimal.

Given the potential for flawed decision-making at mental signal $\tilde{s}_\emptyset$, the agent perceives that there is a benefit from contemplating. In particular, if contemplation is successful, she receives a payoff of 1. In contrast, if she does not contemplate, she receives a payoff of 1/2. Her optimal contemplation decision at $\tilde{s}_\emptyset$ solves

$$\max_{\rho \in [0,1]} \rho + (1 - \rho)(1/2) - \kappa \frac{\rho^2}{2}$$

which has solution $\rho^*(\tilde{s}_\emptyset) = \min\left\{\frac{1}{2\kappa}, 1\right\}$.

As ϕ increases, there is a natural sense in which the agent's intuition improves:

it becomes more likely that the agent's intuition coincides with the meaning of the actual signal realization. Similarly, as κ decreases there is a natural sense in which the agent's contemplative ability improves: the cost of contemplation at any given level, ρ, decreases. Moreover, the agent's ex-ante utility from experiment σ is increasing as ϕ increases or κ decreases. Specifically, ex-ante indirect utility, $V(\phi, \kappa, \sigma)$, can be computed to be

$$V(\phi, \kappa, \sigma) = \begin{cases} 1 - \frac{\kappa}{2} & \text{if } \kappa \leq 1/2 \\ \phi + (1-\phi)\left[\frac{1}{2} + \frac{1}{8\kappa}\right] & \text{if } \kappa \geq 1/2, \end{cases}$$

which is decreasing in κ and (weakly) increasing in ϕ. Hence, the agent values σ more as either her intuition or contemplative ability improves. In the next section we define general notions of better intuition and greater contemplative ability that ensure these predictions carry over to *any* decision problem the agent may face.

4 General Implications

In this section, we present comparative static predictions of our model in arbitrary decision problems. Our main focus is on how the *value of information* depends on the primitives of our model. The value of information for some experiment σ is the difference between the payoff the agent receives when choosing optimally under σ given (Γ, C), and the payoff she receives when using her prior.

We explore, in turn, how the intuition-generating process (Γ), the cognitive cost function (C), and the experiment (σ) affect the value of information. First, we provide an order on the quality of intuition which ensures that an agent always values information more as her intuition improves. Second, we provide an order on cognitive cost functions such that an agent with greater contemplative ability always values information more. Finally, we explore conditions on the statistical properties of both experiments and intuition-generating processes that allow for value of information comparisons across experiments.

4.1 Intuition

We first investigate the role that intuition plays in our model. We define the following order on the quality of the agent's intuition.

Definition 1 (Better Intuition)**.** *Take two intuition-generating processes,* Γ *and* Γ'*. We say that* Γ *displays better intuition at* σ *than* Γ' *if* Γ'_σ *is a garbling of* Γ_σ*. We say that* Γ *displays better intuition than* Γ' *if the former displays better intuition at any* $\sigma \in \mathcal{E}$*.*

This definition is natural in that the agent's intuition improves as it becomes less noisy in the Blackwell sense. Recall the intuition-generating process of the detailed example described in Section 3.3, Γ^ϕ (see Figure 3). For this intuition-generating process, there is a sense in which the agent's intuition improves as ϕ increases for any experiment $\sigma \in \mathcal{E}$. Our definition precisely captures this idea: for $\phi, \phi' \in [0, 1]$ with $\phi > \phi'$ it is straightforward to show that $\Gamma^{\phi'}_\sigma$ is a garbling of Γ^ϕ_σ. Moreover, an agent facing the decision problem in Section 3.3 becomes better off as ϕ increases. It turns out that if the agent's intuition improves in the sense of Definition 1 then she becomes better off in *any* decision problem.

Proposition 1. *Fix $\sigma \in \mathcal{E}$ and take two intuition-generating processes, Γ and Γ', where Γ displays better intuition at σ than Γ'. Then the agent values σ more under Γ than under Γ' for all (A, u, C, μ).*

Proposition 1 implies that better intuition leads to (weakly) higher ex-ante utility, *independent* of the decision problem the agent faces. Therefore, the better her intuition the more the agent values σ. The reason for this result is that the agent can more effectively attune her action to the realized state of the world as her intuition improves. Indeed, since Γ'_σ is a garbling of Γ_σ, there exists g such that $\Gamma'_\sigma = g \circ \Gamma_\sigma$. Under Γ_σ, the agent can use g in order to replicate, for each state ω, the distribution of (a, ρ) that she chooses under Γ'_σ. Since this conditional distribution is sufficient for the agent's ex-ante utility, she is better off under Γ_σ when she optimizes. This argument is similar to the reason why a Bayesian agent prefers experiments that are less noisy in the Blackwell sense (see de Oliveira (2018)).

Proposition 1 implies that improving one's intuition can never be harmful. However, this does not imply that improved intuition increases the likelihood of taking a Bayesian optimal action. Example 1 illustrates this point.

Example 1. *Consider a decision problem where $\Omega = \{\omega_1, \omega_2\}$, $\mu(\omega_1) = 1/2$, $A = \Omega$, utility is given by (3), and $\sigma(s_i|\omega_i) = 1$ for $i = 1, 2$. The intuition-generating process has the structure $\Gamma^\phi_\sigma(\tilde{s}_1|s_1) = \phi \in [0, 1]$ and $\Gamma^\phi_\sigma(\tilde{s}_2|s_2) = 1$. We consider two processes: one with $\phi = 1/2$ and one with $\phi = 0$. Clearly, $\Gamma^{1/2}_\sigma$ displays better intuition than Γ^0_σ.*

Suppose the agent's cognitive cost function is such that $C(0) = 0$ and $C(\rho) = \kappa$ for all $\rho \in (0, 1]$. It follows that, at each mental signal realization, the agent ei-

ther does not contemplate ($\rho = 0$) or fully contemplates ($\rho = 1$). This decision is determined by comparing the perceived benefits from contemplation to κ. For Γ_σ^0, only mental signal $\tilde{s}_2$ realizes. In this case, the agent is indifferent between choosing ω_1 and ω_2 and, for any selection, her benefit from contemplation is equal to $(1/2)(1-0)+(1/2)(1-1) = 1/2$. Hence, she fully contemplates if $\kappa < 1/2$. In contrast, under $\Gamma_\sigma^{1/2}$, the agent takes a Bayesian optimal action when $\tilde{s}_1$ realizes. If, however, $\tilde{s}_2$ realizes her benefits from contemplating are given by $(1/3)[1-0]+(2/3)(1-1) = 1/3$. Hence, if $\kappa \in (1/3, 1/2)$ an agent with intuition Γ_σ^0 always contemplates fully leading to the Bayesian optimal action with probability one. Instead, an agent with intuition $\Gamma_\sigma^{1/2}$ never contemplates at mental signal $\tilde{s}_2$. Thus, she does not take a Bayesian optimal action when s_1 induces intuition $\tilde{s}_2$.

In the example above, the less intuitive agent is more likely to take a Bayesian optimal action because she contemplates more intensely than her more intuitive counterpart. Indeed, since a less intuitive agent has access to less information for free (and understands this), one may think she always perceives greater benefits from contemplation than a more intuitive agent. However, this is not the case. Indeed, take Example 1 but instead suppose that the prior has $\mu(\omega_1) = 2/3$. Then, one can easily show that the more intuitive agent always contemplates more than her less intuitive counterpart. Thus, the propensity to take Bayesian optimal actions is ambiguously related to the quality of intuition. Indeed, it is the *amount of uncertainty* the agent perceives after her intuition is generated that drives this propensity, rather than the quality of intuition itself.

4.2 Contemplation

We now investigate the role that contemplation plays in our model. The following order, defined by comparing cognitive cost functions, measures the agent's contemplative ability.

Definition 2 (Greater Contemplative Ability)**.** *Let C and C' denote two cognitive cost functions. Then C displays greater contemplative ability than C' if $C(\rho) \leq C'(\rho)$ for all $\rho \in [0, 1]$.*

Under Definition 2, the agent's contemplative ability improves as the cost of implementing any given contemplation level $\rho \in [0, 1]$ decreases. Recall the cognitive cost function, $C(\rho; \kappa) = \kappa\rho^2/2$, used in the illustrative example described in Section 3.3. There is a natural sense in which the agent's contemplation process improves as κ decreases. Indeed, for $\kappa < \kappa'$, $C(\rho; \kappa)$ dislplays greater contemplative ability than $C(\rho; \kappa')$ in the sense of Definition 2. The following proposition shows that the agent is better off in any decision problem as her contemplative ability improves.

Proposition 2. *Take two cognitive cost functions, C and C', where C displays greater contemplative ability than C'. Then, for any $(A, u, \sigma, \Gamma, \mu)$, the agent is better off with C than with C'.*

Proposition 2 implies that greater contemplative ability leads to (weakly) higher ex-ante welfare. Obviously, an agent with greater contemplative ability can always replicate the strategy of her less able counterpart at lower cognitive cost. Thus, the former's optimal strategy must yield her higher utility.

If we restrict attention to the prominent class of convex cognitive cost functions, the order from Definition 2 is also necessary for the agent to be better off in any decision problem.

Proposition 3. *Take two convex cognitive cost functions, C and C'. The following two statements are equivalent:*

(a) C displays greater contemplative ability than C'.

(b) For all $(A, u, \sigma, \Gamma, \mu)$, the agent is better off under C than under C'.

The proof of Proposition 3 leverages that, with convex cognitive costs, any level of contemplation $\rho \in [0, 1]$ can be an optimal choice in *some* decision problem. The following example illustrates why, when the cognitive cost function is not convex, our order on contemplative ability may not be necessary to ensure the agent is better off in any decision problem.

Example 2. *Take the following two cognitive cost functions: $C(\rho) = \rho$ for all $\rho \in [0, 1]$ and $C'(\rho) = 0$ for $\rho = 0$ and $C'(\rho) = 1/2$ for $\rho > 0$. Clearly, C' is not convex and neither C nor C' displays greater contemplative ability than the other. Moreover, for both C and C', the agent finds it optimal to choose either $\rho = 0$ or $\rho = 1$ at any mental signal in any decision problem. Hence, it is without loss of generality to suppose the agent selects $\rho \in \{0, 1\}$ under either cognitive cost function. Since $C'(\rho) \leq C(\rho)$ for $\rho \in \{0, 1\}$, it follows that the agent is better off with C' for any $(A, u, \sigma, \Gamma, \mu)$.*

The results in Proposition 2 and Proposition 3 do not necessarily imply that an agent with lower cognitive costs contemplates more. However, if the cognitive cost function satisfies increasing differences (see Milgrom and Shannon (1994)), this is the case. Take a cognitive cost function indexed by a parameter $\kappa \in \mathbb{R}$, $C(\rho; \kappa)$, where $C(\rho; \kappa)$ is increasing in κ for all $\rho \in [0, 1]$. This function satisfies *increasing differences in* (ρ, κ) if $C(\rho; \kappa) - C(\rho'; \kappa)$ is increasing in κ for $\rho \geq \rho'$.

Proposition 4. *Suppose $C(\rho;\kappa)$ satisfies increasing differences in (ρ,κ). Then, for any $\sigma:\Omega\to\Delta(S)$ and $\Gamma_\sigma : S\to\Delta(\tilde{S})$, the agent contemplates with greater intensity at each $\tilde{s}\in\tilde{S}$ with κ than with $\kappa'>\kappa$.*

As an example, the quadratic cost function $C(\rho;\kappa)=\kappa\rho^2/2$ where $\kappa\geq 0$ satisfies increasing differences in (ρ,κ). The following example illustrates that the positive association between contemplation intensity and contemplative ability may fail when increasing differences is not satisfied.

Example 3. *Consider the decision problem from Example 1, where $\Omega = \{\omega_1,\omega_2\}$, $\mu(\omega_1)=1/2$, $A=\Omega$, utility given by (3), and $\sigma(s_i|\omega_i)=1$ for $i=1,2$. The agent has prior-based intuition: $\Gamma_\sigma(\tilde{s}|s_i)=1$ for $i=1,2$ (see Figure 2).*

Consider the following class of cognitive cost functions, indexed by the parameter $\kappa\in[0,1]$:

$$C(\rho;\kappa)=\begin{cases}0 & \text{if } \rho\leq 1-\kappa\\ 1/3 & \text{if } \rho>1-\kappa.\end{cases}$$

For a given κ, $C(\rho;\kappa)$ is such that the agent can contemplate up to intensity $1-\kappa$ for free, but incurs a fixed cognitive cost of $1/3$ if she contemplates with greater intensity. We compare the case of $\kappa=1/2$ to $\kappa=1$. Clearly, by Definition 2, $C(\rho;1/2)$ displays greater contemplative ability than $C(\rho;1)$. However, increasing differences fails.[6] *We now determine the level of contemplation chosen by each agent. For $C(\rho;1)$, the agent chooses between fully*

[6]Indeed, for $\rho=1$ and $\rho=1/2$,

$$C(1;1)-C(1/2;1)=0<C(1;1/2)-C(1/2;1/2)=1/3.$$

contemplating ($\rho = 1$) or no contemplation ($\rho = 0$). She chooses to fully contemplate as $1 - 1/3 > 1/2$. Instead, for $C(\rho; 1/2)$, the agent chooses between fully contemplating and the maximum contemplation level which is free (i.e. $\rho = 1/2$). She chooses $\rho = 1/2$ as $(1/2) + (1/2)(1/2) > 1 - 1/3$. Hence, the agent with higher cognitive costs contemplates with greater intensity and, as such, takes a Bayesian optimal action more often. Nonetheless, the agent that displays greater contemplative ability is better off.

4.3 Quality of Information

We now investigate the role that the informativeness of experiments plays in the agent's decision-making process. From Blackwell's Informativeness Theorem, it is known that a Bayesian agent prefers experiment σ to σ' for any (A, u, μ) if and only if σ' is a garbling of σ. Therefore, experiments ordered by garblings fully characterize the set of experiments that are unambiguously welfare-ranked for any Bayesian agent. This, however, may not be the case for our intuitive agent. Our first result characterizes the set of experiments that are uniformly welfare-ranked across all decision problems faced by any intuitive agent.

Proposition 5. *Take two experiments, $\sigma, \sigma' \in \mathcal{E}$, where $\sigma : \Omega \to \Delta(S)$ and $\sigma' : \Omega \to \Delta(S')$. The following two statements are equivalent:*

(a) σ' is an uninformative experiment, that is, $B(s', \sigma') = \mu$ for all $s' \in S'$.

(b) For all (A, u, C, Γ, μ), the agent prefers σ to σ'.

Suppose a third party, who is ignorant of (A, u, C, Γ, μ), wishes to know which of two experiments the agent prefers. Proposition 5 implies that this third party can only know with certainty that the intuitive agent prefers some information to no information. Example 4 provides intuition for why this is the case.

Example 4. *There are two states,* $\Omega = \{\omega_1, \omega_2\}$ *with prior* $\mu(\omega_1) = 1/2$. *The action space is* $A = \Omega$, *with the agent's state-dependent utility given by (3). Consider the following two experiments on signal space* $S = \{s_1, s_2\}$: $\sigma(s_i|\omega_i) = 1$ *for* $i = 1, 2$ *and* $\sigma'(s_i|\omega_i) = 3/4$ *for* $i = 1, 2$. *Clearly,* σ' *is a garbling of* σ *but* σ' *is not uninformative. A Bayesian agent prefers experiment* σ *to* σ' *by Blackwell's Informativeness Theorem. In contrast, consider an intuitive agent with intuition-generating process such that* $\Gamma_\sigma(\tilde{s}_\emptyset|s_i) = 1$ *for* $i = 1, 2$ *and* $\Gamma_{\sigma'}(s_i|s_i) = 1$ *for* $i = 1, 2$. *This process implies that (i)* $\Gamma_\sigma \circ \sigma$ *is an uninformative experiment and (ii)* $\Gamma_\sigma \circ \sigma$ *is a garbling of* $\Gamma_{\sigma'} \circ \sigma' = \sigma'$. *If* $C(\rho) = +\infty$, *then the agent always prefers (sometimes strictly)* σ' *to* σ *even though* σ' *is a garbling of* σ. *Only an uninformative* σ' *ensures that the agent never strictly prefers* σ' *to* σ.

Thus, the set of experiments which are unambiguously welfare-ranked is heavily restricted when nothing is known about the agent's intuition-generating process. As a result, Blackwell's Theorem does not hold for an intuitive agent. A natural question is whether there exists a restriction on the agent's intuition-generating process that restores Blackwell's theorem; that is, when does an intuitive agent prefer σ to σ' in any decision problem if and only if σ' is a garbling of σ? The answer to this question is affirmative. In order to describe this condition, we introduce some notation.

Take any experiment $\sigma : \Omega \to \Delta(S_\sigma)$ with associated intuition-generating process $\Gamma_\sigma : S_\sigma \to \Delta(\tilde{S}_\sigma)$. Let $T_\sigma \equiv S_\sigma \times \tilde{S}_\sigma$. Note that σ and Γ_σ jointly describe an experiment, π_σ, on T_σ. In particular, $\pi_\sigma(s_\sigma, \tilde{s}_\sigma|\omega) = \Gamma_\sigma(\tilde{s}|s)\sigma(s|\omega)$ for all $s_\sigma \in S_\sigma$, $\tilde{s}_\sigma \in \tilde{S}_\sigma$, and $\omega \in \Omega$. That is, $\pi_\sigma(s_\sigma, \tilde{s}_\sigma|\omega)$ is the probability that true signal s_σ and mental signal $\tilde{s}_\sigma$ jointly realize conditional on state ω. As such, π_σ measures the *pooled information* available to the agent. The next condition compares pooled information across experiments.

Definition 3 (Intuitive Sufficiency)**.** *Take two experiments σ and σ'. We say that σ is intuitively sufficient for σ' if there exists a $g : T_\sigma \to \Delta(T_{\sigma'})$ such that*

(a) $\pi_{\sigma'} = g \circ \pi_\sigma$; and

(b) $\sum_{s_{\sigma'} \in S_{\sigma'}} g(s_{\sigma'}, \tilde{s}_{\sigma'}|s_\sigma, \tilde{s}_\sigma)$ is independent of s_σ for all $\tilde{s}_{\sigma'} \in \tilde{S}_{\sigma'}$ and $\tilde{s}_\sigma \in \tilde{S}_\sigma$.

Intuitive sufficiency is closely related to the individual sufficiency condition provided in Bergemann and Morris (2016) and the notion of non-communicating garblings provided in Lehrer, Rosenberg, and Shmaya (2013), both defined in the context of many-player games.[7] Part (a) of intuitive sufficiency states that the pooled experiment for σ' is a garbling of σ. Part (b) of the condition states that, for this garbling, the marginal probability of $\tilde{s}_{\sigma'}$ given $(s_\sigma, \tilde{s}_\sigma)$ is independent of s_σ. Essentially, these two conditions jointly imply that σ is less noisy than σ' and the agent's intuition-generating process preserves this reduction in noise. Specifically, condition (a) ensures that σ' is a garbling of σ. Moreover, condition (a) and (b) jointly ensure that $\Gamma_{\sigma'} \circ \sigma'$ is a garbling of $\Gamma_\sigma \circ \sigma$. The condition, however, is slightly stronger than these two implications. In particular, it also places structure on *how* the intuition-generating process preserves the garbling of σ to σ'. Indeed, the requirement that the

[7]The main difference from Bergemann and Morris (2016) and Lehrer, Rosenberg, and Shmaya (2013) is that we require independence of the marginal of g only for mental signals and not for true signals. This is because, when our agent makes a decision at some true signal, she also knows which mental signal realized.

conditional marginal probability of $\tilde{s}_{\sigma'}$ is independent of s_σ ensures that the agent can (i) replicate $\Gamma_{\sigma'} \circ \sigma'$ from $\Gamma_\sigma \circ \sigma$ using a strategy measurable with respect to only her mental signals $\tilde{s}_\sigma \in \tilde{S}_\sigma$; and (ii) can do so in a way that maintains the (correct) perception that σ is "better" than σ'.

We say that the agent's intuition-generating process, Γ, *preserves the Blackwell order* if σ is intuitively sufficient for σ' whenever σ' is a garbling of σ. Note that there are intuition-generating processes that preserve the Blackwell order, so that this condition is non-vacuous. In particular, it is simple to verify that the ϕ-mixture between perfect and prior-based intuition described in Section 3.3 ensures intuitive sufficiency is satisfied whenever σ' is a garbling of σ. The following proposition shows that preservation of the Blackwell order is sufficient for Blackwell's theorem to hold with an intuitive agent.

Proposition 6. *Suppose that Γ preserves the Blackwell order. The following two statements are equivalent:*

(a) σ' is a garbling of σ.

(b) For any (A, u, C, μ), the agent prefers σ to σ'.

The direction (b) implies (a) is a direct implication of Blackwell's theorem if one takes $C(\rho) = 0$ for all $\rho \in [0, 1]$. Similar to the proof of Proposition 1, (a) implies (b) because intuitive sufficiency ensures that the agent facing σ can always replicate, for each state of the world, the distribution of (a, ρ) generated under σ'. The following example highlights why it is not sufficient to only require that σ' is a garbling of σ and $\Gamma_{\sigma'} \circ \sigma'$ is a garbling of $\Gamma_\sigma \circ \sigma$.

Example 5. *There are two states, $\Omega = \{\omega_1, \omega_2\}$, the prior is uniform, $A = \Omega$, and the agent's state-dependent utility is given by (3). Suppose that the cogni-*

tive cost function has $C(0) = 0$ *and* $C(\rho) = 3/8$ *for* $\rho \in (0, 1]$*. With this cost function, the agent either chooses* $\rho = 0$ *or* $\rho = 1$ *at each mental signal. Figure 4 displays two experiments and their associated intuition-generating processes.*

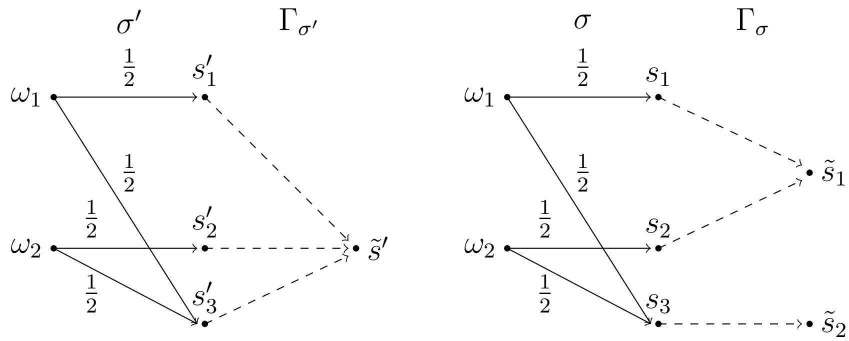

Figure 4: An intuition-generating process impacted by framing effects.

The experiments are equivalent in that they reveal exactly the same information. However, the agent's intuition is generated differently depending on precisely which experiment is used. This captures a situation where the agent's intuition is impacted by framing effects. For example, σ' *may represent a story on the nightly news and* σ *may represent the same story in the newspaper. Since* σ *and* σ' *are equivalent,* σ *is a garbling of* σ'*. Moreover, the agent's intuition is equal to the prior for any mental signal under either experiment. Hence,* $\Gamma_\sigma \circ \sigma$ *is a garbling of* $\Gamma_{\sigma'} \circ \sigma'$*. Nonetheless, the agent's payoff is strictly higher under* σ *than under* σ'*. Indeed,*

$$V(\Gamma, C, \sigma) = (1/2)(1 - 3/8) + (1/2)(1/2) = 9/16 > 1/2 = V(\Gamma, C, \sigma'),$$

where these payoffs arise from the fact that the agent contemplates at $\tilde{s}_1$ *under* σ *but never contemplates under* σ'*. Conversely, it is simple to verify that* σ *is intuitively sufficient for* σ'*, so that* σ *is preferred to* σ' *in any decision problem.*

5 Application: Bayesian Persuasion

We consider the canonical example of a Bayesian persuasion problem, introduced in Kamenica and Gentzkow (2011), with an intuitive receiver. There are many situations in which persuaders need to take into account the first impression or intuition of those people that they are trying to convince. For example, academics need to present their research in a way that ensures other researchers (who may only skim-read articles) are able to intuit the main message of the paper quickly and are convinced that the paper is worth reading in detail. Similarly, newspaper articles are often structured to maximize the amount of information the reader attains by quickly skimming the first few paragraphs (the so-called "onion method").

We illustrate that poor intuition can serve as a form of commitment power to take the sender's least preferred action while contemplation allows for flexibility to renege on this action if the sender transmits sufficiently convincing information. As a consequence, we find that cognitive limitations that are harmful in non-strategic situations can be beneficial in strategic settings.

We now describe the canonical model. A sender designs an experiment $\sigma \in \mathcal{E}$ for a receiver. There are two states, $\Omega = \{\omega_1, \omega_2\}$, with common prior $\mu \equiv \mu(\omega_1) > 1/2$, and two actions, $A = \Omega$. The sender's utility is state-independent and given by

$$\hat{u}(a, \omega) = \begin{cases} 1 & \text{if } a = \omega_2 \\ 0 & \text{if } a = \omega_1. \end{cases}$$

Instead, the receiver wants to take an action that matches the state; that is,

$$u(a,\omega) = \begin{cases} 1 & \text{if } a = \omega \\ 0 & \text{if } a \neq \omega. \end{cases}$$

To guarantee the existence of an optimal experiment for the sender, we assume that the receiver resolves indifference according to the sender's preference.

5.1 Optimal Experiment for a Bayesian Agent

As a benchmark, we derive an optimal experiment that the sender designs for a Bayesian receiver. Following the arguments in Kamenica and Gentzkow (2011), it is without loss of generality to focus on experiments with two signals and we treat an experiment as a distribution of posteriors that average to the prior (the "Bayes plausibility requirement"). For some experiment σ, let $\beta_1 \equiv B(\omega_1|s_1,\sigma)$, $\beta_2 \equiv B(\omega_2|s_2,\sigma)$, and $\alpha \equiv Pr(s_2|\sigma) = \mu\sigma(s_2|\omega_1)+(1-\mu)\sigma(s_2|\omega_2)$. The Bayes plausibility requirement is equivalent to $\alpha\beta_2+(1-\alpha)(1-\beta_1) = 1-\mu$. We assume the sender uses s_1 to convince the receiver to take action ω_1 (i.e $\beta_1 \geq 1/2$) and uses s_2 to convince the receiver to take action ω_2 (i.e. $\beta_2 \geq 1/2$). The sender's optimization problem is

$$\max_{\alpha,\beta_1,\beta_2} \alpha \text{ subject to } \beta_1 \geq 1/2,\ \beta_2 \geq 1/2, \text{ and Bayes plausibility.}$$

This program has the solution $\alpha = 2(1-\mu)$, $\beta_1 = 1$, and $\beta_2 = 1/2$. This solution determines the Bayesian optimal experiment which we denote by σ^{BP}.

Our intuitive receiver may face frictions in terms of her intuition and/or her ability to contemplate. However, in the absence of either of these frictions it is clear that the intuitive receiver will be indistinguishable from a Bayesian receiver and σ^{BP} is an optimal experiment. Proposition 7 states this observation without proof.

Proposition 7. *An optimal experiment for the principal is σ^{BP} if either the agent has perfect intuition or has maximal contemplative ability (that is, $C(\rho) = 0$ for all $\rho \in [0, 1]$).*

Note that the experiment σ^{BP} gives the receiver a payoff of μ which is the same payoff that she would attain under no information. It follows from Proposition 5 that, for any experiment, an intuitive receiver garners higher payoff than μ.

The Bayesian optimal experiment may not effectively convince the receiver when frictions on both intuition and contemplation are present. This is because action ω_1 is (weakly) optimal for all signal realizations. As such, an intuitive receiver never contemplates if her intuition directs her to take this action. The sender needs to take this into consideration when trying to convince the receiver.

5.2 An Intuitive Agent

We utilize the ϕ-mixture of perfect and prior-based intuition described in Section 3.3. We interpret a receiver with this intuition-generating process to be *conservative* since at each true signal she ignores its informational content with probability $1 - \phi$ and intuits to take an action consistent with the prior. Since this action is ω_1, the smaller is ϕ the more *commitment power* the receiver has to take the sender's least-preferred action.

We assume the agent's cognitive cost function takes the quadratic form $C(\rho;\kappa) = \kappa\rho^2/2$ where $\kappa \geq 0$. Since $C(\rho;\kappa)$ satisfies increasing differences in (ρ,κ), the lower is κ the more the receiver contemplates for a given perceived benefit from contemplation (see Proposition 4). Hence, κ captures the receiver's *flexibility* to deviate from what her intuition prescribes.

Under this cognitive process, it is without loss of generality for the sender to use an experiment that has only two signals.

Lemma 1. *Take any experiment $\sigma \in \mathcal{E}$. There exists an experiment $\sigma' : \Omega \to \Delta(\{s_1, s_2\})$ such that the sender's payoff is the same under both σ and σ'.*

5.3 An Optimal Experiment

We now state the optimal experiment chosen by the sender. Given Lemma 1, we can restrict attention to σ that generate only two posterior beliefs, $\beta_1 \equiv B(\omega_1|s_1,\sigma)$ or $\beta_2 \equiv B(\omega_2|s_2,\sigma)$. Let α be the probability that s_2 realizes under σ. Suppose, without loss of generality, the sender uses s_i to convince the receiver to take action ω_i (i.e. $\beta_i \geq 1/2$ for $i = 1,2$). Given the receiver's intuition-generating process, she will only perceive a benefit from contemplation if mental signal $\tilde{s}_\emptyset$ realizes. The following proposition describes the sender's optimal experiment, $(\alpha^*, \beta_1^*, \beta_2^*)$, as well as the receiver's optimal contemplation decision at $\tilde{s}_\sigma$, $\rho^*(\tilde{s}_\emptyset)$.

Proposition 8. *Let $\underline{\kappa}(\phi) \equiv \frac{2(1-\mu)(1-\phi)}{2-\phi} < \bar{\kappa}(\phi) \equiv \frac{2(1-\mu)(1-\phi)}{\phi}$. Then $\beta_1^* = 1$, $\beta_2^* = (1-\mu)/\alpha^*$, and*

(a) if $\kappa \leq \underline{\kappa}(\phi)$, $\alpha^* = 2(1-\mu) - \kappa$ and $\rho^*(\tilde{s}_\emptyset) = 1$, or

(b) if $\kappa \in (\underline{\kappa}(\phi), \bar{\kappa}(\phi))$, $\alpha^* = 1 - \mu + \frac{\kappa\phi}{2(1-\phi)}$ and $\rho^*(\tilde{s}_\emptyset) = \frac{2(1-\mu)-\alpha^*}{\kappa} \in (0,1)$, or

(c) if $\kappa \geq \bar{\kappa}(\phi)$, $\alpha^* = 2(1-\mu)$ *and* $\rho^*(\tilde{s}_\emptyset) = 0$.

Proposition 8 shows that there are three distinct regimes. Figure 5 illustrates these three regimes and the thresholds that define them. In any regime, $\beta_1^* = 1$. Therefore, the sender's optimal experiment is more convincing (i.e. reveals more information) the smaller is α^*. Part (a) of the proposition shows that, when κ and ϕ are sufficiently small, the sender's experiment actually becomes *more* informative (i.e. α^* decreases) as the agent's contemplative ability worsens (i.e. κ increases). This is because (i) when ϕ is small, the sender's benefits from incentivizing contemplation at mental signal $\tilde{s}_\emptyset$ are large, and (ii) κ is so small that it is worth incentizing the receiver to fully contemplate at $\tilde{s}_\emptyset$ (i.e. choose $\rho(\tilde{s}_\emptyset) = 1$). As the receiver's κ increases, so must the perceived benefits from contemplation (i.e. the experiment must become more informative).

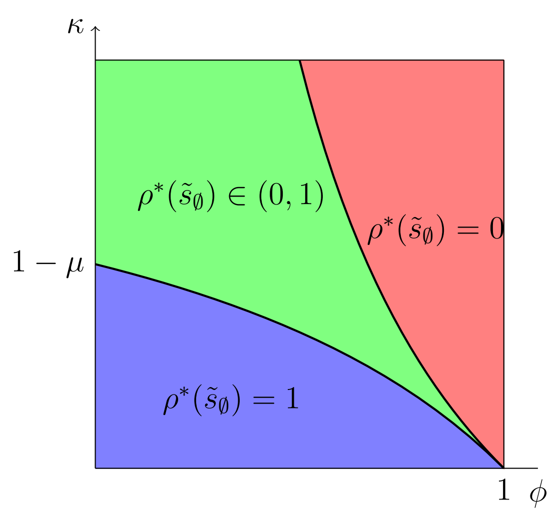

Figure 5: The sender's optimal experiment for all (ϕ, κ).

Part (c) shows that, when either κ or ϕ are sufficiently large, the sender chooses the Bayesian-optimal experiment σ^{BP}, where $\rho^*(\tilde{s}_\emptyset) = 0$ (see the discussion at the end of Section 5.1). This is because (i) when ϕ is large, the receiver intuits the true content of signals with high probability, and (ii) large κ implies it is not worthwhile for the sender to incentivize contemplation at $\tilde{s}_\emptyset$.

Finally, the regime described in part (b) is an intermediate case of those described in (a) and (c). In this region, the sender trades off the gains from reducing the informational content of the experiment (towards to σ^{BP}) against the costs of the receiver not valuing contemplation and, hence, using her intuition to make decisions. Thus, the optimal experiment induces an interior level of contemplation at $\tilde{s}_\emptyset$. It is interesting to note that the sender's optimal experiment actually becomes *more informative* (i.e. α^* decreases) as the receiver's quality of intuition *decreases* (i.e. ϕ decreases). The reason is that, as ϕ decreases, the receiver gains more *commitment power* to take the sender's least-preferred action, ω_1. Since κ is not too large, the receiver is sufficiently *flexible* in that she will contemplate and take a Bayesian-optimal action if the sender reveals convincing information in favor of action ω_2. The sender finds it optimal to reveal more information to undo this increased commitment power from lower ϕ.

The next proposition describes how the receiver's equilibrium payoff varies with ϕ and κ.

Proposition 9. *Let $V(\phi, \kappa)$ denote a (ϕ, κ)-receiver's equilibrium payoff where $\phi < 1$ and $\kappa > 0$. Then,*

(a) For $\kappa \leq \underline{\kappa}(\phi)$, V is strictly increasing in both ϕ and κ.

(b) For $\kappa \in (\underline{\kappa}(\phi), \bar{\kappa}(\phi))$, V is strictly decreasing in κ, and

(i) if $\kappa \leq 1-\mu$*, V is strictly decreasing in* ϕ*; or*

(ii) if $\kappa > 1-\mu$*, there exists* $\bar{\phi} \in (0, \bar{\kappa}^{-1}(\kappa))$ *such that V is increasing for* $\phi < \bar{\phi}$ *and decreasing for* $\phi > \bar{\phi}$*.*

(c) For $\kappa \geq \underline{\kappa}(\phi)$*, V is independent of both* ϕ *and* κ*.*

Figure 6 graphically displays how the receiver's utility varies with ϕ and κ. Part (a) of Proposition 9 shows that, when κ is sufficiently small, the receiver's equilibrium payoff is *increasing* as her contemplative ability declines (i.e. κ increases). This stands in stark contrast to Proposition 2 which states that, in non-strategic settings, an intuitive agent is *never* better off as her contemplative ability worsens. Here, instead, the receiver's payoff increases as the sender's optimal experiment reveals more information as κ increases in this region. Note that V is increasing in ϕ, which is consistent with Proposition 1, as the optimal experiment is independent of ϕ in this region.

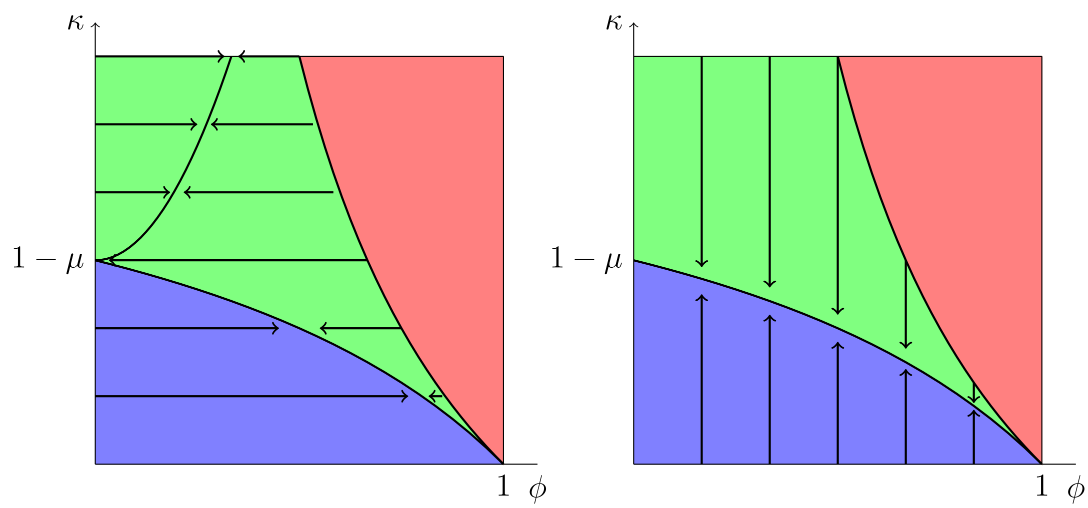

Figure 6: The figure illustrates how receiver utility varies with ϕ and κ. Utility increases in the direction an arrow points. The left panel shows how utility varies with ϕ. The right panel shows how utility varies with κ.

Part (b) of the proposition shows that, for any κ, there is a region over which the receiver's equilibrium utility increases as the quality of her intuition worsens (i.e. ϕ decreases). This fact stands in stark contrast to Proposition 1, which states that in a non-strategic context, any intuitive agent is *always* better off as the quality of her intuition improves. Here, instead, the receiver's payoff increases because the informativeness of the optimal experiment increases as ϕ decreases in this region (see the discussion below Proposition 8). As long as her intuition is not too poor and her cognitive costs are not too large, she reaps benefits from this improved information.

In any non-strategic context, the receiver would always want to have $\phi = 1$ (i.e. perfect intuition) and $\kappa = 0$ (i.e. maximal contemplative ability). However, the preceding analysis shows that the receiver can benefit from either poor intuition or low contemplative ability as both serve to increase commitment power to take undesirable actions from the sender's perspective. Indeed, the levels of ϕ and κ that maximize V are $\phi = 1 - \frac{\sqrt{3}}{2} \approx 0.13$ and $\kappa = \frac{2(1-\mu)\sqrt{3}}{2+\sqrt{3}} \approx 0.9(1-\mu)$. Hence, the receiver desires a quality of intuition which is far from perfect in this application.

6 Conclusion

We have introduced a theory of intuition and contemplation for an agent facing a decision problem under uncertainty. Our model is general and portable to any setting in which an individual needs to extract information from a source. We have provided natural orders on an agent's quality of intuition and her contemplative ability. These orders assure that an agent becomes better off in any non-strategic setting as either her intuition or her contemplative ability improves. Moreover, we have derived versions of Blackwell's Informativeness Theorem for intuitive agents.

Finally, we have demonstrated that embedding our intuitive agent in a strategic setting generates interesting insights. In particular, in a Bayesian persuasion setting an intuitive receiver can be better off the more cognitively limited she is. As such, bounds on either intuition or the contemplative process that hurt the individual in non-strategic settings may actually benefit her when information is strategically designed. This finding arises because poor intuition and low contemplative ability can both serve as sources of commitment power.

Bibliography

Alaoui, Larbi and Antonio Penta. 2016. "Endogenous Depth of Reasoning." *The Review of Economic Studies* 83 (4):1297–1333.

———. 2018. "Cost-Benefit Analysis in Reasoning." *Mimeo* .

Ambuehl, Sandro and Shengwu Li. 2018. "Belief Updating and the Demand for Information." *Games and Economic Behavior* 109:21–39.

Beauchêne, Dorian, Jian Li, and Ming Li. 2019. "Ambiguous Persuasion." *Journal of Economic Theory* 179:312–365.

Bergemann, Dirk and Stephen Morris. 2016. "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Economics* 11 (2):487–522.

Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *The Annals of Mathematical Statistics* :265–272.

Caplin, Andrew and Mark Dean. 2015. “Revealed Preference, Rational Inattention, and Costly Information Acquisition.” *American Economic Review* 105 (7):2183–2203.

de Clippel, Geoffroy and Xu Zhang. 2019. “Non-Bayesian Persuasion.” .

de Oliveira, Henrique. 2018. “Blackwell’s Informativeness Theorem Using Diagrams.” *Games and Economic Behavior* 109:126–131.

De Oliveira, Henrique, Tommaso Denti, Maximilian Mihm, and Kemal Ozbek. 2017. “Rationally Inattentive Preferences and Hidden Information Costs.” *Theoretical Economics* 12 (2):621–654.

Enke, Benjamin, Uri Gneezy, Brian J Hall, David Martin, Vadim Nelidov, Theo Offerman, and Jeroen van de Ven. 2020. “Cognitive Biases: Mistakes or Missing Stakes?” *CESifo Working Paper* .

Entman, Robert M. 1993. “Framing: Toward Clarification of a Fractured Paradigm.” *Journal of Communication* 43 (4):51–58.

Ergin, Haluk and Todd Sarver. 2010. “A Unique Costly Contemplation Representation.” *Econometrica* 78 (4):1285–1339.

Gennaioli, Nicola and Andrei Shleifer. 2010. “What Comes to Mind.” *The Quarterly Journal of Economics* 125 (4):1399–1433.

Jakobsen, Alex. 2020. “Coarse Bayesianism.” *Working Paper* .

Kamenica, Emir and Matthew Gentzkow. 2011. “Bayesian Persuasion.” *American Economic Review* 101 (6):2590–2615.

Kominers, Scott Duke, Xiaosheng Mu, and Alexander Peysakhovich. 2018. "Paying (for) Attention: The Impact of Information Processing Costs on Bayesian Inference." *SSRN Working Paper 2857978* .

Lehrer, Ehud, Dinah Rosenberg, and Eran Shmaya. 2013. "Garbling of Signals and Outcome Equivalence." *Games and Economic Behavior* 81:179–191.

Maćkowiak, Bartosz, Filip Matejka, and Mirko Wiederholt. 2018. "Survey: Rational Inattention, a Disciplined Behavioral Model." *CEPR Discussion Paper No. DP13243* .

Martin, Daniel. 2017. "Strategic Pricing with Rational Inattention to Quality." *Games and Economic Behavior* 104:131–145.

Matyskova, Ludmila and Alfonso Montes. 2018. "Bayesian Persuasion with Costly Information Acquisition." *CERGE-EI Working Paper Series* (614).

Milgrom, Paul and Chris Shannon. 1994. "Monotone Comparative Statics." *Econometrica* :157–180.

Mullainathan, Sendhil. 2002. "Thinking Through Categories." *NBER Working Paper* .

Price, Vincent, David Tewksbury, and Elizabeth Powers. 1997. "Switching Trains of Thought: The Impact of News Frames on Readers' Cognitive Responses." *Communication research* 24 (5):481–506.

Ravid, Doron. 2019. "Bargaining with Rational Inattention." *SSRN Working Paper 2957890* .

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3):665–690.

Tirole, Jean. 2009. "Cognition and Incomplete Contracts." *American Economic Review* 99 (1):265–94.

Tversky, Amos and Daniel Kahneman. 1981. "The Framing of Decisions and the Psychology of Choice." *Science* 211 (4481):453–458.

A Appendix

Proof of Proposition 1

Fix an experiment $\sigma : \Omega \to \Delta(S)$. Take two intuition-generating processes, Γ and Γ', where $\Gamma_\sigma : S \to \Delta(\tilde{S})$ and $\Gamma'_\sigma : S \to \Delta(\tilde{S}')$ and suppose that Γ displays better intuition at σ than Γ'. Take an optimal solution in pure strategies under Γ': $a'_s \in a^*(B(s,\sigma))$ for $s \in S$, $a'_{\tilde{s}'} \in a^*(B(\tilde{s}', \Gamma'_\sigma \circ \sigma))$ for $\tilde{s}' \in \tilde{S}'$, and $\rho'_{\tilde{s}'}$ solves equation (2) for all $\tilde{s}' \in \tilde{S}'$. This strategy generates ex-ante utility under Γ' equal to

$$\sum_{\omega \in \Omega} \mu(\omega) \sum_{s \in S} \sigma(s|\omega) \sum_{\tilde{s}' \in \tilde{S}'} \Gamma'_\sigma(\tilde{s}'|s) \left[\rho'_{\tilde{s}'} u(a'_s, \omega) + (1 - \rho'_{\tilde{s}'}) u(a'_{\tilde{s}'}, \omega) - C(\rho'_{\tilde{s}'})\right]. \tag{4}$$

Since Γ displays better intuition at σ than Γ', there exists a $g : \tilde{S} \to \Delta(\tilde{S}')$ such that $\Gamma'_\sigma = g \circ \Gamma$. Consider the following strategy under Γ: choose a'_s for each $s \in S$, and at $\tilde{s}$, choose $(\rho'_{\tilde{s}'}, a'_{\tilde{s}'})$ with probability $g(\tilde{s}'|\tilde{s})$. This strategy generates ex-ante utility under Γ equal to

$$\begin{aligned}
&\sum_{\omega \in \Omega} \mu(\omega) \sum_{s \in S} \sigma(s|\omega) \sum_{\tilde{s} \in \tilde{S}} \Gamma_\sigma(\tilde{s}|s) \sum_{\tilde{s}' \in \tilde{S}'} g(\tilde{s}'|\tilde{s}) \left[\rho'_{\tilde{s}'} u(a'_s, \omega) + (1 - \rho'_{\tilde{s}'}) u(a'_{\tilde{s}'}, \omega) - C(\rho'_{\tilde{s}'})\right] \\
=& \sum_{\omega \in \Omega} \mu(\omega) \sum_{s \in S} \sigma(s|\omega) \sum_{\tilde{s}' \in \tilde{S}'} \left[\rho'_{\tilde{s}'} u(a'_s, \omega) + (1 - \rho'_{\tilde{s}'}) u(a'_{\tilde{s}'}, \omega) - C(\rho'_{\tilde{s}'})\right] \sum_{\tilde{s} \in \tilde{S}} \Gamma_\sigma(\tilde{s}|s) g(\tilde{s}'|\tilde{s}) \\
=& \sum_{\omega \in \Omega} \mu(\omega) \sum_{s \in S} \sigma(s|\omega) \sum_{\tilde{s}' \in \tilde{S}'} \Gamma'_\sigma(\tilde{s}'|s) \left[\rho'_{\tilde{s}'} u(a'_s, \omega) + (1 - \rho'_{\tilde{s}'}) u(a'_{\tilde{s}'}, \omega) - C(\rho'_{\tilde{s}'})\right],
\end{aligned}$$

where the last equality follows from the fact that

$$\sum_{\tilde{s} \in \tilde{S}} g(\tilde{s}'|\tilde{s}) \Gamma_\sigma(\tilde{s}|s) = g \circ \Gamma_\sigma(\tilde{s}'|\tilde{s}) = \Gamma'_\sigma(\tilde{s}'|s).$$

Since this strategy gives utility under Γ equal to (4), an optimal strategy under Γ must give weakly larger utility than that under Γ'. ■

Proof of Proposition 2

Fix $\sigma : \Omega \to \Delta(S)$ and $\Gamma_\sigma : S \to \Delta(\tilde{S})$. Take two cognitive cost functions, C and C', such that $C(\rho) \leq C'(\rho)$ for all $\rho \in [0,1]$. Take an optimal pure strategy under Γ': $a'_s \in a^*(B(s,\sigma))$ for $s \in S$, $a'_{\tilde{s}} \in a^*(B(\tilde{s}, \Gamma \circ \sigma))$ for $\tilde{s} \in \tilde{S}$, and $\rho'_{\tilde{s}}$ solves equation (2) for all $\tilde{s} \in \tilde{S}$. This same strategy is feasible under C but yields lower cognitive costs. Thus, when the agent optimizes with C, ex-ante utility is higher under C than under C'. ■

Proof of Proposition 3

(a) $\Rightarrow$ (b): This is a direct implication of Proposition 2.

(b) $\Rightarrow$ (a): Suppose that (a) is violated; that is, there exists a $\hat{\rho} \in (0,1]$ such that $C(\hat{\rho}) > C'(\hat{\rho})$. We construct a profile $(A, u, \sigma, \Gamma, \mu)$ such that the agent is strictly better off under C' than under C. Suppose that $\mu(\omega_i) = 1/n$ for $i = 1, \ldots n$ and that σ is fully revealing; i.e. $\sigma(s_i|\omega_i) = 1$ for all $i = 1, \ldots n$. The action space is $A = \{\omega_1, \omega_2\}$ and the agent has utility given by

$$u(a,\omega) = \begin{cases} nx & \text{if } a = \omega \\ 0 & \text{otherwise} \end{cases}$$

where $x \geq 0$. The intuition-generating process is prior-based intuition; that is, $\Gamma_\sigma(\tilde{s}|s_i) = 1$ for $i = 1, \ldots, n$. We show that, by varying x, we can make the agent with C choose $\hat{\rho}$ at $\tilde{s}$. Define $\underline{z} \equiv \lim_{\rho \to \hat{\rho}^-} dC(\rho)/d\rho$ to be the left-derivative of C at $\hat{\rho}$ and $\bar{z} \equiv \lim_{\rho \to \hat{\rho}^+} dC(\rho)/d\rho$ to be the right-derivative of C at $\hat{\rho}$. Since C is convex and weakly increasing, $0 \leq \underline{z} \leq \bar{z}$. Choose $x \in [\underline{z}, \bar{z}]$. In the considered decision problem, the benefits from contemplation are $\rho \times 2x + (1 - \rho)x$ which implies that the marginal benefit from contemplation is x. For $\rho < \hat{\rho}$, the marginal cost of contemplation is less than $\underline{z}$ which itself is less than x. Hence, an optimal ρ must be weakly larger than $\hat{\rho}$. Similarly, for $\rho > \hat{\rho}$, the marginal cost of contemplation is greater than $\bar{z}$ which is itself larger than x. Hence, $\hat{\rho}$ is optimal under C. Since $C(\hat{\rho}) > C'(\hat{\rho})$, the agent must be strictly better off under C' than under C. ■

Proof of Proposition 4

Fix a $\sigma : \Omega \to \Delta(S)$ and $\Gamma_\sigma : S \to \Delta(\tilde{S})$. The contemplation decision at some $\tilde{s} \in \tilde{S}$ is to choose $\rho \in [0, 1]$ to maximize

$$f(\rho, -\kappa) \equiv m_0 + m_1\rho - C(\rho; \kappa); \quad \text{where}$$

$$m_0 \equiv U(a_{\tilde{s}} | B(\tilde{s}, \Gamma_\sigma \circ \sigma) \text{ and } m_1 \equiv \sum_{s \in S} Pr(s|\tilde{s}, \sigma, \Gamma_\sigma) U(a_s | B(s, \sigma)) - U(a_{\tilde{s}} | B(\tilde{s}, \Gamma_\sigma \circ \sigma)$$

are independent of ρ and κ. Note that $f(\rho, -\kappa)$ satisfies increasing differences in $(\rho, -\kappa)$. Indeed, it is simple to verify that $f(\rho, -\kappa) - f(\rho', -\kappa)$ is increasing in $-\kappa$ for $\rho \geq \rho'$ if and only if $C(\rho, \kappa) - C(\rho', \kappa)$ is increasing in κ. Since $f(\rho, -\kappa)$ is supermodular in ρ as $\rho \in [0, 1]$, the result follows immediately from Theorem 5 in Milgrom and Shannon (1994). ■

Proof of Proposition 5

(a) ⇒ (b): Let $a'_\mu \in a^*(\mu)$ be the action the agent chooses under the prior. Since σ' is uninformative, ex-ante utility under σ' is equal to $U(a'_\mu|\mu)$. Under $\Gamma_\sigma \circ \sigma : \Omega \to \Delta(\tilde{S})$, suppose that, at every mental signal $\tilde{s}$ the agent chooses $\rho = 0$ and the action a'_μ. Clearly, this gives ex-ante utility $U(a'_\mu|\mu)$. Hence, when the agent optimizes at σ, she cannot be worse off than under σ'.

(b) ⇒ (a): Suppose that the agent values σ more than σ' for any (A, u, C, Γ, μ). Then she must be better off if $\Gamma_\sigma \circ \sigma$ is an uninformative experiment (i.e. prior-based intuition at σ) and $\Gamma_{\sigma'} \circ \sigma' = \sigma'$ (i.e. perfect intuition at σ'). Suppose $C(\rho) = +\infty$ for all $\rho \in (0, 1]$ such that the agent is unable to contemplate. Because she is better off with σ over σ' for any (A, u, μ) given this Γ and C, by Blackwell (1953), it follows that σ' must be a garbling of $\Gamma_\sigma \circ \sigma$. Finally, since $\Gamma_\sigma \circ \sigma$ is uninformative, σ' must also be uninformative. ■

Proof of Proposition 6

(a) ⇒ (b): Fix two experiments $\sigma : \Omega \to \Delta(S)$ and $\sigma' : \Omega \to \Delta(S')$, with corresponding intuition-generating processes $\Gamma_\sigma : S \to \Delta(\tilde{S})$ and $\Gamma_{\sigma'} : \tilde{S} \to \Delta(\tilde{S}')$. Take an optimal pure strategy under σ': $a'_{s'} \in a^*(B(s', \sigma'))$ for $s' \in S'$, $a'_{\tilde{s}'} \in a^*(B(\tilde{s}', \Gamma_{\sigma'} \circ \sigma'))$ for $\tilde{s}' \in \tilde{S}'$, and $\rho'_{\tilde{s}'}$ solves equation (2) for all $\tilde{s}' \in \tilde{S}'$. This strategy generates ex-ante utility under σ' equal to

$$\sum_{\omega\in\Omega} \mu(\omega) \sum_{s'\in S'} \sum_{\tilde{s}'\in\tilde{S}'} \pi_{\sigma'}(s', \tilde{s}'|\omega) \left[\rho'_{\tilde{s}'} u(a'_{s'}, \omega) + (1 - \rho'_{\tilde{s}'}) u(a'_{\tilde{s}'}, \omega) - C(\rho'_{\tilde{s}'})\right]. \quad (5)$$

Since σ is intuitively sufficient for σ', there exists a $g : T_\sigma \to \Delta(T_{\sigma'})$ such that $\pi_{\sigma'} = g \circ \pi_\sigma$ and $\sum\limits_{s' \in S'} g(s', \tilde{s}'|s, \tilde{s})$ is independent of s for all $\tilde{s}' \in \tilde{S}'$ and $\tilde{s} \in \tilde{S}$. Let $\tilde{g}(\tilde{s}'|\tilde{s})$ denote this sum. Consider the following strategy under σ: first, at each mental signal $\tilde{s}$, the agent samples $\tilde{s}'$ according to distribution $\tilde{g}(\tilde{s}'|\tilde{s})$. If $\tilde{s}'$ realizes, she chooses $\rho'_{\tilde{s}'}$. If contemplation is unsuccessful, she chooses $a'_{\tilde{s}'}$. Instead, if contemplation is successful and she observes signal s, she chooses $a'_{s'}$ with probability $\frac{g(s', \tilde{s}'|s, \tilde{s})}{\sum_{s' \in S'} g(s', \tilde{s}'|s, \tilde{s})} = \frac{g(s', \tilde{s}'|s, \tilde{s})}{\tilde{g}(\tilde{s}'|\tilde{s})}$. Note that this strategy is measurable with respect to information that the agent has at each point in time when making decisions under σ. This strategy generates ex-ante utility under σ equal to

$$\begin{aligned}
&\sum_{\omega \in \Omega} \mu(\omega) \sum_{s \in S} \sum_{\tilde{s} \in \tilde{S}} \pi_\sigma(s, \tilde{s}|\omega) \sum_{\tilde{s}' \in \tilde{S}'} \tilde{g}(\tilde{s}'|\tilde{s}) \\
&\quad \times \left[\rho'_{\tilde{s}'} \left[\sum_{s' \in S'} \frac{g(s', \tilde{s}'|s, \tilde{s})}{\tilde{g}(\tilde{s}'|\tilde{s})} u(a'_{s'}, \omega) \right] + (1 - \rho'_{\tilde{s}'}) u(a'_{\tilde{s}'}, \omega) - C(\rho'_{\tilde{s}'}) \right] \\
= &\sum_{\omega \in \Omega} \mu(\omega) \sum_{s' \in S'} \sum_{\tilde{s}' \in \tilde{S}'} \left[\rho'_{\tilde{s}'} u(a'_{s'}, \omega) + (1 - \rho'_{\tilde{s}'}) u(a'_{\tilde{s}'}, \omega) - C(\rho'_{\tilde{s}'}) \right] \sum_{s \in S} \sum_{\tilde{s} \in \tilde{S}} g(s', \tilde{s}'|s, \tilde{s}) \pi_\sigma(s, \tilde{s}|\omega) \\
= &\sum_{\omega \in \Omega} \mu(\omega) \sum_{s' \in S'} \sum_{\tilde{s}' \in \tilde{S}'} \pi_{\sigma'}(s', \tilde{s}'|\omega) \left[\rho'_{\tilde{s}'} u(a'_{s'}, \omega) + (1 - \rho'_{\tilde{s}'}) u(a'_{\tilde{s}'}, \omega) - C(\rho'_{\tilde{s}'}) \right],
\end{aligned}$$

where the last equality follows from the fact that

$$\sum_{s \in S} \sum_{\tilde{s} \in \tilde{S}} g(s', \tilde{s}'|s, \tilde{s}) \pi_\sigma(s, \tilde{s}|\omega) = g \circ \pi_\sigma(s', \tilde{s}'|\omega) = \pi_{\sigma'}(s', \tilde{s}'|\omega).$$

Since this strategy gives utility under σ equal to (5), an optimal strategy under σ must give weakly larger utility than that under σ'.

(b) $\Rightarrow$ (a): Suppose the agent prefers σ to σ' for every (A, u, C, μ). This implies that the agent prefers σ to σ' when $C(\rho) = 0$ for all $\rho \in [0, 1]$. Under this cost function, she achieves the same payoff as a Bayesian agent. Hence, σ is preferred to σ' for a Bayesian agent with any (A, u, μ). By Blackwell (1953), this implies that σ' is a garbling of σ. ■

Proof of Lemma 1

Take an arbitrary experiment, $\sigma : \Omega \to \Delta(S)$. Define $S_{\omega_1} \equiv \{s \in S : B(\omega_1|s,\sigma) > 1/2\}$ and $S_{\omega_2} \equiv S \setminus S_{\omega_1}$. Consider the following two-signal experiment, $\sigma' : \Omega \to \Delta(\{s'_1, s'_2\})$, where $\sigma'(s'_1|\omega) = \sum_{s \in S_{\omega_1}} \sigma(s|\omega)$ and $\sigma'(s'_2|\omega) = \sum_{s \in S_{\omega_2}} \sigma(s|\omega)$. Note that, in both experiments, the probability the receiver takes action ω_2 without contemplation is the same. Moreover, with probability $1-\phi$ the receiver observes mental signal $\tilde{s}_\emptyset$ in both experiments. Hence, if the optimal choice of contemplation is the same at $\tilde{s}_\emptyset$ across σ and σ', the sender's payoff is also the same. At this mental signal, the benefits from contemplation are the same across both experiments. Indeed, for a given ρ, the benefits from contemplation under σ are

$$
\begin{aligned}
&\rho \left[\sum_{s \in S_{\omega_1}} B(\omega_1|s,\sigma) \left(\sum_{\omega \in \Omega} \sigma(s|\omega)\mu(\omega) \right) + \sum_{s \in S_{\omega_2}} B(\omega_2|s,\sigma) \left(\sum_{\omega \in \Omega} \sigma(s|\omega)\mu(\omega) \right) \right] + (1-\rho)\mu \\
&= \rho \left[\sum_{s \in S_{\omega_1}} \sigma(s|\omega_1)\mu(\omega_1) + \sum_{s \in S_{\omega_2}} \sigma(s|\omega_2)\mu(\omega_2) \right] + (1-\rho)\mu \\
&= \rho \left[\sigma'(s'_1|\omega_1)\mu(\omega_1) + \sigma'(s'_2|\omega_2)\mu(\omega_2) \right] + (1-\rho)\mu \\
&= \rho \left[B(\omega_1|s'_1,\sigma') \sum_{\omega \in \Omega} \sigma'(s'_1|\omega)\mu(\omega) + B(\omega_2|s'_2,\sigma') \sum_{\omega \in \Omega} \sigma'(s'_2|\omega)\mu(\omega) \right] + (1-\rho)\mu
\end{aligned}
$$

which is the benefit of contemplation at $\tilde{s}_\emptyset$ under σ'. Hence, the optimal choice of ρ at $\tilde{s}_\emptyset$ is the same under both σ and σ'. ■

Proof of Proposition 8

Recall that the receiver only perceives a benefit from contemplation at mental signal $\tilde{s}_\emptyset$. Given an arbitrary $(\alpha, \beta_1, \beta_2)$, where $\beta_i \geq 1/2$, $i = 1, 2$, the receiver chooses $\rho^*(\tilde{s}_\emptyset)$ as the solution to

$$\max_{\rho \in [0,1]} \rho \left[(1-\alpha)\beta_1 + \alpha\beta_2\right] + (1-\rho)\left[(1-\alpha)\beta_1 + \alpha(1-\beta_1)\right] - \kappa\frac{\rho^2}{2}$$

which, in a small abuse of notation, has solution $\rho^*(\alpha, \beta_2) = \min\left\{\frac{\alpha}{\kappa}(2\beta_2 - 1), 1\right\}$. The optimization problem of the sender is

$$\max_{\alpha, \beta_1, \beta_2} \alpha \left[\phi + (1-\phi)\rho^*(\alpha, \beta_2)\right]$$

subject to $\beta_1 \geq 1/2$, $\beta_2 \geq 1/2$, $\alpha\beta_2 + (1-\alpha)(1-\beta_1) = 1-\mu$. Note first that $\beta_1 = 1$ at an optimum. If not, the sender can jointly increase β_1 and α so that Bayes plausibility remains satisfied. This (weakly) increases $\rho^*(\alpha, \beta_2)$, which implies the value of the objective strictly increases. Hence, $\alpha\beta_2 = 1-\mu$ and we can re-write the sender's problem as

$$\max_{\alpha \in [(1-\mu), 2(1-\mu)]} \alpha \left[\phi + (1-\phi)\min\left\{\frac{2(1-\mu)-\alpha}{\kappa}, 1\right\}\right].$$

The objective is either linear or quadratic in α. Therefore, there are three candidate solutions: (1) $\alpha = 2(1-\mu) - \kappa$ (i.e. the α such that $\rho^*(\alpha) = \frac{2(1-\mu)-\alpha}{\kappa} = 1$), (2) $\alpha = 1 - \mu + \frac{\kappa\phi}{2(1-\phi)}$ (i.e. the unconstrained solution to the problem with $\rho^*(\alpha) = \frac{2(1-\mu)-\alpha}{\kappa}$), and (3) $\alpha = 2(1-\mu)$ (i.e. α attains its upper bound $\rho^*(\alpha) = 0$). (2) strictly dominates (3) if and only if

$$1 - \mu + \frac{\kappa\phi}{2(1-\phi)} < 2(1-\mu) \quad \Leftrightarrow \quad \kappa < \frac{2(1-\mu)(1-\phi)}{\phi} \equiv \underline{\kappa}(\phi).$$

Instead, (2) strictly dominates (1) if and only if

$$1 - \mu + \frac{\kappa\phi}{2(1-\phi)} > 2(1-\mu) - \kappa \quad \Leftrightarrow \quad \kappa > \frac{2(1-\mu)(1-\phi)}{2-\phi} \equiv \underline{\kappa}(\phi).$$

Since $\underline{\kappa}(\phi) < \bar{\kappa}(\phi)$, it follows that

$$\alpha^* = \begin{cases} 2(1-\mu) - \kappa & \text{if } \kappa \leq \underline{\kappa}(\phi) \\ 1 - \mu + \frac{\kappa\phi}{2(1-\phi)} & \text{if } \kappa \in (\underline{\kappa}(\phi), \bar{\kappa}(\phi)) \\ 2(1-\mu) & \text{if } \kappa \geq \bar{\kappa}(\phi), \end{cases}$$

and substituting α^* into $\rho^*(\alpha)$ gives the desired result. ■.

Proof of Proposition 9

Case (a): Suppose that $\kappa \leq \underline{\kappa}(\phi)$. In this case the optimal experiment induces $\rho^*(\tilde{s}_\emptyset) = 1$. The receiver's equilibrium utility, $V(\phi, \kappa)$, is given by

$$V(\phi, \kappa) = (1-\alpha^*)\phi + \alpha^*\phi\beta_2^* + (1-\phi)\left[(1-\alpha^*)\phi + \alpha^*\phi\beta_2^* - \frac{\kappa}{2}\right].$$

Substituting $\alpha^*\beta_2^* = 1 - \mu$, $\alpha^* = 2(1-\mu) - \kappa$ and rearranging gives

$$V(\phi, \kappa) = \mu + \frac{\kappa(1+\phi)}{2}.$$

Clearly, this is increasing in κ and ϕ.

Case (b): Suppose that $\kappa \in (\underline{\kappa}(\phi), \bar{\kappa}(\phi))$. In this case, the optimal experiment induces $\rho^*(\tilde{s}_\emptyset) = (2(1-\mu) - \alpha^*)/\kappa$ where $\rho^*(\tilde{s}_\emptyset) \in (0,1)$. The receiver's equilibrium utility can be written as

$$\begin{aligned} V(\phi,\kappa) &= (1-\alpha^*)\phi + \alpha^*\phi\beta_2^* + (1-\phi)\left[\rho^*(\tilde{s}_\emptyset)\left((1-\alpha^*)\phi + \alpha^*\phi\beta_2^*\right) + (1-\rho^*(\tilde{s}_\emptyset))\mu - \kappa\frac{\rho^*(\tilde{s}_\emptyset)^2}{2}\right] \\ &= \mu + \phi\kappa\rho^*(\tilde{s}_\emptyset) + (1-\phi)\kappa\frac{\rho(\tilde{s}_\emptyset)^2}{2}. \end{aligned}$$

Substituting $\alpha^* = 1 - \mu + \frac{\kappa\phi}{2(1-\phi)}$ into $\rho^*(\tilde{s}_\emptyset)$ and rearranging gives

$$V(\phi,\kappa) = \frac{\phi(1-\mu)}{2} + \frac{(1-\mu)^2(1-\phi)}{2\kappa} - \frac{3}{8}\frac{\kappa\phi^2}{1-\phi}.$$

Clearly, $V(\phi,\cdot)$ is decreasing in κ. Computing the partial derivative of V with respect to ϕ gives

$$\frac{\partial V}{\partial \phi} = \frac{1-\mu}{2} - \frac{(1-\mu)^2}{2\kappa} - \frac{3}{8}\frac{\kappa\phi(2-\phi)}{(1-\phi)^2}.$$

Define $\underline{\phi}(\kappa) \equiv \max\{0, \underline{\kappa}^{-1}(\kappa)\}$. Note that $\frac{\partial V}{\partial \phi}$ is strictly decreasing in ϕ and is strictly negative at $\bar{\kappa}^{-1}(\kappa)$. Hence, if it is negative at $\phi = 0$, this implies that the receiver's equilibrium is decreasing for all $\phi \in \left(\underline{\phi}(\kappa), \bar{\kappa}^{-1}(\kappa)\right)$. This is true if $\frac{1-\mu}{2} - \frac{(1-\mu)^2}{2\kappa} \leq 0$ or $\kappa \leq 1-\mu$. Instead, for $\kappa > 1-\mu$, we have $\underline{\phi}(\kappa) = 0$ and, by the intermediate value theorem, there exists a unique $\bar{\phi} \in (0, \bar{\kappa}^{-1}(\kappa))$ such that $V(\cdot,\kappa)$ is increasing for $\phi < \bar{\phi}$ and decreasing for $\phi > \bar{\phi}$.

Case (c): Suppose that $\kappa \geq \bar{\kappa}(\phi)$. In this case the sender uses the Bayesian optimal experiment and $\rho^*(\tilde{s}_\emptyset) = 0$. It follows that the receiver's equilibrium utility is given by

$$V(\phi,\kappa) = (1-\alpha^*)\phi + \alpha^*\phi\beta_2^* + (1-\phi)\left(1-\alpha^* + \alpha^*(1-\beta_2^*)\right).$$

Substituting $\alpha^*\beta_2^* = 1-\mu$ and $\alpha^* = 2(1-\mu)$ and rearranging yields

$$V(\phi,\kappa) = (1-2(1-\mu))\phi + \phi(1-\mu) + (1-\phi)\mu = \mu$$

which is independent of ϕ and κ. ■