ISSN 2279-9362

Collegio Carlo Alberto

Identity and the cost of information

Daniele Pennesi

No. 610
June 2020

Carlo Alberto Notebooks

www.carloalberto.org/research/working-papers

Identity and the cost of information*

Daniele Pennesi†

December 2020

Abstract

This paper proposes a rational theory of information avoidance based exclusively on the desire to protect identity, a person's sense of self. Identity determines prescriptions: appropriate behaviors given the identity. When there is uncertainty about the consequences of a behavior, the possibility to acquire information creates a trade-off. Information is valuable because it is instrumental, but it is costly because it can recommend a behavior that violates the prescriptions. If the instrumental value of information is lower than its cost, ignorance is optimal. The theory rationalizes information avoidance observed in different domains without resorting to cognitive constraints, motivated beliefs, or time-inconsistency. A key prediction of the theory, a potential desire to have fewer options, provides an identity-based explanation for excess entry in competitive environments, under-use of flexible work arrangements, and ethnic-specific differences in human capital accumulation. The preference for commitment is also consistent with a form of information-driven moral licensing, in which the individual commits to prescriptively inferior behaviors. Optimal information disclosure to an identity-caring individual may be distant from full information.

KEYWORDS: IDENTITY, SELF-IMAGE, INFORMATION AVERSION, PREFERENCE FOR COMMITMENT
JEL CLASSIFICATION: D01, D83, D91

1 Introduction

The desire to protect identity, a person's sense of self, is a major driver of human behavior (Akerlof and Kranton, 2000; Bénabou and Tirole, 2011). When deciding, identity-caring individuals value the material consequence of an action as well as how the action stands in relation to "who they are". The most direct and unequivocal way to protect identity is to follow the *prescriptions*: appropriate behaviors given the identity.[1] An individual willing to protect her

*I thank seminar partecipants at SAET Conference 2019 (Ischia), Center for Logic, Language and Cognition (Torino), and Bologna.

†ESOMAS Department, University of Torino. Corso Unione 218bis, Torino (Italy). Email: daniele.pennesi@unito.it

[1] According to Akerlof and Kranton (2000) "Prescriptions indicate the behavior appropriate for people in different social categories in different situations. The prescriptions may also describe an ideal for each category in

self-image of being an altruist, for example, will follow the prescription to make a donation when asked. In case there is uncertainty about the outcomes of a decision and acquiring information resolves part or all of the uncertainty, the desire to protect identity generates a trade-off. On the one hand, information has an instrumental value, since it helps making better decisions. On the other hand, information can be "cognitively costly", since it can recommend a behavior that violates the prescriptions.[2] When the cost outweighs the instrumental value of information, the individual will prefer to remain ignorant. To protect her altruistic identity, an individual can prefer avoid learning the quality of a charity (e.g. high or low) prior to make a donation. The prescription is to donate, but learning that the charity is of low quality makes not donating the optimal choice, a behavior that will violate the prescriptions. Anticipating the cost of learning, she may prefer to "close her eyes and do good".[3]

In this paper, I consider the far-reaching implications of the above intuition and develop a theory of information aversion based exclusively on the desire to protect identity. I consider a dynamic choice problem where an individual selects her future opportunities (a menu of actions whose payoffs depend on an unknown state of nature), acquire information or remain ignorant, and lastly selects an action from the menu based on the acquired information. The individual is "rational and sophisticated":[4] she is able to anticipate her desire to preserve identity; she processes information in a Bayesian way; (in the baseline model) she uses information to select the best action[5] and, she values actions according to their expected utility.

The key assumption is that the desire to preserve identity generates an *action-driven cost of information*. Information becomes costly when it creates a divergence between the prescribed action and the optimal action given information (the one that will be implemented). When the cost of information exceeds its instrumental value, information avoidance becomes optimal. In

terms of physical characteristics and other attributes....agents follow prescriptions, for the most part, to maintain their self-concepts...violating the prescriptions evokes anxiety and discomfort in oneself..." and according to Stets and Burke (2000) "In identity theory, the core of an identity is the categorization of the self as an occupant of a role, and the incorporation, into the self, of the meanings and expectations associated with that role... These expectations and meanings form a set of standards that guide behavior."

[2]Violating the prescription is costly, for example, because it can generate shame or a sense of guilt. In a social situation, violations of the prescription can lead to retaliation from others.

[3]Empirical evidence shows that few donors acquire information before making a donation (e.g. Niehaus, 2014).

[4]Empirical evidence that partly supports the sophistication assumption can be found in Kahan (2013). In his experiment subjects with a higher ability to protect their identity scored higher in a cognitive reflection test, suggesting that identity-protective behavior requires cognitive abilities.

[5]In an extension of the baseline model illustrated in Section 7.1, I consider the possibility that information is only partially used to select the best action.

a sense, the prescribed action represents the menu-dependent ideal or reference behavior and deviations from the reference due to information are costly (to the same extent deviations from the reference are costly in models of reference-dependent preferences, e.g. Loomes and Sugden, 1986; Köszegi and Rabin, 2006). However, costly deviations from the prescription/reference do not affect the value of an action, but only the value of a menu. The theory rationalizes information avoidance concerning the effects of harmful consumption (Viscusi, 1990), willful ignorance of work incentives (Huck, Szech, and Wenner, 2018), avoidance of routine health screenings (Courtenay, 2000) and pluralistic ignorance (Bursztyn, González, and Yanagizawa-Drott, 2020). Differently from the approach of Bénabou and Tirole (2011), in the present theory identity protection induces information avoidance without resorting to a combination of bounded memory about identity and time inconsistency. The former aspect, bounded memory, seems particularly hard to justify for "stable identities" such as gender or political identity. More generally, the present theory rationalizes information avoidance by a rational cost-benefit analysis, hence in the absence of cognitive constraints, motivated beliefs or any other non-standard assumption on the evaluation of actions.[6] The action-dependent cost of information in the present theory makes it behaviorally distinguishable from models based on motivated beliefs or anticipatory feelings. In the present theory, if information modifies the beliefs but the information-optimal action is still consistent with the prescriptions, information is costless. As a consequence, "dangerous thoughts" become taboos only if they trigger a costly switch in behavior. With motivated beliefs or anticipatory feelings, even the information that alters beliefs but not the recommended action can be costly.

A further prediction of the theory is that an identity-caring individual can have a strict preference for having fewer options before acquiring information. Indeed, the option value of larger menus can be outweighed by the fact that information is less likely to make the prescribed behavior suboptimal in smaller menus. Therefore, an identity-caring individual may find optimal to eliminate actions that makes information costly (see Section 5). This could also lead to eliminate prescriptively optimal actions when commitment to the prescription is unfeasible. Hence generating an information-driven form of "moral licensing". For example, the protec-

[6] The proposed "behavioral" rationalizations of information avoidance or aversion are numerous: cognitive constraints (Sims, 2003), motivated beliefs (Bénabou, 2015; Köszegi, 2006), anticipatory feelings (Caplin and Leahy, 2001) and (Brunnermeier and Parker, 2005), disappointment aversion (Andries and Haddad, 2019), and sophisticated time-inconsistency (Carrillo and Mariotti, 2000). In their review, Golman, Hagmann, and Loewenstein (2017) list 13 non-strategic motivations for information avoidance.

tion of an altruistic identity is consistent with escaping social dilemmas (e.g. DellaVigna, List, and Malmendier, 2012; Andreoni, Rao, and Trachtman, 2017). Entering the dilemma requires choosing between donating or not. Ideally, the individual would like to commit to donate. If this is unfeasible, the remaining choice is between entering the dilemma or eschewing it. If the former choice implies costly information arrival, for example learning the quality of the charity, it could be optimal to eschew the dilemma (see Theorem 2). A preference for having fewer options also provides an identity-based rationalization for why some entrepreneurs overinvest in certain markets before information arrives. An entrepreneur who sees herself as *bold*, for example, can rationally decide to overinvest in a market as a means to protect her identity, and then to learn the level of future demand. This prediction does not require overconfidence (Camerer and Lovallo, 1999), time-inconsistency (Brocas and Carrillo, 2004) or strategic considerations (e.g. Spence, 1977), and contradicts the classical result linking delayed information arrival to underinvestment (Arrow and Fisher, 1974). A similar argument rationalizes why males select into competitive environment more often than females, but only in patriarchal societies (Niederle and Vesterlund, 2007; Gneezy, Leonard, and List, 2009), why fathers underuse flexible work arrangements (Williams, Blair-Loy, and Berdahl, 2013) and, why different ethnic groups make specific choices concerning human capital accumulation (Austen-Smith and Fryer Jr, 2005).

I then study incentive compatible optimal disclosure to an identity-caring individual. For example, the problem faced by charity that has to decide how much information to deliver about overheads. I show that the optimal information policy may be far from full disclosure and that it varies sharply with the underlying decision problem. When the objective of the sender is to maximize the instrumental value of information for the receiver, the optimal information makes the receiver indifferent between learning and remaining ignorant. When the objective of the sender is to maximize the welfare of the receiver, the optimal information satisfies an intuitive "marginal cost equal marginal gain" condition. The latter result has an additional interpretation, it characterizes the optimal "attention strategy" of an identity-caring individual when information acquisition is flexible (as in the rational inattention models).

I provide two extensions of the baseline model. In the first (Section 7.1), I assume that the choice of an action from the menu publicly signals identity. Therefore, I relax the assumption that choices form menus are fully responsive to information. In this case, the material value of information decreases, whereas the cost of information remains the same. As a consequence, information avoidance and commitment are, *ceteris paribus*, more likely to occur with respect to the baseline model. This means that the baseline model represents a lower bound for the willingness to avoid information. If information is avoided when the actual behavior does not signal identity, it is avoided, *ceteris paribus*, when it does signal identity. In the second extension (Section 7.2), I relax the assumption that prescriptions are fully internalized. I endow the individual with a basic preference over actions that may be different from the preference generated by the identity. In this case, deviating from the prescription can be costly even in absence of information acquisition, but information exacerbates the cost of deviating. With respect to the baseline model, the extension is consistent with a preference for committing to prescriptively suboptimal actions even when committing to the prescription is feasible.

The theory also sheds new light on the relationship between information processing and preference for flexibility. The works of Arrow and Fisher (1974), Jones and Ostroy (1984) and Demers (1991) argue that a Bayesian individual anticipating the arrival of information will value flexibility. This is because the information will be used to optimize choices from the menu. A preference for larger menus also holds if the acquisition of information is Bayesian but cognitively costly, as in rational inattention models (Sims, 2003; Caplin and Dean, 2015; Pennesi, 2015; Oliveira, Denti, Mihm, and Ozbek, 2017). On the contrary, a sophisticated individual may prefer smaller menus when she anticipates a non-Bayesian reaction to information (Epstein, 2006). It seems that the value of flexibility and Bayesianism are intertwined. The present work breaks the duality and uncovers a novel mechanism leading a Bayesian individual to possibly value commitment: the desire to preserve identity.

The paper is structured as follows. Section 2 introduces the framework. Section 2.1 motivates the approach through examples. Section 3 introduces the model and Sections 4 and 5 contain its applications. Section 6 studies optimal disclosure to an identity-caring individual. Section 7 contains the two extensions. Appendix A contain additional material and Appendix B all the proofs.

2 Background and examples

In this section, I introduce the framework in which I develop the models. There is a finite set of states of the world Ω representing uncertainty. An action is a function from states to consequences $f : \Omega \to X$, where X is a compact, convex subset of a normed vector space. The set of all acts is denoted by $\mathcal{F}$. Constant acts $f(\omega) = x$ for all $\omega \in \Omega$ are identified with the element of X they deliver. In the information aversion model the individual evaluates menus, nonempty and finite subsets of $\mathcal{F}$, and I denote by $\mathcal{A}$ the set of all menus. To simplify notation, I will write f for the singletons $\{f\}$ and $f \cup g$ for the doubletons $\{f, g\}$. As customary in menu choice models, the implicit dynamic of the decision process is the following (see Fig. 1): the individual selects a menu, acquires information, forms a posterior and then selects an action from the menu before uncertainty resolves. Lastly, uncertainty resolves and the payoff materializes.

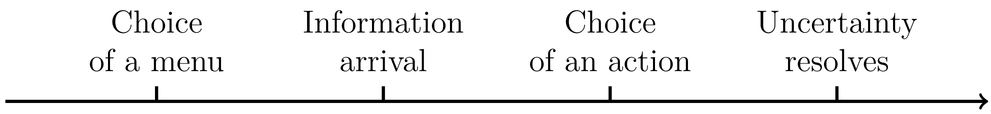

Figure 1: Dynamic of the choice process

Information. I model information as a “Bayesian experiment”, a distribution over posterior that is consistent with the prior. Formally, given a topological space M, I denote by ΔM the set of probabilities defined on the Borel σ-algebra of M. For a prior $\hat{p} \in \Delta\Omega$ an *experiment* μ is a probability $\mu \in \Delta\Delta\Omega$ satisfying the martingale property of Bayesian updating $\hat{p} = \int_{\Delta\Omega} p d\mu(p)$ (see, for example, Gollier, 2004). The experiment μ specifies the probability of deriving each posterior p by Bayesian updating the prior $\hat{p}$.[7] I denote by μ_0 the uninformative experiment, i.e. $\mu_0 = \delta_{\hat{p}}$ and by $\bar{\mu}$ the fully informative experiment, i.e. $\bar{\mu}(\delta_\omega) = \hat{p}(\omega)$ and $\bar{\mu}(p) = 0$ if $p \in \Delta\Omega$ and $p \neq \delta_\omega$. Although I assume that the only source of information is μ, allowing for flexible information acquisition is possible. Indeed, the solution to Problem P in Section 6 can be interpreted as the optimal “attention strategy” of an identity-caring individual who can acquire any information (see Remark 2).

[7] An equivalent approach that is more common in the game-theoretic literature is to assume the existence of a set of signals S correlated with the true state of the world and a function mapping signals to distributions (posteriors) over the states.

	ω_h	ω_m	ω_l
d	9	0	0
m	4	4	0
n	2	2	2

Table 1: Actions and payoffs

2.1 A motivating example: willful ignorance in charitable giving

Growing empirical evidence shows that individuals often avoid information concerning the effects of their choices on others (e.g. Dana, Weber, and Kuang, 2007; Niehaus, 2014; Grossman and van der Weele, 2016; Kandul and Ritov, 2017). Theoretical explanations based on identity or self-image preservation requires a combination of imperfect memory and self-control (Bénabou and Tirole, 2011) or explicit signaling (Grossman and van der Weele, 2016). Here, I argue that identity preservation generates information avoidance in social dilemmas as the result of a rational cost-benefit analysis. Consider three actions: a large donation d, an average donation m and no-donation n to a charity of unknown quality (either high ω_h, medium ω_m or low ω_l). The payoffs (in utils) are reported in Table 1. A large donation is optimal only if the charity is of high quality and no-donation is optimal only if the charity is of low quality. Perfect information $\mu = \bar{\mu}$ can be acquired, for example, by browsing the charity website or by meeting a fundraiser. Suppose the individual sees herself as an altruist, hence her identity prescribes to make a large donation d. Information may reveal that the charity is of low (or medium) quality, so that the optimal action given information is to not donate n (or to make an average donation m), a violation of the prescription. Information avoidance is then optimal if the material value of information is smaller than its cost. The former is the value of tailoring the optimal choice to the true state of the world (d if ω_h, m if ω_m and n if ω_l). The latter arises by the discrepancy between the information-optimal actions and the prescription. That is, choosing m rather than d if ω_m and n rather than d if ω_m. For example, Niehaus (2014) found that only 3 percent of donors acquire information prior to making a donation (see also Kandul and Ritov, 2017, for laboratory evidence).

A second and less intuitive consequence of a desire to protect identity is a potential preference for having fewer options before acquiring information. In models of costly (Sims, 2003; Oliveira et al., 2017; Pennesi, 2015) or non-costly information acquisition (Arrow and Fisher, 1974; Dillenberger, Lleras, Sadowski, and Takeoka, 2014), larger menus are always weakly preferred. Indeed, an individual expecting to receive information prefers larger menus because she will use the information to optimize her choice *from* the menu,[8] and because the cost of information is independent of the menu. Suppose the individual could commit to $d \cup m$, hence eliminating n, the possibility to not donate. In this case, acquiring information when facing $d \cup m$ can be superior to avoiding information and select an action from $d \cup m \cup n$. This is because eliminating n reduces the threats to identity deriving from information acquisition (learning that the true state is ω_l, hence suggesting n over d). Once a treat to identity is eliminated, it could optimal to learn the quality of the charity to optimize the choice between making a large or an average donation. Consistent with this conclusion, Andreoni and Serra-Garcia (2016) found that delaying charitable giving increases the number of individual agreeing to donate (from 30 to 50 percent). By pledging a future donation, the individual commits to donate, and she privately acquires information concerning the charity's quality.

3 The baseline model: fully internalized prescriptions

In this section, I formalize the intuitions provided with the previous examples and introduce the baseline model.

Identity and prescriptions. Identity determines prescriptions or norms: behaviors that are appropriate given the identity. Using the words of Akerlof and Kranton (2000):

> "We often say that people have notions–norms–of how they and others should behave. Should could imply ethical or moral views. However, we apply a more expansive meaning of should. How people should behave can refer to a social code, which can be largely internalized and even largely unconscious...The tastes derive from norms, which we define as the social rules regarding how people should behave in different situations. These rules are sometimes explicit, sometimes implicit, largely internalized, and often deeply held."

In the baseline model, I assume that prescriptions are *fully internalized* (see Akerlof and Kranton, 2000, p. 728) and they determine the individual's preference over actions. I will relax

[8] Preference for flexibility posits that if $F \subseteq G$, then G is weakly preferred to F.

such an assumption in section 7.2, where I allow for partially internalized prescriptions. In this case, the individual's ex-ante preference over actions can be different from the order induced by the prescriptions. In a situation of choice under uncertainty, fully internalized prescriptions determine the prior $\hat{p} \in \Delta\Omega$ and the utility u over payoffs. In particular, the identity prescribes the action f over g if and only if $\mathbb{E}_{\hat{p}}[u(f)] \geq \mathbb{E}_{\hat{p}}[u(g)]$. Consider the menu $d \cup n$ of the motivating example. The fact that $\hat{p}$ and u represent the prescriptions allows us interpreting $\mathbb{E}_{\hat{p}}[u(d)] = \max_{f \in d \cup n} \mathbb{E}_{\hat{p}}[u(f)]$ as the value of "following the prescription" d when facing a menu $d \cup n$. After all, when prescriptions are fully internalized and absent information acquisition, the value of a menu is given its "identity utility" (see Akerlof and Kranton, 2010, p. 18).

The value of information. What is the material value of learning μ when facing a menu F? I assume that information, if acquired, determines the optimal choice from the menu F (I will relax this assumption in Section 7.1). Therefore, the action selected *from* F is the one that maximizes $\mathbb{E}_q[u(f)]$, where q is the realized posterior. From the ex-ante point of view, the material value of μ is given by:

$$\int_{\Delta\Omega} \max_{f \in F} \mathbb{E}_p[u(f)] d\mu(p)$$

The expression is standard (see, for example, Arrow and Fisher, 1974; Gollier, 2004; Dillenberger et al., 2014) and represents the ex-ante average expected value that can be obtained from F after information acquisition. For example, consider the actions d, m, n with payoffs given in Table 1. If $u(x) = x$, $\hat{p}(\omega) = \frac{1}{3}$ for all $\omega \in \Omega$ and $\mu = \bar{\mu}$ the perfectly revealing information, the material value of $\bar{\mu}$ when facing $F = d \cup m \cup n$ is $\int_\Omega \max_{f \in d \cup m \cup n} f(\omega) d\hat{p}(\omega) = \frac{1}{3} \cdot 9 + \frac{1}{3} \cdot 4 + \frac{1}{3} \cdot 2 = 5$.

The cost of information. Motivated by the example in Section 2.1 and the fact that "violating the prescriptions evokes anxiety and discomfort in oneself" (Akerlof and Kranton, 2000), I model the cost of information as a function of the divergence between the prescribed action and the optimal action given information. Identity prescribes to donate, hence to choose $d \in d \cup m \cup n$, but for a posterior q with $\mu(q) > 0$, it is true that $\mathbb{E}_q[u(n)] > \mathbb{E}_q[u(d)]$. So the ex-post optimal action is n. Acquiring information entails the possibility that the ex-post optimal action is different from the prescription, a threat to the identity that generates a cognitive cost (in utils the cost is $\mathbb{E}_q[u(n)] - \mathbb{E}_q[u(d)] > 0$). Since different posteriors in the support of μ can rank different actions $f \in d \cup m \cup n$ better than d, the cost of information is (an increasing function of) *the average cost of identity threats* in $d \cup m \cup n$, $I(\mu, d \cup m \cup$

$n) = \int_{\Delta\Omega} \max_{f \in d\cup m\cup n} \mathbb{E}_q[u(f(\omega))d\mu(q) - \mathbb{E}_{\hat{p}}[u(d)]$. With the payoffs of Table 1, and assuming $u(x) = x$ and $\mu = \bar{\mu}$, the perfectly revealing information, $I(\bar{\mu}, d \cup m \cup n) = 5 - 3 = 2$. For a general $F \in \mathcal{A}$, the cost of information is

$$\phi(I(\mu, F)) = \phi\left(\int_{\Delta\Omega} \max_{g\in F} \mathbb{E}_q[u(g)]d\mu(q) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]\right)$$

where ϕ is an increasing and proper convex function with $\phi(0) = 0$. The term $\phi(I(\mu, F))$ is always greater than zero[9] and it is a well-defined cost of information. Indeed:

Proposition 1. *For all $F \in \mathcal{A}$, $\phi(I(\cdot, F)) : \Gamma(\hat{p}) \to [0, \infty]$ is a canonical cost function (see Definition 4), that is Normalized, (Blackwell's) Monotone and Convex.*

Normalization means that the uninformative experiment is costless ($\phi(I(\mu_0, F)) = 0$ for all $F \in \mathcal{A}$), Monotonicity means that more informative experiments (in the sense of Blackwell) are more costly and Convexity means that the marginal cost of information is increasing (see Appendix A for formal definitions). In addition, $\phi(I(\cdot, F))$ is related to well-known information costs. For example, when $X \subseteq \mathbb{R}_{++}$, $u(x) = \ln(x)$, $\hat{p}(\omega) > 0$ for all $\omega \in \Omega$ and

$$F = \left\{ c \in X^\Omega : c > 0, \sum_{\omega\in\Omega} \pi_\omega c_\omega = W \right\}, \quad \text{for some } W > 0$$

then:

$$I(\mu, F) = \int_{\Delta\Omega} H(\hat{p}) - H(p)d\mu(p) = \int_{\Delta\Omega} D(p\|\hat{p})d\mu(p)$$

where $H(q) = -\sum_{\omega\in\Omega} q(\omega)\ln q(\omega)$ is the entropy of $q \in \Delta\Omega$ and $D(p\|\hat{p}) = \sum_{\omega\in\Omega} p(\omega)\ln\left(\frac{p(\omega)}{\hat{p}(\omega)}\right)$ is the Kullback-Leibler divergence between p and $\hat{p}$ (see Fact 2 in Appendix A). This corresponds to the typical cost of information used in the rational inattention theory (Sims, 2003; Matějka and McKay, 2014). In Appendix A, I provide an alternative example in which the cost of information is a function of the average variance reduction between the prior and the posteriors.

The baseline model. The information aversion model combines the previous elements:

Definition 1. *An Information Aversion (IA) model is a tuple $(u, \hat{p}, \mu, \phi,)$ where $u : X \to \mathbb{R}$ is non-constant, $\hat{p} \in \Delta\Omega$, $\mu \in \Gamma(\hat{p})$, and the value of a menu $F \in \mathcal{A}$ is given by $V(F) =$*

[9] To see that $\phi(I(\mu, F)) \geq 0$ for all $F \in \mathcal{A}$, take any F, by the condition $\hat{p} = \int_{\Delta\Omega} q d\mu(q)$, the convexity of $q \mapsto \max_{g\in F} \mathbb{E}_q[u(g)]$ implies that $\phi(I(\mu, F)) = \phi\left(\int_{\Delta\Omega} \max_{g\in F} \mathbb{E}_q[u(g)]d\mu(q) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]\right) \geq \phi\left(\max_{g\in F} \int_{\Delta\Omega} \mathbb{E}_q[u(g)]d\mu(q) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]\right) = 0$, since $\phi(0) = 0$.

$\max\{V(F|\mu), V(F|\mu_0)\}$ *where:*

$$V(F|\mu) = \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \phi\left(I(\mu, F)\right) \tag{I}$$

with

$$I(\mu, F) = \int_{\Delta\Omega} \max_{g\in F} \mathbb{E}_q[u(g)]d\mu(q) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$$

for some weakly increasing and proper convex $\phi : [0, \infty) \to [0, \infty]$ *with* $\phi(0) = 0$*, and where:*

$$V(F|\mu_0) = \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] \tag{U}$$

In the baseline IA model, the individual values a menu F by taking the maximum between its value under ignorance $V(F|\mu_0)$ and its value when acquiring μ, $V(F|\mu)$. The latter is given by the difference between the material value of information and its cost. In $V(\cdot|\mu)$ the prescribed action acts as a menu-dependent reference and costly deviations from the reference are determined by information.[10] Notice that information is costless if the choice from the menu is trivial,[11] $\phi(I(\mu, f)) = 0$ for all $f \in \mathcal{F}$. In my interpretation, information cannot challenge identity in a singleton (I relax this property in the extension of Section 7.2, where information can be costly even for singleton menus). By the Bayesian consistency condition, it follows that $V(f|\mu) = V(f|\mu_0) = V(f)$ for all $f \in \mathcal{F}$. The next result follows from simple duality arguments (see De Lara and Gossner, 2020) and characterizes a sufficient condition for the average cost of identity threats to be null in the IA model:

Proposition 2. *The following two conditions are equivalent:*

1. *For* $F \in \mathcal{A}$, $I(\mu, F) = 0$.

2. *There exists* $f \in \text{argmax}_{g\in F} \mathbb{E}_{\hat{p}}[u(g)]$ *such that* $f \in \text{argmax}_{g\in F} \mathbb{E}_p[u(g)]$, μ*-a.s.*

Moreover, 1. or 2. implies $V(F|\mu) = V(F|\mu_0)$.

Information can be costly only when it recommends, with strictly positive probability, an *action* different from the prescription. Information that modifies the posterior beliefs but not

[10]Differently from models of reference-dependence, here the deviations from the reference are on average positive. Hence there is no room for loss-aversion, indeed the domain of ϕ is $[0, \infty)$.

[11]Indeed, $\phi(I(\mu, f)) = \phi\left(\int_{\Delta\Omega} \mathbb{E}_p[u(f)]d\mu(p) - \mathbb{E}_{\hat{p}}[u(f)]\right) = \phi\left(\mathbb{E}_{\hat{p}}[u(f)] - \mathbb{E}_{\hat{p}}[u(f)]\right) = \phi(0) = 0$, where the second equality follows from the fact that $\hat{p} = \int_{\Delta\Omega} q d\mu(q)$ and the last equality from the property of ϕ.

the recommended action is costless. Therefore, in the IA model prohibited thoughts induced by information are not necessarily costly (so they do not necessarily become *taboos*). For the altruistic individual, learning that the quality of the charity is low may not be costly. It becomes a cost when, given information, not donating becomes the optimal action, hence superior to the prescription. Conditions 1. and 2. in Proposition 2 imply that costless information is also valueless. The result of Proposition 2 allows us discriminating between the baseline IA and models with belief-dependent utilities (e.g. Köszegi, 2006). Information changing the posterior but not the recommended action can be costly when the beliefs enter the utility directly. Such distinction disappears when each posterior maps to an action that is different from the prescription. In that case, "prohibited thoughts" corresponds to "prohibited actions".

The assumption that information informs second-period choices makes the IA model an upper bound to the willingness to learn. As I will show in Section 7.1, when signaling identity requires acting in (partial) conformity with the prescription, the value of information is lower while its cost remains. Therefore, if information is not acquired when it is fully exploited in the second-period, it will not be acquired, *ceteris paribus*, when it is only partially exploited.

Remark 1. I do not explicitly model second-period choices. The assumption is that it is fully or partially informed by μ. Therefore, from the ex-ante point of view, second-period choice is random (except in the case in which it maximizes the prescription regardless of the acquired information and corresponds to $\lambda = 0$ in the model of Section 7.1). In particular, the probability of selecting $f \in F$ in the baseline IA model is given by $\mu \{p \in \Delta\Omega : \mathbb{E}_p[u(f)] \geq \mathbb{E}_p[u(g)] \ \ \forall g \in F\}$. This random choice model is studied in Lu (2016). A similar random choice rule arises in the extension presented in Section 7.1.

3.1 Parametric examples

In this section, I consider some parametrizations of the function ϕ.

Example 1 (Linear cost of information). Suppose that $\phi(x) = \kappa x$ for some $\kappa \geq 0$. Then

$$V(F|\mu) = \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - k\left(\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)].\right)$$

If $0 \leq k \leq 1$ the function $V(F|\mu)$ becomes a particular case of the Dillenberger et al.

(2014) subjective learning model. In this case, information is always (weakly) beneficial and $V(F) = V(F|\mu)$ for all $F \in \mathcal{A}$. Indeed $V(F|\mu) = (1-\kappa)\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) + k \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] \geq \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$ if and only if $I(\mu, F) \geq \kappa I(\mu, F)$, that is always true if $\kappa \leq 1$. If $\kappa > 1$, $V(F|\mu) = \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] - (\kappa-1)I(\mu, F)$ and the marginal cost of information is always greater than its marginal value. Therefore information is always avoided, $V(F) = V(F|\mu_0)$ for all $F \in \mathcal{A}$. Figure 2 (left panel) provides a graphical illustration of the two cases. Information is acquired if the cost of information (dashed lines) is under the 45° line. When $\kappa < 1$ (blue dashed line) this is always true. If $\kappa > 1$, the cost of information (dashed red line) is always higher than its value.

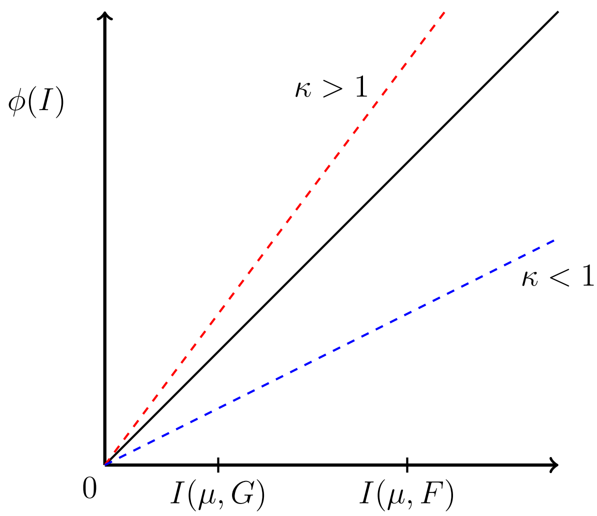

Figure 2: Linear and piece-wise linear cost of information

A less extreme case is the following:

Example 2 (Piece-wise linear cost). Suppose that, for some $I^* \geq 0$ and some $\kappa \geq 0$,

$$\phi(x) = \begin{cases} 0 & \text{if } x \leq I^* \\ \kappa(x - I^*) & \text{if } x > I^* \end{cases}$$

In this case, information is costless for small "average identity threats", see Fig. 2 (right panel) and grows linearly after a threshold. The resulting model is:

$$V(F|\mu) = \begin{cases} \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p), & \text{if } I(\mu, F) > I^* \\ \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \kappa(I(\mu, F) - I^*), & \text{if } I(\mu, F) \geq I^* \end{cases}$$

If $\kappa \leq 1$, information is always acquired, $V(F|\mu) \geq V(\mu_0|F)$ for all $F \in \mathcal{A}$. For $\kappa > 1$, information for F is avoided if[12] $I(\mu, F) \geq \frac{\kappa I^*}{\kappa - 1}$.

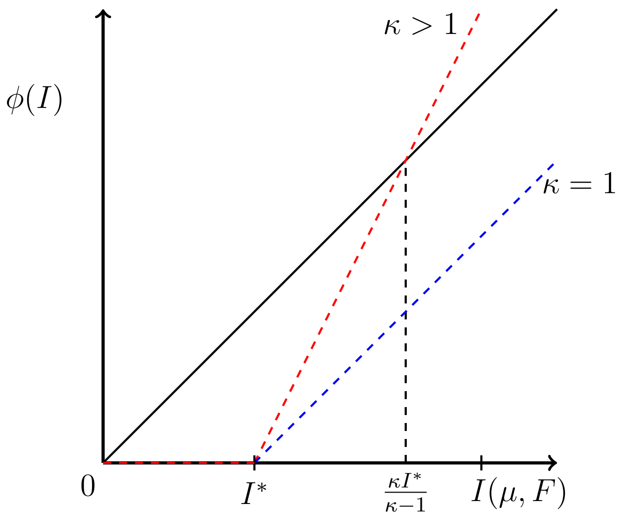

Figure 3: Piece-wise linear cost of information

Example 3 (Quadratic cost). Suppose that $\phi(x) = \frac{1}{\kappa}x^2$ for some $\kappa \geq 0$. Then $V(F|\mu_0) \geq V(F|\mu)$ if and only if $I(\mu, F) \geq \kappa$, hence information for F is acquired when it is smaller than the parameter κ. See Fig. 4.

3.2 Information and preference reversal

The first set of results studies the conditions under which preferences under ignorance differ from preferences conditional on information acquisition. I define the benefit-cost ratio between F and G:

$$\lambda_{F,G} = \frac{\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \int_{\Delta\Omega} \max_{f\in G} \mathbb{E}_p[u(f)]d\mu(p)}{I(\mu, F) - I(\mu, G)}$$

[12]Indeed, there is information avoidance for F if $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \kappa(I(\mu, F) - I^*) \geq \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$, or $I(\mu, F) \leq \kappa(I(\mu, F) - I^*)$.

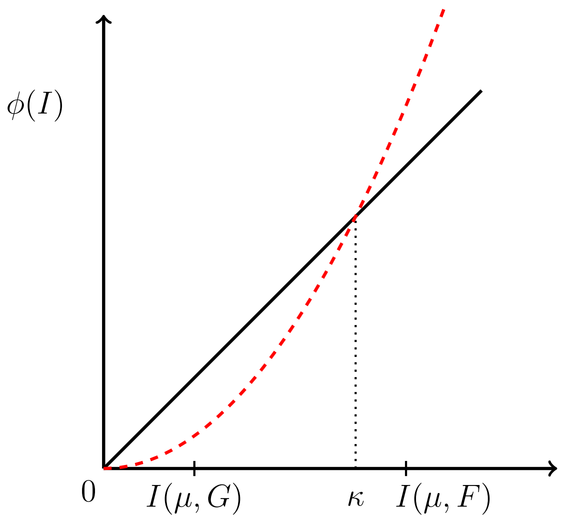

Figure 4: Quadratic cost: information is acquired for G but not for F.

if $I(\mu, F) \neq I(\mu, G)$ and $\lambda_{F,G} = 0$ otherwise. Note that $\lambda_{F,G} = 1$ if $\max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] = \max_{g\in G} \mathbb{E}_{\hat{p}}[u(g)]$. Since ϕ is a convex function, its subdifferential at $x \in \mathbb{R}_+$ is defined as:

$$\partial\phi(x) = \{\lambda \in \mathbb{R} : \phi(y) - \phi(x) \geq \lambda(y - x),\ \forall y \in \mathbb{R}_+\}$$

The next result provides two conditions that are sufficient to observe a preference reversal induced by information acquisition:

Theorem 1. *For all $F, G \in \mathcal{A}$ such that $V(F|\mu_0) \geq V(G|\mu_0)$:*

1. *If $I(\mu, F) > I(\mu, G)$ and $\lambda \geq \lambda_{F,G}$ for some $\lambda \in \partial\phi(I(\mu, G))$ then $V(G|\mu) \geq V(F|\mu)$.*

2. *If $\int_{\Delta\Omega} \max_{f\in G} \mathbb{E}_p[u(f)]d\mu(p) > \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p)$ and $\lambda \leq \lambda_{F,G}$ for some $\lambda \in \partial\phi(I(\mu, G))$, then $V(G|\mu) \geq V(F|\mu)$.*

In the first case, the average cost of identity threats in F is higher than that in G, $I(\mu, F) > I(\mu, G)$. By the initial assumption it follows that $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) > \int_{\Delta\Omega} \max_{f\in G} \mathbb{E}_p[u(f)]d\mu(p)$. Hence F is superior to G from both the prescriptive and the material viewpoint. However, if the marginal cost of information at G is high enough, the larger average cost of identity threats in F may overcome its higher material value. So, under information acquisition is optimal to choose G.

In the second case, F is prescriptively superior to G, but G has a strictly greater material value. Therefore, G has a higher average cost of identity threats than F, $I(\mu, G) > I(\mu, F)$.[13] However, if the marginal cost of information at G is "small enough", G is preferred to F when information is acquired. Clearly, the reversal due to information acquisition may not be observable from choices. Indeed, even if $V(F|\mu_0) \geq V(G|\mu_0)$ and $V(G|\mu) \geq V(F|\mu)$, it is possible that $V(F) \geq V(G)$.

Example 4. Consider the case $F = d$ and $G = m \cup n$ with the payoffs of Table 1, $\hat{p}(\omega) = \frac{1}{3}$ for all $\omega \in \Omega$, $u(x) = x$ and $\mu = \bar{\mu}$. In this case, $V(d) = \mathbb{E}_{\hat{p}}[d] = 3$ and $\mathbb{E}_{\hat{p}}[m] = V(G|\mu_0) = \frac{8}{3}$. Moreover, $\int_\Omega \max_{f \in m \cup n} f(\omega) d\hat{p}(\omega) = \frac{10}{3} > 3 = \mathbb{E}_{\hat{p}}[d]$, hence $\lambda_{d, m\cup n} = \frac{1}{2}$. Consider $\phi(x) = \frac{1}{\kappa}x^2$, by point 1. of Theorem 1, if $\frac{2}{\kappa} I(\bar{\mu}, m \cup n) = \frac{2}{\kappa}\frac{2}{3} \leq \frac{1}{2}$, namely if $\frac{1}{\kappa} \leq \frac{3}{8}$, $V(m \cup n|\bar{\mu}) \geq V(d)$.[14] Although making a large donation d is prescriptively superior to both m and n, the material value of $G = m \cup n$ is larger than that of d. On the other side $I(\bar{\mu}, d) = 0 < I(\bar{\mu}, m \cup n) = \frac{2}{3}$. Hence, if the cost of information at $I(\bar{\mu}, m \cup n)$ is small enough, it is optimal to "violate" the prescription and choose from the menu $m \cup n$ rather than committing to d (indeed $V(m \cup n|\bar{\mu}) > V(d) > V(m \cup n|\mu_0)$).

4 Application: information avoidance

In this section, I apply the IA model to rationalize various instances of *information avoidance* observed empirically. The latter is defined as the active avoidance of costless and instrumental information (see e.g. Golman et al., 2017). The main tool employed in the applications is the next result:

Proposition 3. *For all $F \in \mathcal{A}$, there is information avoidance for F, i.e. $V(F|\mu_0) > V(F|\mu)$ if and only if $I(\mu, F) < \phi(I(\mu, F))$. Moreover, if there is information avoidance for F, then $\lambda > 1$ for all $\lambda \in \partial\phi(I(\mu, F))$.*

Proposition 3, whose proof is immediate,[15] shows that information is avoided if and only if the net instrumental value of information is strictly smaller than its cost. As discussed after Proposition 2, in the baseline IA model prohibited thoughts are not necessarily taboos.

[13] Indeed, $I(\mu, G) = \int_{\Delta\Omega} \max_{f\in G} \mathbb{E}_p[u(f)]d\mu(p) - \max_{f\in G} \mathbb{E}_{\hat{p}}[u(f)] \geq \int_{\Delta\Omega} \max_{f\in G} \mathbb{E}_p[u(f)]d\mu(p) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] > \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] = I(\mu, F)$.

[14] Indeed $V(m \cup n|\bar{\mu}) = \frac{10}{3} - \frac{1}{\kappa}(\frac{10}{3} - \frac{8}{3})^2 \geq 3$ since $\frac{10}{3} - \frac{1}{\kappa}(\frac{2}{3})^2 \geq \frac{10}{3} - \frac{3}{8}\frac{4}{9} = 3.1666$.

[15] Suppose that, for some $\lambda \in \partial\phi(I(\mu, F))$, $\lambda \leq 1$. Then by the definition of subdifferential $\phi(y) - \phi(I(\mu, F)) \geq \lambda(y - I(\mu, F))$ for all $y \in \mathbb{R}_+$. Taking $y = 0$ implies $-\phi(I(\mu, F)) \geq \lambda(-I(\mu, F)) \geq -I(\mu, F)$. A contradiction to the fact that there is information avoidance for F.

They become so when they induce a violation of the prescription *and* the cost of violating the prescription is higher than its material value. For the altruistic individual, not donating becomes a taboo under two conditions: information makes it the optimal action and the cost of violating the prescription is larger than the material value of taking the optimal action. The next two results illustrate the attitude toward information avoidance or acquisition across menus. Since the function ϕ is convex, the marginal cost of information is increasing. Therefore:

Proposition 4. *If $I(\mu, G) \geq I(\mu, F)$ and there is information avoidance for F, then there is information avoidance for G.*

Given information μ, the set $\mathcal{A}$ of menus can be partitioned in two subsets: a set containing the menus for which information is acquired (i.e. $F \in \mathcal{A}$ such that $I(\mu, F) \geq \phi(I(\mu, F)))$ and a set of menus for which information is avoided (i.e. $F \in \mathcal{A}$ such that $I(\mu, F) < \phi(I(\mu, F)))$

As a corollary of the previous result:

Corollary 1. *If $V(F|\mu_0) \geq V(G|\mu_0)$ and there is information avoidance for F, then there is information avoidance for $F \cup G$.*

If F is prescriptively superior to G, then $I(\mu, F \cup G) \geq I(\mu, F)$. Hence the cost of information in $F \cup G$ is always weakly higher than that in F. Indeed there are more identity threats in $F \cup G$ than in F, but the prescriptive value of F and $F \cup G$ is the same. The fact that information is not valuable for F means that the marginal cost of information at F is higher than its marginal value. Since the marginal cost of information is increasing information cannot be valuable for $F \cup G$.

Masculinity and preventive health care. Gender is a strong predictor of life-expectancy: men die earlier than women, independently of other socio-economic factors, such as income, ethnicity, etc. Among the causes of this disparity, it is well-established that men are less likely than women to get routine medical tests and engage in preventive health care (Courtenay, 2000). This form of information avoidance appears different from alternative instances of health-related information avoidance. For example, Oster, Shoulson, and Dorsey (2013) find that individuals at-risk of developing the Huntington disease eschew a perfectly revealing genetic test. A suitable rationalization of the behavior is optimism maintenance (Brunnermeier and Parker, 2005). However, the same rationalization seems less appropriate to explain why men avoid routine medical tests. At-risk individuals in the study of Oster et al. (2013) are likely to know their condition, while it is less plausible to assume a similar awareness for the average man. An alternative explanation relies on identity protection. Getting a medical test threatens the image that *real men* are "independent, self-reliant, strong, robust and tough" (Courtenay, 2000). The IA model offers a precise description of the mechanism linking a masculine identity to the avoidance of health-related information. According to the IA model, it is not the mere act of getting tested to threaten masculinity, but the impact that health-related information will have on identity. After a diagnosis of high blood pressure, the recommendation will be to consume less red meat or take a leave from work, behaviors that threat to the (stereotypical) masculine identity.

Information avoidance and "self-control". Information avoidance is also observed in individuals experiencing self-control problems. For example, Viscusi (1990) found that smokers overestimated the risk of developing lung cancer even if the available information is virtually free and abundant. A potential explanation derives from sophistication (Carrillo and Mariotti, 2000). Information avoidance acts as a commitment device in an intra-personal game. When multiple sequential selves have conflicting preferences, the current self may prefer to remain ignorant so that future selves will not exploit public information when deciding.[16] The IA model offers an alternative rationalization of information avoidance about the effect of harmful consumption. As argued in Akerlof and Kranton (2010, pp. 19-20), smoking is a defining behavior within certain social categories, for example, men in the first decades of the last century or women after the 1970s. Indeed, smoking became an acceptable behavior for women and a symbol of empowerment and "the change in gender norms was the single most important reason for the increase in women's smoking in the United States" (Akerlof and Kranton, 2010). If smoking is a prescription, the IA model can rationalize information avoidance about its adverse effect. Suppose that are two actions, smoking s or not n and identity prescribes s. If the condition $I(s \cup n) < \phi(I(s \cup n))$ in Proposition 3 is satisfied there is information avoidance. Learning that the probability of developing lung cancer is different than expected can suggest that smoking s (the prescribed action) is inferior to not smoking n, hence threatening the identity. If the material value of learning the true probability of developing smoking-related

[16] For example, a smoker may avoid information on the probability to develop lung cancer because knowing that the latter is smaller than expected can lead her future selves to smoke more.

disease is lower than the cost of violating the prescription, the individual will prefer to remain ignorant.

Pluralistic ignorance. Individuals often privately reject a norm, but publicly accept it because they incorrectly believe that most other people accept it (Prentice and Miller, 1993). Pluralistic ignorance can arise from social identity protection: conforming to the prescriptions is a signal of belonging to the reference group. This is also a case of information avoidance, since learning the actual support of a norm in the reference group is reasonably simple. The IA model provides the following rationalization of pluralistic ignorance: learning that the prescription is less supported than thought may lead the individual to modify his behavior. For example, learning that most other husbands (in Saudi Arabia) support women working outside the home (WWOH), will make the prescription (to restrict WWOH) suboptimal, hence generating a threat to the identity (Bursztyn et al., 2020). As a consequence, information avoidance occurs at the individual level and it collectively generates pluralistic ignorance.

Work incentives. In a laboratory experiment, Huck et al. (2018) study subjects' performance in a real-effort task in which the performance pay is uncertain (either a high or a low piece-rate). They found that one-third of their subjects, given the choice, prefer to remain ignorant about the actual piece-rate. Huck et al. (2018) rationalize their findings with the optimal expectation model of Brunnermeier and Parker (2005), where ignorance helps workers to maintain motivation. The IA model provides an alternative rationalization: suppose that the subject can choose to exert either high effort h or low effort l. She sees herself as a *hard worker*, hence her identity prescribes to exert high effort. If the condition of Proposition 3 holds, the IA model is consistent with information aversion $V(h \cup l|\mu_0) > V(h \cup l|\mu)$. If information about the piece-rate recommends low effort l, the prescribed action h becomes suboptimal, hence her identity is threatened. In case the value of l is inferior to the cost of violating the identity, she may prefer to avoid information.

5 Information and commitment

In the previous section, I showed that information avoidance could be optimal for an identity-caring individual. In this section, I consider an alternative strategy that an identity-caring individual may follow to maximize her welfare. Under fairly general conditions, an identity-caring individual may prefer to commit before acquiring information. By eliminating actions that make information costly, the individual finds it optimal to acquire information and choose from a restricted menu. Moreover, if commitment to the prescriptively optimal action is unfeasible, the individual may prefer commitment to prescriptively sub-optimal actions, hence generating an information-driven form of "moral licensing".

Commitment and information acquisition. Consider the actions with payoffs of Table 1). Suppose that $\hat{p}(\omega) = \frac{1}{3}$, $u(x) = x$ and that information is perfectly revealing $\mu = \bar{\mu}$. Then, $V(d \cup m|\mu) = \frac{13}{3} - \phi(\frac{13}{3} - 3) \geq 3 = V(d \cup m \cup n|\mu_0) \geq 5 - \phi(5 - 3) = V(d \cup m \cup n|\mu)$, whenever $\frac{4}{3} \geq \phi(\frac{4}{3})$ and $\phi(2) \geq 2$. Intuitively, eliminating the option to not donate n reduces the cost of information and makes information acquisition optimal when deciding between a large d or an average m donation (if ϕ increases rapidly between $\frac{4}{3}$ and 2). Hence, the IA model is consistent with the following inequalities:

$$V(d \cup m|\mu) \geq \max\{V(d \cup m|\mu_0), V(d \cup m \cup n, \mu_0), V(d \cup m \cup n|\mu)\}$$

The next proposition provides a sufficient condition for observing the previous inequality:

Proposition 5. *If $V(F|\mu_0) \geq V(G|\mu_0)$, information is valuable for F and $\lambda \geq 1$ for some $\lambda \in \partial\phi(I(\mu, F))$, then $V(F|\mu) \geq \max\{V(F|\mu_0), V(F \cup G, \mu_0), V(F \cup G|\mu)\}$.*

When G does not add prescriptive value to F and the cost of information is steep enough at F, it is optimal to commit to F and acquire information, rather than having the flexibility of $F \cup G$. When there are only two actions, the condition for commitment to be valuable is weaker. Suppose that identity prescribes d over n, so that $V(d) = V(d \cup n|\mu_0)$. Information avoidance for $d \cup n$ means that the instrumental value of having n is lower than the cost of its threats to the identity (that prescribes d). Therefore, $V(d \cup n|\mu) < V(d \cup n|\mu_0) = V(d)$ and commitment to d is optimal.

Corollary 2. *If $V(f) \geq V(g)$, there is information avoidance for $f \cup g$ if and only if $V(f) > V(f \cup g|\mu)$.*

The results in Proposition 5 and Corollary 2 provide a unified rationalization to behaviors observed across choice domains in which identity matters (see Section 5.1). Moreover, Proposition 5 and Corollary 2 establish that commitment can be optimal when information is not physically avoidable, for example, under mandatory disclosure laws. In that case, commitment is the only way to reduce the cost of information.

Informational moral licensing. The next result concerns a preference for commitment to prescriptively suboptimal actions. For example, an altruistic individual decides to eschew a social dilemma. Suppose that an action $f \in F$ is prescriptively superior to all actions $g \in G$. However, the average cost of identity threats in $F \cup G$ is strictly greater than that in G, then it could be optimal to commit to G. In the following theorem $\lambda_{F\cup G,G}$ is the benefit-cost ratio defined in Section 3.2.

Theorem 2. *Suppose that* $V(F|\mu_0) \geq V(G|\mu_0)$ *and* $I(\mu, F \cup G) > I(\mu, G)$*, if* $\lambda \in \partial\phi(I(\mu, G))$ *is such that* $\lambda \geq \lambda_{F\cup G,G}$ *then* $V(G|\mu) \geq V(F \cup G|\mu)$*.*

An interesting particular case of Theorem 2 occurs when $G = g$. In that case, commitment to g could be strictly preferred to having the flexibility of $F \cup g$. Although $F \cup g$ has a (weakly) larger material value of g, it is possible that g becomes too costly in F and the individual prefers to commit to a prescriptively suboptimal action g. Theorem 2 can rationalize the findings of DellaVigna et al. (2012) and Andreoni et al. (2017), in which individuals intentionally avoid meeting a fundraiser, hence eschewing social dilemmas. Consider the choice faced by the individual: meeting the fundraiser implies acquiring information and then decide to make a donation or not. Although the prescription is to donate d, the cost of acquiring information may be so high that avoiding the fundraiser becomes optimal, leading the individual to commit to not donate. For example, with the payoffs of Table 1 and the assumptions, $u(x) = x$, $\mu = \bar{\mu}$, and $\hat{p}(\omega) = \frac{1}{3}$ for all $\omega \in \Omega$, $V(n) = 2 \geq \frac{13}{3} - \phi(\frac{13}{3} - 3) = V(d \cup n|\mu)$ when $\phi(\frac{4}{3}) \geq \frac{7}{3}$. Ideally, the individual would commit to d, but in case the unique choice is between $d \cup n$ and n, commitment to n may be optimal. The presence of n makes information costly because it is superior to the prescription d in states ω_m and ω_l. If such a cost is too high compared to the additional material value of $d \cup n$, the individual will commit to n. Since prescriptions are fully internalized, the individual would ideally commit to the prescriptively optimal action (i.e. d), because $V(d) > V(n) > V(d \cup n|\mu)$. Hence, commitment to a prescriptively suboptimal action should be observed only if commitment to the prescriptively optimal action is unfeasible. In section 7.2, where preference over actions can differ from the prescription, I generalize the result above allowing for $V(n) \geq V(d) \geq V(d \cup n)$.

A model that predicts a preference for commitment in the context of social dilemmas was introduced by Dillenberger and Sadowski (2012). They applied it to explain the experimental evidence collected by Dana, Cain, and Dawes (2006), where some subjects costly but quietly avoid playing the dictator in a dictator game. In Dillenberger and Sadowski (2012) a preference for commitment is motivated by the desire to avoid the shame of acting selfishly. In the present context, costly information drives commitment to prescriptively suboptimal actions.

5.1 Applications

Masculinity and the flexibility stigma. Consider the following decision problem: a worker is about to become a father and has to decide whether to ask for a flexible work arrangement (a menu F) or to maintain the current work schedule (a menu G).[17] In case of a flexible work arrangement e.g. flex-time, he will face two actions: to look after his newborn c or to delegate d (e.g. to the mother or a nanny). Hence, the menu corresponding to the flexible work arrangement is $F = c \cup d$. In case he maintains the current work schedule, he can only delegate, hence $G = d$. Uncertain is about how stressful is to look after his newborn, it can be either high ω_h or low ω_l.[18] In this case, information about the true state of the world will arrive (inexorably) after the delivery. Standard models of costly information acquisition (e.g. rational inattention) predict that the worker will *never* strictly prefer to maintain the current work schedule G prior to information arrival. Since information will be exploited to optimize the choice between c and d, flexibility is always weakly valuable, hence $F = c \cup d$ is preferred to $G = d$. However, suppose that the worker wants to preserve a (traditional) masculine and breadwinner identity. In this case, child-caring represents a threat to the identity, hence the prescribed behavior is to delegate d (see, for example, Williams et al., 2013; Rudman and Mescher, 2013; Vandello, Hettinger, Bosson, and Siddiqi, 2013). Learning that childcare is actually superior to delegating represents a threat to the identity. If the cost of information for $c \cup d$ is higher than its value, the worker will prefer to maintain his current working schedule so committing to d (by Corollary 2). Empirical evidence consistent with a dislike for flexible work arrangements in males is large (e.g. Williams et al., 2013). Among the various reasons for the

[17] To abstract from monetary considerations, I assume that the two work arrangements are economically equivalent.

[18] To avoid trivialities, I assume that child care c is superior to delegate d if the state is ω_l and d is superior to c in state ω_h.

so-called *flexibility stigma* for men, the role of masculinity protection is well-established. Asking for flexibility, especially for family caregiving, is seen as an impermissible lack of commitment, if not a feminine behavior[19] (Rudman and Mescher, 2013; Vandello et al., 2013).

Excess entry without overconfidence. Consider the following decision problem: an entrepreneur has to decide whether to enter e a new market or not n. Uncertainty concerns the level of future demand, it can be either high ω_h or low ω_l. Before deciding, she can either make an irreversible investment H or not N. By making a irreversible investment H, she commits to enter the market in the second period $H = e$. By making no initial investment N, the entrepreneur maintains flexibility $N = e \cup n$. Models of costly information acquisition predict that the entrepreneur will *never* strictly prefer to make an irreversible investment prior to acquire information about future demand. This is the classical result linking under-investment and delayed information arrival due to Arrow and Fisher (1974) (see also Gollier, 2004). However, suppose that the entrepreneur sees herself as a *bold* (for example, Brocas and Carrillo, 2004) and wants to preserve her self-image. The latter prescribes to enter the market. The IA model is consistent with $V(e) > V(e \cup n|\mu)$, if the condition of Corollary 2 is satisfied. Having n available makes information too costly with respect its material value, hence making a high initial investment $H = e$ becomes optimal. Note that this rationalization of excess entry/overinvestment *does not* require overconfidence (Camerer and Lovallo, 1999), time-inconsistency (Brocas and Carrillo, 2004) or strategic aspects à la Spence (1977). A fully rational entrepreneur can strictly prefer entering a new market because of a desire to protect her identity of being bold.

Gender competitiveness. Women are more likely then men to shy away from competition, even if there are no gender differences in performance (Niederle and Vesterlund, 2007). Interestingly, the opposite result holds in *matrilineal* societies (Gneezy et al., 2009), indicating that cultural aspects play a role. The baseline IA model provides an identity-based explanation for the gender-specific attitude to compete. Suppose that a male faces a real effort task. He can choose to exert either high effort h or low effort l. Uncertainty pertains to his ability

[19] The following anecdotal evidence comes from Williams (2006) and concerns a union arbitration of "Tractor Supply Co. (2001), in which an employer posted notice of 2 hours of mandatory overtime....[the worker] refused to stay at work past his regular shift because he had to get home to care for his grandchild. When his supervisor asked why he would not stay, he replied that it was none of his business. The supervisor said that accommodations could be made for reasonable excuses and then asked again why he could not stay. The worker again said it was none of his business. The supervisor ordered him to stay, and he was fired for insubordination"

(relative to his competitors). Information about ability is often acquired through trials sessions before the experiment starts. Entering the competitive environment (e.g. a tournament) implies competing with others, hence exerting high effort, $F = h$. Entering the non-competitive environment (e.g. a flat-rate payment scheme) leaves the flexibility to compete or not (e.g. exert either high effort or low effort), $G = h \cup l$. Suppose his self-image prescribes to compete. Therefore, Proposition 3 is consistent with $V(h) > V(h \cup l|\mu)$. If information about his relative ability recommends to shy away from competition (exert low effort l), competition h becomes suboptimal, hence identity is threatened. When the cost is higher than the instrumental value of information, entering the competitive environment can be strictly optimal. An equivalent argument holds for women willing to protect their identity in a matrilineal society. Sorting into competitive environment is often ascribed to overconfidence of men (Niederle and Vesterlund, 2007), but information processing is rational in the IA model. Although the prescription may be determined by the incorrect assumption that men outperform women, the posterior is formed by Bayes' rule. Hence, a man who rationally anticipates that competing is suboptimal can still decide to enter the competitive environment.

Identity and human capital accumulation. The potential preference for commitment predicted by the IA model can also rationalize ethnic-specific human capital investment decisions. Suppose that a student belonging to a minority (e.g. a black student) has to decide how much human capital she wants to accumulate, for simplicity, either high or low. A low accumulation of human capital corresponds to a menu G of potential future occupations, while a high level of human capital corresponds to a menu $F = G \cup w$ (where w represents a high-wage occupation e.g. becoming a lawyer). Uncertainty pertains to the features of the labor market. Suppose that the identity prescribes an occupation $b \in G$, for example, the high-wage occupation w is perceived as a "white job".[20] The result of Proposition 5 is consistent with $V(G|\mu) > V(G \cup w|\mu)$. Identity-protection implies that the cost of w, in terms of identity violations, overwhelms the value of w in terms of ex-post optimality. Hence, the student may strictly prefer commitment to a low level of human capital G (e.g. Austen-Smith and Fryer Jr, 2005).

[20] In their two audience signaling dilemma, Austen-Smith and Fryer Jr (2005) argue that "signals that induce high wages can be signals that induce peer group rejection".

6 Optimal disclosure to an identity-caring individual

In this section, I investigate the possibility to provide information to an identity-caring individual in an incentive compatible way. I consider a sender, a receiver and an experiment $\mu \in \Gamma(\hat{p})$. The experiment is called an information policy. A sender could be, for example, a regulator imposing information disclosure to sellers in a given market: in this case, the receiver is the consumer. Alternatively, the sender could be a charity that has to decide how much information to deliver about its efficiency. The sender decides the informativeness of the signal in the sense of Blackwell's (see Definition 3). A more informative experiment has a dual effect on the receiver's welfare: on one side, it has a higher material value. On the other side, it is more costly, because the cost of information $\phi(I(\cdot, F))$ is increasing with respect to the Blackwell's order (see Proposition 1).

Proposition 6. *For all $\mu, \nu \in \Gamma(\hat{p})$:*

(a). For all $F \in \mathcal{A}$: If $\nu \trianglerighteq \mu$ and if $\lambda \geq 1$ for some $\lambda \in \partial\phi(I(\mu, F))$, then $V(F|\mu) \geq V(F|\nu)$.

(b). If $\nu \trianglerighteq \mu$ and information μ is avoided for F, then ν is avoided for F.

By point (a), if the cost of information at $I(\mu, F)$ is "steep" enough, a more informative ν is welfare-reducing. By point (b), given a menu F, the set of experiments $\Gamma(\hat{p})$ can be partitioned in two subsets: a set containing the information μ that is acquired and a set containing the information ν that is avoided. Therefore, the IA model predicts that information is not always acquired (Proposition 3) and that more information can be welfare decreasing (Proposition 6).

I will consider two problems faced by the sender. First, to find the information policy with the highest instrumental value that is also incentive compatible. Second, to find the information policy that maximizes the welfare of the receiver. The first problem is the following:

$$\begin{aligned} \mu_F \in \operatorname*{argmax}_{\mu \in \Gamma(\hat{p})} & \int_{\Delta\Omega} \max_{f \in F} \mathbb{E}_p[u(f)] d\mu(p) \\ \text{s.t.} \quad & I(\mu_F, F) \geq \phi(I(\mu_F, F)) \end{aligned} \tag{Q}$$

If $\mu_F^\dagger \in \Gamma(\hat{p})$ is a solution to problem Q, then $\mu_F^\dagger \in \operatorname{argmin}\{I(\bar{\mu}, F), I(\nu_F, F)\}$ where ν_F is such that $I(\nu_F, F) = \phi(I(\nu_F, F))$. So the solution to problem Q is either the fully informative policy $\bar{\mu}$ or the policy making the individual indifferent between information acquisition and ignorance (indeed, by point 2. of Proposition 6, more informative policies are not incentive compatible).

In the second problem, the Sender has to solve the following program, for a fixed $F \in \mathcal{A}$:

$$\mu_F \in \underset{\mu \in \Gamma(\hat{p})}{\operatorname{argmax}} V(F|\mu_F) \tag{P}$$

The following result characterizes the solutions to P:

Theorem 3. *If $\mu_F^* \in \operatorname{argmin}_{\mu_1, \bar{\mu}} \{I(\mu_1, F), I(\bar{\mu}, F)\}$ where μ_1 is such that $1 \in \partial\phi(I(\mu_1, F))$ and $\bar{\mu}$ is the fully informative experiment, the information policy $\mu_F^* \in \Gamma(\hat{p})$ solves Problem P. If μ_F^* solves Problem P and $t \mapsto V(F|\mu_F^*) > V(F|\mu_0)$ then $\mu_F^* \in \operatorname{argmin}_{\mu_1, \bar{\mu}} \{I(\mu_1, F), I(\bar{\mu}, F)\}$*

The welfare maximizing solution is either the fully informative policy $\bar{\mu}$, or the policy for which the marginal value and the marginal cost of information are equal. Theorem 3 shows that full information may not be optimal for an identity-caring receiver.

Figure 5 provides a graphical illustration of the results. For the menu G, the cost function

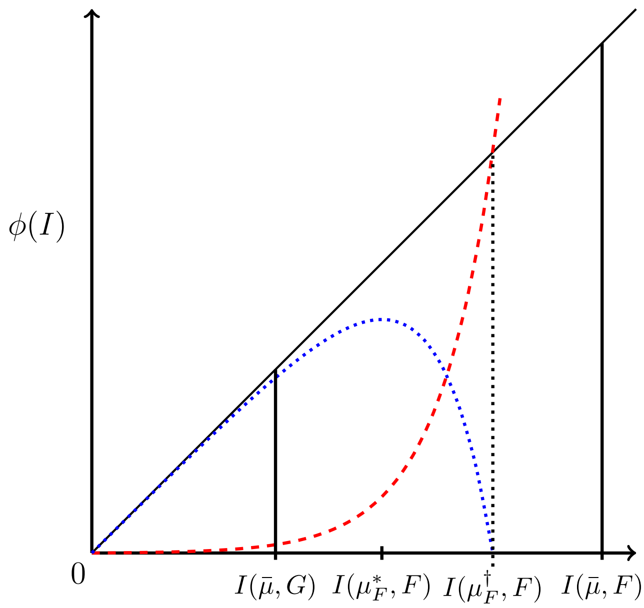

Figure 5: Optimal disclosure.

is such that $\bar{\mu} = \mu_G^* = \mu_G^\dagger$, hence full information is optimal in both problems. For the menu F, the two solutions are different. In particular, $I(\bar{\mu}, F) > I(\mu_F^\dagger, F) > I(\mu_F^*, F)$. In general:

Corollary 3. *If μ_F^* solves problem P and $\mu_F^\dagger$ solves problem Q, then $I(\bar{\mu}, F) \geq I(\mu_F^\dagger, F) \geq I(\mu_F^*, F)$.*

A related result is proved in Schweizer and Szech (2018). They consider the optimal design of payoff-relevant (medical) tests when individuals have anticipatory utility. They show that the optimal test should be perfectly informative if the true state is the preferred state of the agent and noisy otherwise. In the IA model, rather than preferred state, the agent has a preferred action, but the intuition of Schweizer and Szech (2018) remains valid. The key difference is that in the IA model there is no role for utility from anticipation, since the individual is Bayesian and an expected utility maximizer. An alternative result concerning the provision of information to behavioral individuals highlights the importance of recommendations (e.g. Lipnowski and Mathevet, 2018). Given an information policy μ, a recommendation is an information policy ν that induces the same choices (ex-post), but it is less informative than μ (in the Blackwell's sense). In the IA model, a recommendation cannot be strictly optimal because it lowers the cost of information to the same extent it lowers the material value of information.

Remark 2. The solution to Problem P has an alternative interpretation. Consider an IA model where information acquisition is flexible: the individual optimally selects her "attention strategy" $\mu \in \Gamma(\hat{p})$ as to maximize $V(F|\mu)$:

$$\max_{\mu \in \Gamma(\hat{p})} \int_{\Delta\Omega} \max_{f \in F} \mathbb{E}_p[u(f)] d\mu(p) - \phi(I(\mu, F)) \tag{1}$$

The optimal attention strategy is the solution to Problem P and it is characterized in Theorem 3. As in the case of rational inattention models, full information may be "too costly". However, differently from rational inattention theory, the cost of information depends on the menu.

7 Extensions

7.1 Choice as a signal of identity

The underlying assumption in the baseline IA model is that the second-period choice is fully responsive to information, even if the optimal ex-post choice violates the prescriptions. In this section, I relax this assumption, and I allow the ex-post choice to (partly) follow the prescription. This situation happens when actual behavior is used to (publicly) signal identity, for instance, through self-mutilation, like tattooing or circumcision (see Akerlof and Kranton, 2000). The following modification of the IA model captures the anticipated value of (partially) behave according to the prescriptions in the second-period:

Definition 2. *An General Information Aversion (GIA) model is a tuple* $(u, \hat{p}, \mu, \phi, \gamma)$ *where* $u : X \to \mathbb{R}$ *is non-constant,* $\hat{p} \in \Delta\Omega$*,* $\mu \in \Gamma(\hat{p})$*,* $\gamma \in [0, 1]$ *and the value of a menu* $F \in \mathcal{A}$ *is given by* $V^\gamma(F) = \max\{V^\gamma(F|\mu), V(F|\mu_0)\}$ *where* $V(F|\mu_0)$ *is as in Definition 1 and*

$$V^\gamma(F|\mu) = \int_{\Delta\Omega} \max_{f \in F} \mathbb{E}_p[u(f)] d\mu_\gamma(p) - \phi\left(I(\mu, F)\right) \qquad \text{(GIA)}$$

where $\mu_\gamma = \gamma\mu + (1-\gamma)\mu_0$ *and* $I(\mu, F) = \int_{\Delta\Omega} \max_{g \in F} \mathbb{E}_q[u(g)] d\mu(q) - \max_{f \in F} \mathbb{E}_{\hat{p}}[u(f)]$ *for some weakly increasing and convex* $\phi : [0, \infty) \to [0, \infty]$ *with* $\phi(0) = 0$.

With respect to the IA model, the extra parameter γ represents the responsiveness of the second period choice to information. The case of $\gamma = 1$, corresponds to the baseline IA model, where second-period choice is fully responsive to information. If $\gamma = 0$, the function $V^0(F|\mu)$ becomes $V^0(F|\mu) = \max_{f \in F} \mathbb{E}_{\hat{p}}[u(f)] - \phi\left(I(\mu, F)\right)$, and the second-period choice is completely insensitive to information. This is as if, the individual will follow the prescription regardless of the arriving information.[21] In this case, information is always weakly avoided, $V^0(F|\mu) \leq V(F|\mu_0)$. The intermediate cases $\gamma \in (0, 1)$ corresponds to a second-period choice that is partially responsive to information. The parameter $1 - \gamma$ can be interpreted as the probability that second-period choice will follow the prescriptions. The lower the γ, the lower is the material value of information. Moreover, for all $F \in \mathcal{A}$, $V^\gamma(F|\mu) \geq V^{\gamma'}(F|\mu)$ if $\gamma \geq \gamma'$, so the higher the signaling value of following the prescription, the lower the welfare of acquiring information.

Since the cost of information is the same as in the IA model Proposition 2 holds unchanged. Concerning the conditions for a preference reversal, let define

$$\lambda^\gamma_{F,G} = \frac{\gamma\left(I(\mu, F) - I(\mu, G)\right) + \max_{f \in F} \mathbb{E}_{\hat{p}}[u(f)] - \max_{f \in G} \mathbb{E}_{\hat{p}}[u(f)]}{I(\mu, F) - I(\mu, G)}$$

if $I(\mu, F) \neq I(\mu, G)$ and $\lambda^\gamma_{F,G} = 0$ otherwise. Clearly $\lambda^\gamma_{F,G} \leq \lambda_{F,G}$ for all $F, G \in \mathcal{A}$. Intuitively, with respect to $\gamma = 1$, when $\gamma \leq 1$ the cost of information remains the same but the material

[21] An earlier version of this paper provided an axiomatic characterization of this model. The axioms are similar to those of anticipated regret model of Sarver (2008) adapted to the present setting.

value of information is lower.

Theorem 4. *For all $F, G \in \mathcal{A}$ such that $V(F|\mu_0) \geq V(G|\mu_0)$:*

1. *If $I(\mu, F) > I(\mu, G)$ and $\lambda \geq \lambda^{\gamma}_{F,G}$ for some $\lambda \in \partial\phi(I(\mu, G))$ then $V(G|\mu) \geq V(F|\mu)$.*

2. *If $\int_{\Delta\Omega} \max_{f\in G} \mathbb{E}_p[u(f)]d\mu(p) > \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p)$ and $\lambda \leq \lambda^{\gamma}_{F,G}$ for some $\lambda \in \partial\phi(I(\mu, G))$, then $V(G|\mu) \geq V(F|\mu)$.*

The next result shows that, *ceteris paribus*, when actions have a signaling role information avoidance is more likely to occur:

Proposition 7. *For all $F \in \mathcal{A}$, $V(F|\mu_0) > V^{\gamma}(F|\mu)$ if and only if $\gamma I(\mu, F) < \phi(I(\mu, F))$. If there is information avoidance for F, then $\lambda > \gamma$ for all $\lambda \in \partial\phi(I(\mu, F))$. Moreover, if $V(F|\mu_0) \geq V(G|\mu_0)$ and information is not acquired for F, then it is not acquired for $F \cup G$.*

A lower marginal value of information $\gamma \leq 1$ makes it harder to satisfy the condition for optimality of information acquisition. As for the baseline model, commitment may be valuable:

Proposition 8. *For any $\gamma \in [0, 1]$:*

1. *If $V(F|\mu_0) \geq V(G|\mu_0)$ then information would not be acquired for $F \cup G$ if and only if $V(F|\mu_0) > V^{\gamma}(F \cup G|\mu)$.*

2. *If $V(F|\mu_0) \geq V(G|\mu_0)$, information is valuable for F and $\lambda \geq \gamma$ for some $\lambda \in \partial\phi(I(\mu, F))$, then $V^{\gamma}(F|\mu) \geq \max\{V(F|\mu_0), V(F \cup G, \mu_0), V^{\gamma}(F \cup G|\mu)\}$.*

The points of the previous proposition suggest that, *ceteris paribus*, commitment is more likely to occur when actions signal identity. Consider point 2., for example, the condition $\lambda \geq \gamma$ is easier to meet with respect to the corresponding condition ($\lambda \geq 1$) in Proposition 5, because $0 \leq \gamma \leq 1$. In conclusion, allowing for a signaling value of the second-period choice exacerbates both the tendency to avoid information and the willingness to commit.

7.2 Partially internalized prescriptions

In this section, I relax the assumption that the prescriptions are fully internalized. I assume that the prescriptions determine a ranking over action through a function $v : \mathcal{F} \to \mathbb{R}$: $v(f) \geq v(g)$ if and only if the identity prescribes f over g. The individual has a prior $\hat{p}$ that generates an order over actions not necessarily aligned with the prescription order represented by v. I require the following technical assumption:

Assumption 1. *For all $f \in \mathcal{F}$, $\mathbb{E}_{\hat{p}}[u(f)] \geq v(f)$.*

Assumption 1 is clearly satisfied if the prescriptions are perfectly internalized, i.e. if $v(f) = \mathbb{E}_{\hat{p}}[u(f)]$. Alternatively, it is satisfied if $v(f) = \mathbb{E}_{\tilde{p}}[u(f)]$ and $\hat{p}$ first-order stochastically dominates $\tilde{p}$. Assumption 1 is rather innocuous, as it is essentially a renormalization of the utility u.[22] The partially internalized model has the following functional form: the value of a menu $F \in \mathcal{F}$ is given by: $V^v(F) = \max\{V^v(F|\mu), V^v(F|\mu_0)\}$ where:

$$V^v(F|\mu) = \int_{\Delta\Omega} \max_{f \in F} \mathbb{E}_p[u(f)] d\mu(p) - \phi\left(I^v(\mu, F)\right) \tag{2}$$

with

$$I^v(\mu, F) = \int_{\Delta\Omega} \max_{g \in F} \mathbb{E}_q[u(g)] d\mu(q) - \max_{f \in F} v(f)$$

for some weakly increasing and convex $\phi : [0, \infty) \to [0, \infty]$ with $\phi(0) = 0$, and where:

$$V^v(F|\mu_0) = \max_{f \in F} \mathbb{E}_{\hat{p}}[u(f)] - \phi\left(\max_{g \in F} \mathbb{E}_{\hat{p}}[u(g)] - \max_{f \in F} v(f)\right) \tag{3}$$

The assumption is that, in absence of information acquisition, the choice from menus is determined by $\mathbb{E}_{\hat{p}}[u]$ subject to the cost of deviating from the prescription. The role of v is to determine a menu-dependent "reference" in terms of ideal behavior. Note that the cost of information $\phi(I^v(\mu, F))$ is always greater than zero[23] and it is increasing in the Blackwell's informativeness order. However it is not normalizes, so it is not a canonical cost of information. When the prescription is fully internalized, $v(f) = \mathbb{E}_{\hat{p}}[u(f)]$, the model corresponds to the IA model. As for the IA model, in this extension the value of an actions f is independent of information acquisition, indeed, $V^v(f|\mu) = V^v(f|\mu_0) = V^v(f)$. However, the preference over action may be different from the ranking induced by prescription. For example, $v(f) \geq v(g)$ but $V^v(f) < V^v(g)$ (see Proposition 10 for a sufficient condition). The following result provides

[22] If v is bounded and one defines $u' = u + \max_{f \in \mathcal{F}} v(f)$, then $\mathbb{E}_{\hat{p}}[u']$ is ordinally equivalent to $\mathbb{E}_{\hat{p}}[u]$ and is such that $\mathbb{E}_{\hat{p}}[u'] = \mathbb{E}_{\hat{p}}[u(f)] + \max_{f \in \mathcal{F}} v(f) \geq v(f)$ for all $f \in \mathcal{F}$ if $\mathbb{E}_{\hat{p}}[u(f)] \geq 0$.

[23] For all $F \in \mathcal{A}$ and by the condition $\hat{p} = \int_{\Delta\Omega} q d\mu(q)$, the convexity of $q \mapsto \max_{g \in F} \mathbb{E}_q[u(g)]$ implies that $\phi(I^v(\mu, F)) = \phi\left(\int_{\Delta\Omega} \max_{g \in F} \mathbb{E}_q[u(g)] d\mu(q) - \max_{f \in F} v(f)\right) \geq \phi\left(\max_{g \in F} \int_{\Delta\Omega} \mathbb{E}_q[u(g)] d\mu(q) - \max_{f \in F} v(f)\right) \geq \phi(0) = 0$, where the last inequality follows from Assumption 1.

necessary and sufficient conditions for the average threat to identity to be null in the partially internalized model:

Proposition 9. *If $I^v(\mu, F) = 0$ then there exists $f \in \operatorname{argmax}_{g\in F} \mathbb{E}_{\hat{p}}[u(g)]$ such that $f \in \operatorname{argmax}_{g\in F} \mathbb{E}_p[u(g)]$, μ-a.s. and $\max_{f\in F} v(f) = \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$. If there exists $f \in \operatorname{argmax}_{g\in F} \mathbb{E}_{\hat{p}}[u(g)]$ such that $f \in \operatorname{argmax}_{g\in F} \mathbb{E}_p[u(g)]$, μ-a.s. and $v(f) = \mathbb{E}_{\hat{p}}[u(f)]$, then $I^v(\mu, F) = 0$. Moreover, if $I^v(F, \mu) = 0$ then $V^v(F|\mu) = V^v(F|\mu_0)$.*

The average threat to identity is null if an action is optimal according to the prescription *and* for any posterior. In the following results concerning preference reversal, I will use the following benefit-cost ratio: $\lambda^v_{F,G} = \frac{\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)]d\mu(q) - \int_{\Delta\Omega} \max_{g\in G} \mathbb{E}_q[u(g)]d\mu(q)}{I^v(\mu,F)-I^v(\mu,G)}$:

Proposition 10. *For all $F, G \in \mathcal{A}$ such that $\max_{f\in F} v(f) \geq \max_{g\in G} v(g)$:*

1. *If $I^v(\mu, F) > I^v(\mu, G)$, and $\lambda \geq \lambda^v_{F,G}$ for some $\lambda \in \partial\phi(I^v(\mu, G))$ then $V^v(G|\mu) \geq V^v(F|\mu)$.*

2. *If $\int_{\Delta\Omega} \max_{g\in G} \mathbb{E}_p[u(g)]d\mu(p) > \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p)$, and $\lambda \leq \lambda^v_{F,G}$ for some $\lambda \in \partial\phi(I^v(\mu, G))$ and then $V^v(G|\mu) \geq V^v(F|\mu)$.*

The previous result is qualitatively similar to Theorem 1, preference reversal occurs depending on the slope of the function ϕ.

Before considering information avoidance in the extended model V^v, I define the baseline deviation from the prescription at F as

$$D(F) = \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] - \max_{f\in F} v(f)$$

It measures the distance (in utils) between the prescriptive behavior and the actual behavior when facing F. Notice that $D(F) = 0$ when prescriptions are fully internalized, $v(f) = \mathbb{E}_{\hat{p}}[u(f)]$. In general, Assumption 1 implies $D(f) \geq 0$.

Proposition 11. *For all $F \in \mathcal{A}$, there is information avoidance for F if $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) > \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$ and $\lambda > 1$ for some $\lambda \in \partial\phi(D(F))$. If there is information avoidance for F, then $\lambda > 1$ for all $\lambda \in \partial\phi(I^v(\mu, F))$.*

A necessary condition for information avoidance for F is $I(\mu, F) < \phi(I^v(\mu, F)) + \phi(D(F))$. Namely, if the net material value of information $I(\mu, F)$ is smaller than the total cost of information. The latter includes the cost of deviating from the prescription $\phi(D(F))$ that is, however, independent of information.

Proposition 11 can explain the use of moral wiggle rooms (e.g. Dana et al., 2007). Consider a choice between two actions f, g whose payoffs represent uncertain allocations. For simplicity, assume two states of the world $\Omega = \{\omega_1, \omega_2\}$ and two individuals. The payoff of an action f is an allocation of monetary amounts to the individuals $f(\omega) = (x, y)$, where x is the amount for the decision-maker and y for the other agent. Consider the following actions:

	ω_1	ω_2
f	$(6,5)$	$(6,1)$
g	$(5,1)$	$(5,5)$

In their experiment, Dana et al. (2007) observed that g was the modal choice when knowing that the state of the world was ω_2. Namely $u(5,5) \geq u(6,1)$ for 74% of the subjects. In their main treatment, however, the subjects did not know the state of the world but can learn it at no cost. In this case, only 56% of their subjects decide to learn the state of the world. The V^v model can explain the use of moral wiggle room. Suppose that $v(g) \geq v(f)$, $\hat{p}(\omega_1) = \frac{1}{2}$ and information is perfectly revealing, as in Dana et al. (2007), then:

Fact 1. If $u(5,5) > u(6,1)$ and $\lambda \in \partial\phi(\frac{1}{2}u(6,5) + \frac{1}{2}u(6,1) - v(g))$ is such that $\lambda > 1$, then the individual will prefer to remain ignorant, $V^v(f \cup g|\mu_0) > V^v(f \cup g|\mu)$ and will choose f from $f \cup g$.

By the condition $u(5,5) > u(6,1)$, $\int_{\Delta\Omega} \max_{h\in f\cup g} \mathbb{E}_p[u(h)]d\bar{\mu}(p) = \frac{1}{2}u(6,5) + \frac{1}{2}u(5,5) > \frac{1}{2}u(6,5) + \frac{1}{2}u(6,1) = \max_{h\in f\cup g} \mathbb{E}_{\hat{p}}[u(h)]$ and Proposition 11 applies. Even if the prescription is g, the individual will prefer to avoid information and select f from the menu $f \cup g$. A similar result holds for the baseline model with the assumption that the identity prescribes f over g.

The next result provides a sufficient condition for optimality of commitment and information acquisition:

Proposition 12. *If* $\max_{f\in F} v(f) \geq \max_{g\in G} v(g)$*, information is valuable for* F *and* $\lambda \geq 1$ *for some* $\lambda \in \partial\phi(I(\mu, F))$*, then* $V^v(F|\mu) \geq \max\{V^v(F|\mu_0), V^v(F \cup G, \mu_0), V^v(F \cup G|\mu)\}$.

As for the baseline model, commitment to a prescriptively optimal menu can be strictly preferred to having flexibility. For singleton menus the sufficient condition is simpler:

Corollary 4. *If $v(f) \geq v(g)$ and $\mathbb{E}_{\hat{p}}[u(f)] \geq \mathbb{E}_{\hat{p}}[u(g)]$ then information would not be acquired for $f \cup g$, if and only if $V^v(f) > V^v(f \cup g|\mu)$.*

If the prescription and the preference are aligned and information is not valuable for $f \cup g$, commitment to f is strictly preferred to having $f \cup g$.

Moral licensing II. In the V^v model is possible to observe a stronger form of moral licensing with respect to Theorem 2. Suppose that $\max_{f \in F} v(f) \geq v(g)$, hence there is an action in F that is at least as good as g from the prescriptive viewpoint. However, the individual strictly prefer to commit to g:

$$V^v(g) > V^v(f), \quad V^v(g) > V(f \cup g|\mu_0), \quad V^v(g) > V^v(f \cup g|\mu)$$

This is a particular case of the following proposition:

Proposition 13. *Suppose that $\max_{f \in F} v(f) \geq v(g)$ and $I^v(\mu_0, F \cup g) > I^v(\mu_0, F) > I^v(\mu_0, g)$, then there exists $\lambda^* > 0$ such that, if $\lambda \in \partial\phi(I^v(\mu_0, g))$ and $\lambda \geq \lambda^*$ then*

$$V^v(g) \geq \max\left\{V^v(F|\mu_0), V^v(F \cup g|\mu), V^v(F \cup g|\mu_0)\right\}.$$

Differently from Theorem 2, the generalized model allows for $V^v(g) > V^v(f)$ and $V^v(g) > V^v(f \cup g|\mu_0)$ which is not possible in the baseline model. As discussed after Theorem 2, with fully internalized prescriptions the individual would ideally commit to f, so that he commits to g only if commitment to f is unfeasible. When prescriptions are not fully internalized, commitment to a prescriptively suboptimal action g can be observed even when commitment to f is feasible. Such a difference, that discriminates fully from partially internalized prescriptions, can be tested with a laboratory experiment. It also suggests an additional treatment that experiments on moral licensing do not typically consider: the possibility to commit to the prescriptively optimal action.

8 Related literature

Akerlof and Kranton (2000, 2005, 2010) introduced identity consideration in economics with the objective of integrating results in social psychology with the traditional analysis of individual behavior based only on material incentives. The original work of Akerlof and Kranton (2000) describes a general model of behavior where the utility of an individual is a function of her actions, the action of others and her identity. Differently from the approach of Akerlof and Kranton, I assume (except in Section 7.2) that the value of an action only depends on its material payoffs, and it is not a function of non-standard elements. Identity is a multidimensional and dynamic concept. An individual has multiple coexisting identities depending on gender, occupation, political view, social class, etc. that may be more or less salient for a given choice.[24] What matters for the present work is that identity, whether mono or multidimensional, is internalized by the individual and generates a ranking among actions. With respect to Akerlof and Kranton (2000), I do not consider how individuals select their identities, since the focus of the paper is on the effect of identity on information acquisition. A recent empirical approach that examines identity choice is Atkin, Colson-Sihra, and Shayo (2021). The paper that is most closely related, in spirit, to the current work is Bénabou and Tirole (2011). In their work, the individual has imperfect memory and uses past actions to (self-) signal her identity. Combined with temptation in the form of quasi-hyperbolic discounting their model accommodates a variety of empirical facts including a form of information avoidance. The present work takes an alternative approach and shows that identity-protection does not require behavioral biases, such as imperfect memory and temptation. As argued before, imperfect memory of "stable identities" such as gender or political identities, appears hard to justify. Moreover, to the best of my knowledge empirical evidence of the relation between the present bias and the desire to protect one's identity is lacking. The individual in the IA model knows who she is, processes information in a Bayesian way, values actions according to their expected utility and does not suffer from temptation. These assumptions also differentiate the present work from the numerous alternative rationalizations of information aversion or avoidance (see Hertwig and Engel, 2016; Golman et al., 2017; Sunstein, 2020, for recent treatments). These

[24] Recently, Gennaioli and Tabellini (2019) provides a model of behavior blending stereotypes and social identity (in the sense of Tajfel and Turner, 1979) to explain political conflict and how the salience of certain political issues affects identification and behavior.

include cognitive constraints, motivated beliefs, anticipatory feelings, disappointment aversion and sophisticated time-inconsistency. The theory of rational inattention developed in Sims (2003) posits that cognitive limitations forces the individual to optimally acquire information subject to a cost. In this model and its generalizations (Oliveira et al., 2017; Pennesi, 2015), the cost of information depends on the informativeness of the acquired signal, hence independent of the menu faced by the individual. Therefore, flexibility is weakly valuable, an assumption that seems unable to capture identity-protective behavior (see the example in Section 2.1).

When beliefs have a hedonic value, individuals actively manage them, selecting and distorting information or memories in a self-serving manner (Bénabou, 2015; Köszegi, 2006). Therefore, information has a direct effect on welfare. First, in the IA model, posteriors are not distorted and information is processed rationally. Second, the cost of information in the IA model is null when the information does not affect optimality of the prescription. A conclusion that may not hold when beliefs enter the utility. An alternative rationalization of willful ignorance is provided in Carrillo and Mariotti (2000). They study intertemporal choice with time-inconsistent preferences, proving that information avoidance can be optimal in an intra-personal game between successive selves of the same individual. In Carrillo and Mariotti (2000), ignorance is instrumental to constraint the choices of the future selves. In the IA model, preferences are time-consistent and commitment may be instrumental to limit the negative effect of information. Moreover, the reason to avoid information is identity protection rather than time-inconsistency. Non-standard preferences are also able to rationalize information aversion: for example, Andries and Haddad (2019) assume a dynamic extension of the disappointment aversion model of Gul (1991) and Pagel (2018) the stochastic reference-dependent model of Köszegi and Rabin (2006, 2007). A related approach is Piermont (2019), who develops a model of image-conscious choice. In his work, the individual cares about the signal that a choice provides in term of possible rationalizations. Piermont (2019), however, does not consider information acquisition.

The paper contributes to the literature studying unobservable information acquisition. There are two main modeling approaches: ex-ante (choices among menus) or ex-post (choices from menus). Models focusing on ex-ante choices are Epstein (2006), Ergin and Sarver (2010), Dillenberger et al. (2014), Pennesi (2015) and Oliveira et al. (2017). Models focusing on ex-post choices are Caplin and Dean (2015), Lu (2016) and Ellis (2018). Lu (2016) studied private information when the available data are stochastic choices from menus. Oliveira et al. (2017) and Pennesi (2015) axiomatize a model of costly information acquisition, in the tradition of rational inattention of Sims (2003), where the individual decides what information to pay attention to subject to a cognitive cost. Both models assume a preference for flexibility, because the cost of information is menu-independent.

Concerning the relation between commitment and information processing, few models allow for a preference for commitment in the presence of objective uncertainty and learning. Epstein (2006) develops a model of non-Bayesian updating that leads to a desire for commitment when the individual anticipates her biased information processing. His setting considers three periods and the signal observed by the decision maker also changes the menu of options. So, the last period choice is made from signal-contingent menus. A similar setting is used in Epstein, Noor, and Sandroni (2008). In my model, the second-period choice is made from the menu selected initially. This makes the comparability with my model difficult. In any case, both models (Epstein, 2006; Epstein et al., 2008) assume Set-Betweenness, but the IA model can violate it. For example, if $V(f) \geq V(g)$ and information is strictly preferred to ignorance for $f \cup g$: $V(f \cup g|\mu) > V(f \cup g|\mu_0) = V(f)$. Most importantly, in the IA model information processing is Bayesian, and the preference for commitment derives from a desire to protect identity. In a different setting with menus of lotteries, Alauoi (2016) proved that the preservation of self-image may lead to a preference for smaller menus. In his approach, lotteries in the menu convey information about the unobservable ability of the decision-maker, and a doubt-prone individual may restrict her options as to remain "ignorant".

9 Conclusion and future research

In this paper, I propose a theory of rational information aversion based on the desire to protect one's identity. Information is costly because it may induce an identity-caring individual to act against the prescriptions. Two main results, selective information avoidance and a preference for commitment find applications in a variety of settings. Given the focus of the present approach, some aspects are left for future research: first, I do not model how the individual selects her identity (see, e.g. Atkin et al., 2021, for an empirical approach). Integrating information acquisition into a model of identity selection may increase the descriptive power of the theory, although it is unclear how information could affect the choice of an identity. Second, I do not (explicitly) consider strategic interactions. In the original model of Akerlof and Kranton (2000), the actions of others enter the utility through their material consequences and identity considerations. The model developed in the present work can be applied to strategic interactions, keeping in mind that the choice of an action depends on its expected utility. In particular, the analysis developed in Section 6, can be extended to model strategic information provision (e.g. persuasion).

A Supplementary material

A.1 Canonical cost of information

In this section, I introduce canonical costs of information (Oliveira et al., 2017) and prove Proposition 1, which establishes that $\phi(I)$ is canonical . Moreover, I prove that common costs of information such as average entropy reduction and average variance reduction are particular cases of $\phi(I)$. The next definition introduces the well-known Blackwell's informativeness order (Blackwell, 1953):

Definition 3. *Given two experiments $\mu, \nu \in \Gamma(\hat{p})$ for some $\hat{p} \in \Delta\Omega$, μ is Blackwell more informative than ν, written $\mu \trianglerighteq \nu$, if*

$$\int_{\Delta\Omega} \phi(p) d\mu(p) \geq \int_{\Delta\Omega} \phi(p) d\nu(p)$$

for all convex and continuous functions $\phi : \Delta\Omega \to \mathbb{R}$.

An experiment μ is Blackwell more informative than ν, if any experimenter (associated to a payoff function ϕ) will prefer μ to ν. For a fixed $F \in \mathcal{A}$, $\max_{f\in F} \mathbb{E}_p[u(f)]$ is convex and continuous, hence $I(\cdot, F)$ increases in the Blackwell's order. Therefore, the maximum I for a given F is given by the perfectly informative signal: $\max_{\mu\in\Gamma(\hat{p})} I(\mu, F) = I(\bar{\mu}, F) = \int_\Omega \max_{f\in F} u(f(\omega))d\hat{p}(\omega) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$.

Definition 4. *A cost function $c : \Gamma(\hat{p}) \to [0, \infty]$ is called canonical if it satisfies:*

(Normalization). *$c(\delta_{\hat{p}}) = 0$.*

(Blackwell Monotonicity). *If $\nu, \mu \in \Gamma(\hat{p})$ and $\mu \trianglerighteq \nu$ implies $c(\mu) \geq c(\nu)$.*

(Convexity). *For all experiments $\nu, \mu \in \Gamma(\hat{p})$ and $\alpha \in (0, 1)$, $c(\alpha\nu + (1-\alpha)\mu) \leq \alpha c(\nu) + (1-\alpha)c(\mu)$.*

The first condition says that the uninformative experiment has no cost. The second says that more informative experiments are more costly. Convexity is a regularity condition. Proposition 1 in Section 3 establishes that the cost of information $\phi(I(\cdot, F))$ is *canonical* for any $F \in \mathcal{A}$.

Proof. Of Proposition 1. For a given $F \in \mathcal{A}$,

$$\begin{aligned}\phi(I(\mu_0, F)) =& \phi\left[\int_{\Delta\Omega} \max_{g\in F} \mathbb{E}_p[u(g)]d\delta_{\hat{p}} - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]\right] \\ =& \phi\left[\max_{g\in F} \mathbb{E}_{\hat{p}}[u(g)] - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]\right] = \phi(0) = 0\end{aligned}$$

proving Normalization. If $\nu \trianglerighteq \mu$,

$$\begin{aligned}\phi(I(\mu, F)) =& \phi\left[\int_{\Delta\Omega} \max_{g\in F} \mathbb{E}_p[u(g)]d\mu(p) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]\right] \\ \geq& \phi\left[\int_{\Delta\Omega} \max_{g\in F} \mathbb{E}_p[u(g)]d\nu(p) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]\right] = \phi(I(\nu, F))\end{aligned}$$

proving Blackwell Monotonicity. By convexity of ϕ:

$$\begin{aligned}\phi(I(\alpha\nu+(1-\alpha)\mu,F)) =&\phi\left[\int_{\Delta\Omega}\max_{g\in F}\mathbb{E}_p[u(g)]d(\alpha\nu+(1-\alpha)\mu)(p)-\max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)]\right]\\ =&\phi\left[\alpha\int_{\Delta\Omega}\max_{g\in F}\mathbb{E}_p[u(g)]d\nu+(1-\alpha)\int_{\Delta\Omega}\max_{g\in F}\mathbb{E}_p[u(g)]d\mu(p)-\max_{f\in F}E_{\hat{p}}[u(f)]\right]\\ \leq&\alpha\phi(I(\nu,F))+(1-\alpha)\phi(I(\mu,F))\end{aligned}$$

proving Convexity. ☐

As argued in Section 3, well-known costs of information arise in the IA model for certain decision problems and specific utility functions.

Relative entropy. Suppose that the space $X\subseteq\mathbb{R}_{++}$, $\hat{p}(\omega)>0$ for all $\omega\in\Omega$ and consider the menu:

$$F=\left\{c\in X^{\Omega}:c>0,\sum_{\omega\in\Omega}\pi_\omega c_\omega=W\right\},\quad\text{for some } W>0 \tag{4}$$

These are state-contingent consumption bundles, π_ω is the "state price" and W is the total wealth. I have the following result:

Fact 2. If F in Eq. (4) and $u(x)=\ln(x)$,

$$I(\mu,F)=\int_{\Delta\Omega}H(\hat{p})-H(p)d\mu(p)=\int_{\Delta\Omega}D(p\|\hat{p})d\mu(p)$$

where $H(q)=-\sum_{\omega\in\Omega}q(\omega)\ln q(\omega)$ is the entropy of $q\in\Delta\Omega$ and $D(p\|\hat{p})=\sum_{\omega\in\Omega}p(\omega)\ln\left(\frac{p(\omega)}{\hat{p}(\omega)}\right)$ is the Kullback-Leibler divergence between p and $\hat{p}$.

Proof of Fact 2. For each $p\in\Delta\Omega$, the individual solves $\max_{c\in F}\sum_{\omega\in\Omega}p(\omega)\ln(c_\omega)-\lambda\left(\sum_{\omega\in\Omega}\pi_\omega c_\omega-W\right)$. The first order conditions are $p(\omega)\frac{1}{c_\omega}=\lambda\pi_\omega$, or $p(\omega)=\pi_\omega\lambda c_\omega$. Summing over Ω, $1=\sum_{\omega\in\Omega}p(\omega)=\lambda\sum_{\omega\in\Omega}\pi_\omega c_\omega=\lambda W$. Then, $\lambda=W^{-1}$. By the FOC, $c^*_\omega=W\frac{p(\omega)}{\pi_\omega}$. Substituting gives $\sum_{\omega\in\Omega}p(\omega)\ln\left(\frac{p(\omega)}{\pi_\omega}\right)+\ln(W)$. Then

$$\begin{aligned}I(\mu,F)=&\int_{\Delta\Omega}\sum_{\omega\in\Omega}p(\omega)\ln\left(\frac{p(\omega)}{\pi_\omega}\right)d\mu(p)+\ln(W)-\sum_{\omega\in\Omega}\hat{p}(\omega)\ln\left(\frac{\hat{p}(\omega)}{\pi_\omega}\right)-\ln(W)\\ &\int_{\Delta\Omega}\sum_{\omega\in\Omega}p(\omega)\ln(p(\omega))d\mu(p)-\int_{\Delta\Omega}\sum_{\omega\in\Omega}p(\omega)\ln(\pi_\omega)d\mu(p)-\sum_{\omega\in\Omega}\hat{p}(\omega)\ln(\hat{p}(\omega))+\sum_{\omega\in\Omega}\hat{p}(\omega)\ln(\pi_\omega)\\ =&H(\hat{p})-\int_{\Delta\Omega}H(p)d\mu(p)\end{aligned}$$

where the last equality follows from $\hat{p} = \int_{\Delta\Omega} p d\mu(p)$. Lastly,

$$
\begin{aligned}
\int_{\Delta\Omega} D(p\|\hat{p}) d\mu(p) &= \int_{\Delta\Omega} \sum_{\omega\in\Omega} p(\omega) \ln\left(\frac{p(\omega)}{\hat{p}(\omega)}\right) d\mu(p) \\
&= \int_{\Delta\Omega} \left[\sum_{\omega\in\Omega} p(\omega) \ln(p(\omega)) - \sum_{\omega\in\Omega} p(\omega) \ln(\hat{p}(\omega))\right] d\mu(p) \\
&= \int_{\Delta\Omega} \left[\sum_{\omega\in\Omega} p(\omega) \ln(p(\omega))\right] d\mu(p) - \sum_{\omega\in\Omega} \left[\int_{\Delta\Omega} p(\omega) d\mu(p) \ln(\hat{p}(\omega))\right] \\
&= \int_{\Delta\Omega} \left[\sum_{\omega\in\Omega} p(\omega) \ln(p(\omega))\right] d\mu(p) - \sum_{\omega\in\Omega} \hat{p}(\omega) \ln(\hat{p}(\omega)) \\
&= \int_{\Delta\Omega} H(\hat{p}) - H(p) d\mu(p)
\end{aligned}
$$

□

Average variance reduction. A different measure of informativeness is the average variance reduction between the prior and the posteriors (see Frankel and Kamenica, 2019). It arises in the IA model when the choice problem is to maximize a quadratic utility. Consider $X = co\,\Omega \subset \mathbb{R}$ and, given $x \in X$ define the following menu:

$$G = \{f_x \in \mathcal{F} : f_x(\omega) = x - \omega\} \tag{5}$$

The action f_x maps a state of the world ω to its distance from x. If the utility is quadratic $u(f_x(\omega)) = -(x-\omega)^2$, the optimal action is to match the state. Then, the following result holds (proof in Appendix B):

Fact 3. If G in Eq. (5) and $u(z) = -z^2$,

$$I(\mu, G) = \int_{\Delta\Omega} (Var(\hat{p}) - Var(p))\, d\mu(p)$$

where $Var(q) = \mathbb{E}_q[(\mathbb{E}_q(\omega) - \omega)^2]$.

Proof of Fact 3. The cost of information $I(\mu, G)$ can be formulated in an alternative form. For a given belief $p \in \Delta\Omega$, define the cost of uncertainty under p as $c(p, G) = \mathbb{E}_p[\max_{f\in G} u(f)] - \max_{f\in G} \mathbb{E}_p[u(f)]$ (see, for example, Frankel and Kamenica, 2019). The first term is the expected value of G if the individual knows the true state of the world before selecting an action. The second term is the expected value of the optimal action according to p. So $c(p, G)$ is the cost of not knowing the true state of the world when the belief is p and the choice set is G. It is immediate to see[25] that:

$$I(\mu, G) = \int_{\Delta\Omega} c(\hat{p}, G) - c(p, G) d\mu(p) \tag{6}$$

[25] $\int_{\Delta\Omega} c(\hat{p}, G) - c(p, G) d\mu(p) = \int_{\Delta\Omega} \mathbb{E}_{\hat{p}}[\max_{f_x\in G} u(f_x)] - \max_{f_x\in G} \mathbb{E}_{\hat{p}}[u(f_x)] - \mathbb{E}_p[\max_{f_x\in G} u(f_x)] + \max_{f_x\in G} \mathbb{E}_p[u(f_x)]$. Since $\mu \in \Gamma(\hat{p})$, $\int_{\Delta\Omega} \mathbb{E}_p[\max_{f_x\in G} u(f_x)] d\mu(p) = \mathbb{E}_{\hat{p}}[\max_{f_x\in G} u(f_x)]$, so that $\int_{\Delta\Omega} c(\hat{p}, G) - c(p, G) d\mu(p) = \int_{\Delta\Omega} \max_{f_x\in G} \mathbb{E}_p[u(f_x)] - \max_{f_x\in G} \mathbb{E}_{\hat{p}}[u(f_x)] = I(\mu, G)$.

Then,

$$I(\mu,G)=\int_{\Delta\Omega}\mathbb{E}_{\hat{p}}[\max_{f_x\in G}-(x-\omega)^2]-\max_{f_x\in G}\mathbb{E}_{\hat{p}}[-(x-\omega)^2]-\mathbb{E}_p[\max_{f_x\in G}-(x-\omega)^2]+\max_{f_x\in G}\mathbb{E}_p[-(x-\omega)^2]d\mu(p)$$

but $\mathbb{E}_p[\max_{f_x\in G}-(x-\omega)^2]=0$ for all $p\in\Delta\Omega$, hence

$$I(\mu,G)=\int_{\Delta\Omega}-\max_{f_x\in G}\mathbb{E}_{\hat{p}}[-(x-\omega)^2]+\max_{f_x\in G}\mathbb{E}_p[-(x-\omega)^2]d\mu(p)$$

since for any $p\in\Delta\Omega$, $\max_{f_x\in G}\mathbb{E}_p[-(x-\omega)^2]$ is attained at $x=\mathbb{E}_p[\omega]$, it follows that

$$I(\mu,G)=\int_{\Delta\Omega}\mathbb{E}_{\hat{p}}[(\mathbb{E}_{\hat{p}}[\omega]-\omega)^2]-\mathbb{E}_p[(\mathbb{E}_p[\omega]-\omega)^2]d\mu(p)$$

□

A.2 Falsifiability

In this section, I relate testable restrictions of the IA model to behavioral axioms introduced in the literature. I consider a binary relation $\succcurlyeq$ defined on $\mathcal{A}$ represented by the baseline IA model: $F\succcurlyeq G$ if and only if $V(F)\geq V(G)$. To simplify the notation, I will occasionally use the following notation $\sigma_F(p)\triangleq\max_{f\in F}\mathbb{E}_p[u(f)]$. Before introducing the axioms, I define the mixture + of two acts, denoted by $\alpha f+(1-\alpha)g$ for all $\alpha\in[0,1]$, as, $(\alpha f+(1-\alpha)g)(\omega)=\alpha f(\omega)+(1-\alpha)g(\omega)$. The mixture of two menus is the menu containing all mixtures of the elements in F and G, $\alpha F+(1-\alpha)G=\{\alpha f+(1-\alpha)g: f\in F,\ g\in G\}$, for all $\alpha\in[0,1]$. Clearly, $\succcurlyeq$ is a weak order and restricted to singletons it satisfies the Anscombe-Aumann axioms of expected utility if and only if u is affine with respect to vector structure of X.

The first axiom is a weakening of the independence axiom: preferences are not reversed after mixing with a singleton menu. It has been introduced by Oliveira et al. (2017) and it is in the same spirit of the weak certainty independence axiom of Maccheroni, Marinacci, and Rustichini (2006). A similar axiom is present in the costly contemplation model of Ergin and Sarver (2010).

Axiom (Weak Singleton Independence)**.** *For all menus $F,G\in\mathcal{A}$, $h,h'\in\mathcal{F}$ and $\alpha\in(0,1)$,*

$$\alpha F+(1-\alpha)h\succcurlyeq\alpha G+(1-\alpha)h\implies\alpha F+(1-\alpha)h'\succcurlyeq\alpha G+(1-\alpha)h'$$

The WSI axiom is related to a "translation invariance" property of the representation V.

Proposition 14. *If $\succcurlyeq$ has an IA representation, then $\succcurlyeq$ satisfies the WSI axiom.*

Proof of Proposition 14. Consider $\int_{\Delta\Omega}\max_{f\in\alpha F+(1-\alpha)h}\mathbb{E}_p[u(f)]d\mu$, by the properties of the support function, it is equal to $\int_{\Delta\Omega}\max_{f\in\alpha F}\mathbb{E}_p[u(f)]+(1-\alpha)\mathbb{E}_p[u(h)]d\mu$. By the Bayesian consistency condition, $\int_{\Delta\Omega}\max_{f\in\alpha F}\mathbb{E}_p[u(f)]+(1-\alpha)\mathbb{E}_{\hat{p}}[u(h)]$. This implies $I(\mu,\alpha F+(1-\alpha)h)=I(\mu,\alpha F)$. Hence, $V(\alpha F+(1-\alpha)h|\mu)=\int_{\Delta\Omega}\max_{f\in\alpha F+(1-\alpha)h}\mathbb{E}_p[u(f)]d\mu-\phi\left(I(\mu,\alpha F+(1-\alpha)h)\right)=$

$V(\alpha F|\mu) + (1-\alpha)\mathbb{E}_{\hat{p}}[u(h)]$. Suppose that $\alpha F + (1-\alpha)h \succcurlyeq \alpha G + (1-\alpha)h$, then

$$V(\alpha F + (1-\alpha)h) = \max\left\{V(\alpha F|\mu) + (1-\alpha)\mathbb{E}_{\hat{p}}[u(h)], V(\alpha F|\mu_0) + (1-\alpha)\mathbb{E}_{\hat{p}}[u(h)]\right\}$$

hence

$$V(\alpha F + (1-\alpha)h) = V(\alpha F) + (1-\alpha)\mathbb{E}_{\hat{p}}[u(h)]$$

If $V(\alpha F + (1-\alpha)h) \geq V(\alpha G + (1-\alpha)h)$, then $V(\alpha F) \geq V(\alpha G)$, adding on both sides $(1-\alpha)\mathbb{E}_{\hat{p}}[u(h')]$ gives the result. □

The next axiom bundles the Anscombe-Aumann monotonicity axiom with a weaker form of strategic rationality. If a menu contains an action f that dominates *in any state* an action g, then adding g to the menu leaves the individual indifferent. The condition is assumed in Dillenberger et al. (2014); Pennesi (2015); Lu (2016); Oliveira et al. (2017).

Axiom (Dominance)**.** *If $f(\omega) \succcurlyeq g(\omega)$ for all $\omega \in \Omega$ and $f \in F$, then $F \sim F \cup g$.*

First, the ranking dictated by the prescriptions satisfies the Anscombe-Aumann monotonicity axiom. Second, regardless of the information one expects to receive, the act f is unambiguously better than g. Hence, adding g to a menu containing f cannot increase or decrease the value of the menu. This is because no posterior will recommend choosing g in the presence of f, implying that information has the same value (and cost) in both $F \cup g$ and in F.

Proposition 15. *If $\succcurlyeq$ has an IA representation, then $\succcurlyeq$ satisfies the Dominance axiom.*

Proof of Proposition 15. If $f(\omega) \succcurlyeq g(\omega)$, then clearly $\mathbb{E}_p[u(f)] \geq \mathbb{E}_p[u(g)]$ for all $p \in \Delta\Omega$. Moreover, for all $p \in \Delta\Omega$, $\max_{f \in F\cup g} \mathbb{E}_p[u(f)] = \max_{f\in F} \mathbb{E}_p[u(f)]$. Hence, $V(F \cup g|\mu) = \int_{\Delta\Omega} \sigma_{F\cup g}(p)d\mu(p) - \phi\left(\int_{\Delta\Omega} \sigma_{F\cup g}(p)d\mu(p) - \sigma_{F\cup g}(\hat{p})\right) = \int_{\Delta\Omega} \sigma_F(p)d\mu(p) - \phi\left(\int_{\Delta\Omega} \sigma_F(p)d\mu(p) - \sigma_F(\hat{p})\right) = V(F|\mu)$. Then, $V(F \cup g) = \max\{V(F\cup g|\mu_0), V(F\cup g|\mu)\} = \max\{V(F|\mu_0), V(F|\mu)\} = V(F)$, hence the result. □

B Proofs

As done in the previous section, I will occasionally use the following notation $\sigma_F(p) \triangleq \max_{f\in F} \mathbb{E}_p[u(f)]$.

Proof of Proposition 2. The equivalence of 1. and 2. follows form De Lara and Gossner (2020) and the fact that $\phi(0) = 0$. To prove that 1. implies $V(F|\mu) = V(F|\mu_0)$: if, for some $F \in \mathcal{A}$, $I(\mu, F) = 0$, then $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) = \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$, hence $V(F|\mu) = \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \phi(I(\mu, F)) = \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \phi(0) = \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$. □

Proof of Theorem 1. The result is a particular case of Theorem 4 with $\gamma = 1$. □

Proof of Proposition 4. Since $\phi(0) = 0$ and ϕ is convex, the ratio $\frac{\phi(x)}{x}$ is weakly increasing in x, for all $x > 0$. By the condition $I(\mu, G) \geq I(\mu, F)$, if there is information avoidance for F, $I(F|\mu) < \phi(I(\mu, F))$ and, by the properties of ϕ, $1 < \frac{\phi(I(\mu,F))}{I(\mu,F)} \leq \frac{\phi(I(\mu,G))}{I(\mu,G)}$, so there is information avoidance for G as well.

□

Proof of Proposition 5. Since information is valuable for F, $V(F|\mu) \geq V(F|\mu_0) = V(F \cup G|\mu_0)$, where the last equality follows from $V(F|\mu_0) \geq V(G|\mu_0)$. If $\lambda \geq 1$ for some $\lambda \in \partial\phi(I(\mu, F))$, then

$$\begin{aligned}
&\phi(I(\mu, F \cup G)) - \phi(I(\mu, F)) \geq \lambda(I(\mu, F \cup G) - I(\mu, F)) \\
=&\lambda \left(\int_{\Delta\Omega} \max_{f \in F \cup G} \mathbb{E}_p[u(f)] d\mu(p) - \max_{f \in F \cup G} \mathbb{E}_{\hat{p}}[u(f)] - \int_{\Delta\Omega} \max_{f \in F} \mathbb{E}_p[u(f)] d\mu(p) + \max_{f \in F} \mathbb{E}_{\hat{p}}[u(f)] \right) \\
=&\lambda \left(\int_{\Delta\Omega} \max_{f \in F \cup G} \mathbb{E}_p[u(f)] d\mu(p) - \int_{\Delta\Omega} \max_{f \in F} \mathbb{E}_p[u(f)] d\mu(p) \right) \\
\geq& \int_{\Delta\Omega} \max_{f \in F \cup G} \mathbb{E}_p[u(f)] d\mu(p) - \int_{\Delta\Omega} \max_{f \in F} \mathbb{E}_p[u(f)] d\mu(p)
\end{aligned}$$

where the first inequality follows from the definition of subgradient, the first equality by definition, the second equality follows from $V(F|\mu_0) \geq V(G|\mu_0)$ and the last inequality from the fact that $\lambda \geq 1$ and the term in the parenthesis is positive. Rearranging gives $V(F|\mu) \geq V(F \cup G|\mu)$.

□

Proof of Corollary 2. If information is not acquired a $f \cup g$, then $V(f \cup g|\mu) < V(f \cup g|\mu_0)$, meaning $\int_{\Delta\Omega} \max_{h \in f \cup g} \mathbb{E}_p[u(h)] d\mu(p) - \phi(I(\mu, f \cup g)) < \max_{h \in f \cup g} \mathbb{E}_{\hat{p}}[u(h)] = \mathbb{E}_{\hat{p}}[u(f)]$, where the last equality follows from the assumption that $\mathbb{E}_{\hat{p}}[u(f)] \geq \mathbb{E}_{\hat{p}}[u(g)]$. For the opposite implication, $\int_{\Delta\Omega} \max_{h \in f \cup g} \mathbb{E}_p[u(h)] d\mu(p) - \phi(I(\mu, f \cup g)) < \mathbb{E}_{\hat{p}}[u(f)] = \max_{h \in f \cup g} \mathbb{E}_{\hat{p}}[u(h)]$, hence, information is not acquired at $f \cup g$. □

Proof of Theorem 2. Consider $\lambda_{F \cup G, G} = \frac{\int_{\Delta\Omega} \sigma_{F \cup G}(p) d\mu(p) - \int_{\Delta\Omega} \sigma_G(p) d\mu(p)}{I(\mu, F \cup G) - I(\mu, G)}$. The same argument used in the proof of Theorem 4 shows that $\lambda_{F \cup G, G} > 0$. If $\lambda' \in \partial\phi(I(\mu, G))$ and $\lambda' \geq \lambda_{F \cup G, G}$ then

$$\begin{aligned}
\phi(I(\mu, F \cup G)) - \phi(I(\mu, G)) \geq& \lambda' \left(I(\mu, F \cup G) - I(\mu, G) \right) \\
\geq\geq& \lambda_{F \cup G, G} \left(I(\mu, F \cup G) - I(\mu, G) \right) \\
=& \int_{\Delta\Omega} \max_{f \in F \cup G} \mathbb{E}_p[u(f)] d\mu(p) - \int_{\Delta\Omega} \max_{f \in G} \mathbb{E}_p[u(f)] d\mu(p)
\end{aligned}$$

where the first inequality follows from the definition of subgradient, the second from the condition and the equality by the definition of $\lambda_{F \cup G, G}$. Rearranging gives $V(G|\mu) \geq V(F \cup G|\mu)$.

□

Proof of Theorem 3. Consider an arbitrary $\mu \in \Gamma(\hat{p})$, and a $\mu_1 \in \Gamma(\hat{p})$ such that $1 \in \phi(I(\mu_1, F))$.

Then:

$$\begin{aligned}V(\mu_1,F)-V(\mu,F) &= \int_{\Delta\Omega}\sigma_F(p)d\mu_1(p)-\phi(I(\mu,F_1))-\int_{\Delta\Omega}\sigma_F(p)d\mu(p)+\phi(I(\mu,F))\\ &= \int_{\Delta\Omega}\sigma_F(p)d\mu_1(p)-\int_{\Delta\Omega}\sigma_F(p)d\mu(p)+\phi(I(\mu,F))-\phi(I(\mu_1,F))\\ &\geq \int_{\Delta\Omega}\sigma_F(p)d\mu_1(p)-\int_{\Delta\Omega}\sigma_F(p)d\mu(p)+I(\mu,F)-I(\mu_1,F)=0\end{aligned}$$

Hence, μ_1 is optimal, otherwise the solution is $\bar{\mu}$. For the existence of μ_1, notice that:

Fact 4. For each $t\in[0,I(\bar{\mu},F)]$, there is $\mu_t\in\Gamma(\hat{p})$ such that $I(\mu_t,F)=t$.

Proof of Fact 4. Let $\mu_t=\frac{t}{I(F,\bar{\mu})}\bar{\mu}+(1-\frac{t}{I(F,\bar{\mu})})\mu_0$. Then $I(\mu_t,F)=\int_{\Delta\Omega}\sigma_F(p)d\mu_t(p)-\sigma_F(\hat{p})=\frac{t}{I(F,\bar{\mu})}(I(F,\bar{\mu})+\sigma_F(\hat{p}))+(1-\frac{t}{I(F,\bar{\mu})})\sigma_F(\hat{p})-\sigma_F(\hat{p})=t$. □

Hence, if $1\in\partial\phi(t)$, μ_t defined above is μ_1.

If μ_1 maximizes $V(F|\cdot)$ and $V(F|\mu_1)>V(F|\mu_0)$, μ_1 also maximizes $V(F|\cdot)-V(F|\mu_0)$. Since $V(F|\cdot)-V(F|\mu_0)$ is increasing in the Blackwell's order, if the maximum is not attained for some μ^* such that $I(\mu^*,F)<I(\bar{\mu},F)$, then the optimal policy is $\bar{\mu}$. If the maximum is not attained for some μ_t^* such that $I(\mu_t^*,F)<I(\bar{\mu},F)$, then it must be the case that $0\in\partial(V(F|\mu_t^*)-V(F|\mu_0))$ seen as a function $t\mapsto V(F|\mu_t)-V(F|\mu_0)$. By the chain rule of differentiation (Clarke, 1983) $0=\left(\int_{\Delta\Omega}\sigma_F(p)d\mu_t^*(p)-\sigma_F(\hat{p})\right)(1-\lambda)$ for some $\lambda\in\partial\phi(I(\mu_t^*,F))$, that implies $\lambda=1$

□

Proof of Theorem 4. For point 1., by the second assumption, F and G cannot be singletons. Consider $\lambda^{\gamma}_{F,G}=\frac{\gamma\int_{\Delta\Omega}\sigma_F(p)d\mu(p)+(1-\gamma)\sigma_F(\hat{p})-\gamma\int_{\Delta\Omega}\sigma_G(p)d\mu(p)-(1-\gamma)\sigma_G(\hat{p})}{I(\mu,F)-I(\mu,G)}$. If $\gamma=0$, $\lambda^0_{F,G}=\frac{\sigma_F(\hat{p})-\sigma_G(\hat{p})}{I(\mu,F)-I(\mu,G)}$, then either $\lambda^0_{F,G}=0$ if $\sigma_F(\hat{p})=\sigma_G(\hat{p})$ or $\lambda^0_{F,G}>0$ if $\sigma_F(\hat{p})>\sigma_G(\hat{p})$. In that both cases the result follows from $\phi(I(\mu,F)-\phi(I(\mu,G))\geq\lambda(I(\mu,F)-I(\mu,G))\geq\lambda^0_{F,G}(I(\mu,F)-I(\mu,G)))=\sigma_F(\hat{p})-\sigma_G(\hat{p})$. Rearranging gives the result. For $\gamma>0$, notice that $\gamma I(\mu,F)=\gamma\int_{\Delta\Omega}\sigma_F(p)d\mu(p)-\gamma\sigma_F(\hat{p})>\gamma\int_{\Delta\Omega}\sigma_G(p)d\mu(p)-\gamma\sigma_G(\hat{p})=\gamma I(\mu,G)$. Adding $\sigma_F(\hat{p})$ to both sides gives $\gamma\int_{\Delta\Omega}\sigma_F(p)d\mu(p)+(1-\gamma)\sigma_F(\hat{p})>\gamma\int_{\Delta\Omega}\sigma_G(p)d\mu(p)-\gamma\sigma_G(\hat{p})+\sigma_F(\hat{p})\geq\gamma\int_{\Delta\Omega}\sigma_G(p)d\mu(p)-\gamma\sigma_G(\hat{p})+\sigma_G(\hat{p})$. Hence $\lambda^{\gamma}_{F,G}>0$. If $\lambda\geq\lambda^{\gamma}_{F,G}$ for some $\lambda\in\partial\phi(I(\mu,G))$, then

$$\begin{aligned}\phi(I(\mu,F))-\phi(I(\mu,G)) \geq&\lambda\left(I(\mu,F)-I(\mu,G)\right)\\ \geq&\lambda^{\gamma}_{F,G}\left(I(\mu,F)-I(\mu,G)\right)\\ =&\gamma\int_{\Delta\Omega}\sigma_F(p)d\mu(p)-\gamma\int_{\Delta\Omega}\sigma_G(p)d\mu(p)+(1-\gamma)\sigma_F(\hat{p})-(1-\gamma)\sigma_G(p)\end{aligned}$$

where the first inequality follows from the definition of subdifferential, the second inequality from the fact that $\lambda\geq\lambda^{\gamma}_{F,G}$ and the strict positivity of the term in the parenthesis and the last equality by definition of $\lambda^{\gamma}_{F,G}$. Rearranging gives $V^{\gamma}(G|\mu)\geq V^{\gamma}(F|\mu)$.

For point 2., by the two initial assumptions F and G cannot both singletons. First notice that $I(\mu,F)=\int_{\Delta\Omega}\sigma_F(p)d\mu(p)-\sigma_F(\hat{p})<\int_{\Delta\Omega}\sigma_G(p)d\mu(p)-\sigma_G(\hat{p})=I(\mu,G)$. Hence $\lambda^{\gamma}_{F,G}>0$

If $\lambda \leq \lambda_{F,G}$ for some $\lambda \in \partial\phi(I(\mu, G))$, then

$$
\begin{aligned}
\phi(I(\mu,F)) - \phi(I(\mu,G)) \geq & \lambda\left(I(\mu,F) - I(\mu,G)\right) \\
\geq & \lambda^{\gamma}_{F,G}\left(I(\mu,F) - I(\mu,G)\right) \\
= & \gamma \int_{\Delta\Omega} \sigma_F(p) d\mu(p) - \gamma \int_{\Delta\Omega} \sigma_G(p) d\mu(p) + (1-\gamma)\sigma_F(\hat{p}) - (1-\gamma)\sigma_G(\hat{p})
\end{aligned}
$$

where the first inequality follows from the definition of subdifferential, the second from the condition $I(\mu, F) < I(\mu, G)$ and the fact that $\lambda \leq \lambda^{\gamma}_{F,G}$ and the strict negativity of the term in parenthesis. The equality follows from the definition of $\lambda^{\gamma}_{F,G}$. Rearranging gives $V^{\gamma}(G|\mu) \geq V^{\gamma}(F|\mu)$. □

Proof of Proposition 6. For a., since $\nu \trianglerighteq \mu$ either $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)] d\nu(p) = \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)] d\mu(p)$ or $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)] d\nu(p) > \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)] d\mu(p)$. In the first case, $V(F|\nu) = V(F|\mu)$. Consider the strict inequality, if $\lambda \in \partial\phi(I(\mu, F))$ and $\lambda \geq 1$, then

$$
\begin{aligned}
\phi(I(\nu,F)) - \phi(I(\mu,F)) \geq & \lambda\left(I(\nu,F) - I(\mu,F)\right) \\
\geq & \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)] d\nu(p) - \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)] d\mu(p)
\end{aligned}
$$

where the first inequality follows from the definition of subgradient and the second from the definition of $I(\mu, F)$ and the fact that $\lambda \geq 1$. Rearranging gives $V(F|\nu) \geq V(F|\mu)$.

For b., since $\phi(0) = 0$ and ϕ is convex, the ratio $\frac{\phi(x)}{x}$ is weakly increasing in x, for all $x > 0$. By the condition $\nu \trianglerighteq \mu$, it follows that $V(F|\nu) \geq V(F|\mu)$, hence $I(\nu, F) \geq I(\mu, F)$. If information μ is not acquired for F, $I(F|\mu) < \phi(I(\mu, F))$ and, by the properties of ϕ, $1 < \frac{\phi(I(\mu,F))}{I(\mu,F)} \leq \frac{\phi(I(\nu,F))}{I(\nu,F)}$, so information ν is not acquired for F as well. □

Proof of Proposition 7. The first part is trivial. For the second, since $\phi(0) = 0$ and ϕ is convex, the ratio $\frac{\phi(x)}{x}$ is weakly increasing in x, for all $x > 0$. By the condition $V(F|\mu_0) \geq V(G|\mu_0)$, it follows that $V(F \cup G|\mu_0) = V(F|\mu_0)$, hence $I(\mu, F \cup G) \geq I(\mu, F)$. By the first part, if information is not acquired for F, $\gamma I(F|\mu) < \phi(I(\mu, F))$ and, by the properties of ϕ, $\gamma < \frac{\phi(I(\mu,F))}{I(\mu,F)} \leq \frac{\phi(I(\mu,F\cup G))}{I(\mu,F\cup G)}$, so information is not acquired for $F \cup G$ as well. Suppose that there is information avoidance for F and $\lambda \leq \gamma$ for some $\lambda \in \partial\phi(I(F, \mu))$. Then, $\phi(y) - \phi(I(\mu, F)) \geq \lambda(y - I(\mu, F))$ for all $y \in \mathbb{R}_+$. Taking $y = 0$ implies $-\phi(I(\mu, F)) \geq \lambda(-I(\mu, F)) \geq -\gamma I(\mu, F)$. A contradiction to the fact that there is information avoidance for F □

Proof of Proposition 8. Point 1. if information is not acquired a $F \cup G$, then $V^{\gamma}(F \cup G|\mu) < V(F \cup G|\mu_0)$, meaning $\int_{\Delta\Omega} \max_{f\in F\cup G} \mathbb{E}_p[u(f)] d\mu_{\gamma}(p) - \phi(I(\mu, F \cup G)) < \max_{f\in F\cup G} \mathbb{E}_{\hat{p}}[u(h)] = \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] = V(F|\mu_0)$, where the last equality follows from the assumption $V(F|\mu_0) \geq V(G|\mu_0)$. For the opposite implication, $V_{\gamma}(F\cup G|\mu) = \int_{\Delta\Omega} \max_{f\in F\cup G} \mathbb{E}_p[u(f)] d\mu_{\gamma}(p) - \phi(I(\mu, F \cup G)) < \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] = \max_{f\in F\cup G} \mathbb{E}_{\hat{p}}[u(f)]$, hence, information is not acquired at $F \cup G$. For point 2., since information is valuable for F, $V^{\gamma}(F|\mu) \geq V(F|\mu_0) = V(F \cup G|\mu_0)$, where the last equality follows from $V(F|\mu_0) \geq V(G|\mu_0)$. Notice that $\lambda \in \partial\phi(x)$ cannot be negative because ϕ is weakly increasing. If $\lambda \geq \gamma$ for some $\lambda \in \partial\phi(I(\mu, F))$, then

$$
\begin{aligned}
&\phi(I(\mu, F \cup G)) - \phi(I(\mu, F)) \geq \lambda(I(\mu, F \cup G) - I(\mu, F)) \\
=&\lambda\left(\int_{\Delta\Omega} \max_{f\in F\cup G} \mathbb{E}_p[u(f)]d\mu(p) - \max_{f\in F\cup G} \mathbb{E}_{\hat{p}}[u(f)] - \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) + \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]\right) \\
=&\lambda\left(\int_{\Delta\Omega} \max_{f\in F\cup G} \mathbb{E}_p[u(f)]d\mu(p) - \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p)\right) \\
\geq&\gamma \int_{\Delta\Omega} \max_{f\in F\cup G} \mathbb{E}_p[u(f)]d\mu(p) - \gamma \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) \\
=&\gamma \int_{\Delta\Omega} \max_{f\in F\cup G} \mathbb{E}_p[u(f)]d\mu(p) - \gamma \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) + (1-\gamma) \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] - (1-\gamma) \max_{f\in F\cup G} \mathbb{E}_{\hat{p}}[u(f)]
\end{aligned}
$$

where the first inequality follows from the definition of subgradient, the first equality by definition, the second equality follows from $V(F|\mu_0) \geq V(G|\mu_0)$, the third inequality from the fact that $\lambda \geq \gamma$ and the term in the parenthesis is positive. The last equality follows from $V(F|\mu_0) \geq V(G|\mu_0)$. Rearranging gives $V^\gamma(F|\mu) \geq V^\gamma(F \cup G|\mu)$.

□

Proof of Proposition 9. The first result is trivial. Suppose that for a $F \in \mathcal{A}$, $I(\mu, F) = 0$. Then $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \max_{f\in F} v(f) = 0$ implies $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] = 0$ since $\mathbb{E}_{\hat{p}}[u(f)] \geq v(f)$ for all $f \in \mathcal{F}$. But the latter inequality allows to use the same argument of De Lara and Gossner (2020). Lastly, if $I^v(\mu, F) = 0$ $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) = \max_{f\in F} v(f)$, hence $V(F|\mu) = \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \phi(I^v(\mu, F)) = \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \phi(0) = \max_{f\in F} v(f)$. However,

$$\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) \geq \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] \geq \max_{f\in F} v(f)$$

hence, $\max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] = \max_{f\in F} v(f)$ and $V(F|\mu_0) = \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] = \max_{f\in F} v(f)$. □

Proof of Proposition 10. The result follows from a straightforward modification of the argument in the proof of Proposition 4 with $\gamma = 1$. For easy of notation, let define $v_F = \max_{f\in F} v(f)$. For point 1., first notice that $I^v(\mu, F) = \int_{\Delta\Omega} \sigma_F(p)d\mu(p) - v_F > \int_{\Delta\Omega} \sigma_G(p)d\mu(p) - v_G = I^v(\mu, G) \geq \int_{\Delta\Omega} \sigma_G(p)d\mu(p) - v_F$, hence $\int_{\Delta\Omega} \sigma_F(p)d\mu(p) > \int_{\Delta\Omega} \sigma_G(p)d\mu(p)$. If $\lambda \geq \lambda^v_{F,G}$ for some $\lambda \in \partial\phi(I(\mu, G))$, then

$$
\begin{aligned}
\phi(I^v(\mu, F)) - \phi(I^v(\mu, G)) \geq& \lambda\,(I^v(\mu, F) - I^v(\mu, G)) \\
\geq& \lambda^v_{F,G}\left(\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - v_F - \int_{\Delta\Omega} \max_{f\in G} \mathbb{E}_p[u(f)]d\mu(p) + v_G\right) \\
=& \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \int_{\Delta\Omega} \max_{f\in G} \mathbb{E}_p[u(f)]d\mu(p)
\end{aligned}
$$

where the first inequality follows from the definition of subdifferential, the second from the condition $v_F \geq v_G$ and the last inequality by the fact that $\lambda \geq 1$. Rearranging gives $V(G|\mu) \geq V(F|\mu)$.

For point 2., by the two initial assumptions F and G cannot both singletons. Let define $\lambda_{F,G} = \frac{\int_{\Delta\Omega} \sigma_F(p)d\mu(p) - \int_{\Delta\Omega} \sigma_G(p)d\mu(p)}{I^v(\mu,F)-I^v(\mu,G)}$. First notice that $I^v(\mu, F) = \int_{\Delta\Omega} \sigma_F(p)d\mu(p) - \sigma_F(\hat{p}) < \int_{\Delta\Omega} \sigma_G(p)d\mu(p) - \sigma_G(\hat{p}) = I^v(\mu, G)$, hence $\lambda_{F,G} > 0$. Suppose that $\lambda_{F,G} > 1$. Then, $\int_{\Delta\Omega} \sigma_F(p)d\mu(p) - \int_{\Delta\Omega} \sigma_G(p)d\mu(p) < I^v(\mu, F) - I^v(\mu, G)$, that implies $\sigma_F(\hat{p}) < \sigma_G(\hat{p})$, a contradiction to the assumption, hence $\lambda_{F,G} \leq 1$. If $\lambda \in \partial\phi(I^v(\mu, G))$, then

$$\begin{aligned}\phi(I^v(\mu, F)) - \phi(I^v(\mu, G)) \geq& \lambda \left(I^v(\mu, F) - I^v(\mu, G)\right)\\ \geq& \lambda_{F,G} \left(I^v(\mu, F) - I^v(\mu, G)\right)\\ =& \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \int_{\Delta\Omega} \max_{f\in G} \mathbb{E}_p[u(f)]d\mu(p)\end{aligned}$$

where the first inequality follows from the definition of subgradient, the second from the condition $I^v(\mu, F) < I^v(\mu, G)$ and the equality by the definition of $\lambda_{F,G}$. Rearranging gives $V(G|\mu) \geq V(F|\mu)$.

□

Proof of Proposition 11. Suppose that, for some $\lambda \in \partial\phi(D(F))$, $\lambda > 1$. Then by the definition of subdifferential

$$\begin{aligned}\phi(I^v(\mu, F)) - \phi(D(F)) \geq& \lambda(I^v(\mu, F) - D(F))\\ =& \lambda \left(\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - v_F - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] + v_F\right)\\ >& \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]\end{aligned}$$

where the first inequality follows from the definition of subdifferential. The equality from the definition of $I^v(\mu, F)$ and $D(F)$. The last inequality from the fact that $\lambda > 1$ and $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu(p) > \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$. Rearranging, implies information avoidance for F.

For the second implication. Suppose that $\lambda \leq 1$ for some $\lambda \in \partial\phi(I^v(\mu, F))$ Then by the definition of subdifferential $\phi(y) - \phi(I^v(\mu, F)) \geq \lambda(y - I^v(\mu, F))$ for all $y \in \mathbb{R}_+$. Taking $y = D(F)$ implies $\phi(D(F)) - \phi(I^v(\mu, F)) \geq \lambda(D(F) - I^v(\mu, F)) \geq D(F) - I^v(\mu, F)$, since $D(F) \leq I^v(\mu, F)$ for all $F \in \mathcal{A}$ and $\lambda \leq 1$. A contradiction to the fact that there is information avoidance for F. □

Proof of Corollary 4. If information is not acquired a $f \cup g$, then $V^v(f \cup g|\mu) < V^v(f \cup g|\mu_0)$, meaning $\int_{\Delta\Omega} \max_{h\in f\cup g} \mathbb{E}_p[u(h)]d\mu(p) - \phi(I^v(\mu, f\cup g)) < \max_{h\in f\cup g} \mathbb{E}_{\hat{p}}[u(h)] - \phi(\max_{h\in f\cup g} \mathbb{E}_{\hat{p}}[u(h)] - v(f)) = \mathbb{E}_{\hat{p}}[u(f)] - \phi(\mathbb{E}_{\hat{p}}[u(f)] - v(f)) = V^v(f|\mu_0)$, where the inequalities follow from the assumption that $\mathbb{E}_{\hat{p}}[u(f)] \geq \mathbb{E}_{\hat{p}}[u(g)]$ and $v(f) \geq v(g)$. For the opposite implication,

$$\int_{\Delta\Omega} \sigma_{f\cup g}(p)d\mu(p) - \phi(I(\mu, f \cup g)) < \mathbb{E}_{\hat{p}}[u(f)] - \phi(\mathbb{E}_{\hat{p}}[u(f)] - v(f)) = \sigma_{f\cup g}(\hat{p}) - \phi(\sigma_{f\cup g}(\hat{p}) - v(f))$$

hence, information is not acquired at $f \cup g$. □

Proof of Proposition 12. Since information is valuable for F, $V(F|\mu) \geq V(F|\mu_0)$. If $\lambda \geq 1$ for some $\lambda \in \partial\phi(I^v(\mu_0, F))$, then

$$\begin{aligned}
&\phi(I^v(\mu_0, F\cup G)) - \phi(I^v(\mu_0, F)) \geq \lambda(I^v(\mu_0, F\cup G) - I^v(\mu_0, F)) \\
=&\lambda\left(\max_{f\in F\cup G}\mathbb{E}_{\hat{p}}[u(f)] - \max_{f\in F\cup G} v(f) - \max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)] + \max_{f\in F} v(f)\right) \\
=&\lambda\left(\max_{f\in F\cup G}\mathbb{E}_{\hat{p}}[u(f)] - \max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)]\right) \\
\geq& \max_{f\in F\cup G}\mathbb{E}_{\hat{p}}[u(f)] - \max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)]
\end{aligned}$$

where the first inequality follows from the definition of subgradient, the first equality by definition, the second equality follows from $\max_{f\in F} v(f) \geq \max_{f\in G} v(f)$ and the last inequality from the fact that $\lambda \geq 1$ and the term in the parenthesis is positive. Rearranging gives $V(F|\mu_0) \geq V(F\cup G|\mu_0)$. Since ϕ is weakly increasing if $\lambda \geq 1$ for some $\lambda \in \partial\phi(I^v(\mu_0, F))$, it must be the case that $\lambda \geq 1$ for all $\lambda \in \partial\phi(I^v(\mu, F))$, since $I^v(\mu, F) \geq I^v(\mu_0, F)$. Hence, the same argument applied to $I^v(\mu, F\cup G)$ and $I^v(\mu, F)$ shows that $V(F|\mu) \geq V(F\cup G|\mu)$.

□

Proof of Theorem 13. Consider $\lambda^v_{F,g} = \frac{\sigma_F(\hat{p}) - \mathbb{E}_{\hat{p}}[u(g)]}{I^v(\mu_0,F) - I^v(\mu_0,g)}$ and $\lambda_{F\cup g,g} = \frac{\sigma_{F\cup g}(\hat{p}) - \mathbb{E}_{\hat{p}}[u(g)]}{I^v(\mu_0,F\cup G) - I^v(\mu_0,g)}$. The same argument used in the proof of Theorem 4 shows that $\lambda^v_{F,g}, \lambda^v_{F\cup g,g} > 0$. Let define $\lambda^* = \max\left\{\lambda^v_{F,g}, \lambda^v_{F\cup G,G}\right\}$. If $\lambda \in \partial\phi(I^v(\mu_0, g))$ is such that $\lambda \geq \lambda^*$:

$$\begin{aligned}
\phi(I^v(\mu_0, F)) - \phi(I^v(\mu_0, g)) \geq& \lambda\left(I^v(\mu_0, F) - I^v(\mu_0, g)\right) \\
\geq& \lambda^*\left(I^v(\mu_0, F) - I^v(\mu_0, g)\right) \\
\geq& \lambda_{F,g}\left(I^v(\mu_0, F) - I^v(\mu_0, g)\right) \\
=& \max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)] - \mathbb{E}_{\hat{p}}[u(g)]
\end{aligned}$$

Rearranging gives $V^v(g) \geq V^v(F|\mu_0)$. By the convexity of ϕ, if $\lambda \in \partial\phi(I^v(\mu, g))$, then $\lambda' \geq \lambda$ for all $\lambda' \in \partial\phi(I^v(\mu_0, F))$, since $I^v(\mu, F) > I^v(\mu, g)$. It follows that $\lambda' \geq \lambda^*$ then

$$\begin{aligned}
\phi(I^v(\mu_0, F\cup g)) - \phi(I^v(\mu_0, F)) \geq& \lambda'\left(I^v(\mu_0, F\cup g) - I^v(\mu_0, F)\right) \\
\geq& \lambda^*\left(I^v(\mu_0, F\cup g) - I^v(\mu_0, F)\right) \\
\geq& \lambda^v_{F\cup g,g}\left(I^v(\mu_0, F\cup g) - I^v(\mu_0, F)\right) \\
=& \max_{f\in F\cup g}\mathbb{E}_{\hat{p}}[u(f)] - \max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)]d\mu(p)
\end{aligned}$$

where the first inequality follows from the definition of subgradient, the second from the condition and the equality by the definition of $\lambda^v_{F\cup g,g}$. Rearranging gives $V^v(F|\mu_0) \geq V^v(F\cup g|\mu_0)$. Since $I^v(\mu, F\cup g) \geq I^v(\mu_0, F\cup g)$, the same argument shows that $V^v(g) \geq V^v(F\cup g|\mu)$. Moreover,

□

References

Akerlof, G. A. and R. E. Kranton (2000). Economics and identity. *The Quarterly Journal of Economics 115*(3), 715–753.

Akerlof, G. A. and R. E. Kranton (2005, March). Identity and the economics of organizations. *Journal of Economic Perspectives 19*(1), 9–32.

Akerlof, G. A. and R. E. Kranton (2010). *Identity Economics: How Our Identities Shape Our Work, Wages, and Well-Being.* Princeton University Press.

Alauoi, L. (2016). Information avoidance and the preservation of self-image.

Andreoni, J., J. M. Rao, and H. Trachtman (2017). Avoiding the ask: A field experiment on altruism, empathy, and charitable giving. *Journal of Political Economy 125*(3), 625–653.

Andreoni, J. and M. Serra-Garcia (2016, November). Time-inconsistent charitable giving. Working Paper 22824, National Bureau of Economic Research.

Andries, M. and V. Haddad (2019). Information aversion. *Journal of Political Economy, Forthcoming.*

Arrow, K. J. and A. C. Fisher (1974). Environmental preservation, uncertainty, and irreversibility. *The Quarterly Journal of Economics 88*(2), 312–319.

Atkin, D., E. Colson-Sihra, and M. Shayo (2021). How do we choose our identity? a revealed preference approach using food consumption. *Journal of Political Economy* (Forthcoming).

Austen-Smith, D. and R. G. Fryer Jr (2005). An economic analysis of “acting white". *The Quarterly Journal of Economics 120*(2), 551–583.

Bénabou, R. (2015). The economics of motivated beliefs. *Revue d'Économie Politique 125*(5), 665–685.

Bénabou, R. and J. Tirole (2011). Identity, morals, and taboos: Beliefs as assets. *The Quarterly Journal of Economics 126*(2), 805–855.

Blackwell, D. (1953, 06). Equivalent comparisons of experiments. *Ann. Math. Statist. 24*(2), 265–272.

Brocas, I. and J. D. Carrillo (2004). Entrepreneurial boldness and excessive investment. *Journal of Economics & Management Strategy 13*(2), 321–350.

Brunnermeier, M. K. and J. A. Parker (2005). Optimal expectations. *American Economic Review 95*(4), 1092–1118.

Bursztyn, L., A. L. González, and D. Yanagizawa-Drott (2020, October). Misperceived social norms: Women working outside the home in Saudi Arabia. *American Economic Review 110*(10), 2997–3029.

Camerer, C. and D. Lovallo (1999). Overconfidence and excess entry: An experimental approach. *The American Economic Review 89*(1), 306–318.

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

Caplin, A. and J. Leahy (2001). Psychological expected utility theory and anticipatory feelings. *The Quarterly Journal of Economics 116*(1), 55–79.

Carrillo, J. D. and T. Mariotti (2000). Strategic ignorance as a self-disciplining device. *The Review of Economic Studies 67*(3), 529–544.

Clarke, F. H. (1983). *Optimization and Nonsmooth Analysis.* J. Wiley, New York.

Courtenay, W. H. (2000). Constructions of masculinity and their influence on men's well-being: a theory of gender and health. *Social Science & Medicine 50*(10), 1385–1401.

Dana, J., D. M. Cain, and R. M. Dawes (2006). What you don't know won't hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and human decision Processes 100*(2), 193–201.

Dana, J., R. A. Weber, and J. X. Kuang (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory 33*(1), 67–80.

De Lara, M. and O. Gossner (2020). Payoffs-beliefs duality and the value of information. *SIAM Journal on Optimization, Forthcoming.*

DellaVigna, S., J. A. List, and U. Malmendier (2012). Testing for altruism and social pressure in charitable giving. *The Quarterly Journal of Economics 127*(1), 1–56.

Demers, M. (1991). Investment under uncertainty, irreversibility and the arrival of information over time. *The Review of Economic Studies 58*(2), 333–350.

Dillenberger, D., J. S. Lleras, P. Sadowski, and N. Takeoka (2014). A theory of subjective learning. *Journal of Economic Theory 153*(0), 287 – 312.

Dillenberger, D. and P. Sadowski (2012). Ashamed to be selfish. *Theoretical Economics 7*(1), 99–124.

Ellis, A. (2018). Foundations for optimal inattention. *Journal of Economic Theory 173*, 56–94.

Epstein, L. G. (2006). An axiomatic model of non-Bayesian updating. *The Review of Economic Studies 73*(2), 413–436.

Epstein, L. G., J. Noor, and A. Sandroni (2008). Non-Bayesian updating: a theoretical framework. *Theoretical Economics 3*(2), 193–229.

Ergin, H. and T. Sarver (2010). A unique costly contemplation representation. *Econometrica 78*(4), 1285–1339.

Frankel, A. and E. Kamenica (2019, October). Quantifying information and uncertainty. *American Economic Review 109*(10), 3650–80.

Gennaioli, N. and G. Tabellini (2019). Identity, Beliefs, and Political Conflict. Technical report.

Gneezy, U., K. L. Leonard, and J. A. List (2009). Gender differences in competition: Evidence from a matrilineal and a patriarchal society. *Econometrica 77*(5), 1637–1664.

Gollier, C. (2004). *The Economics of Risk and Time.* MIT press.

Golman, R., D. Hagmann, and G. Loewenstein (2017). Information avoidance. *Journal of Economic Literature 55*(1), 96–135.

Grossman, Z. and J. J. van der Weele (2016, 12). Self-Image and willful ignorance in social decisions. *Journal of the European Economic Association 15*(1), 173–217.

Gul, F. (1991). A theory of disappointment aversion. *Econometrica 59*(3), 667–686.

Hertwig, R. and C. Engel (2016). Homo ignorans: Deliberately choosing not to know. *Perspectives on Psychological Science 11*(3), 359–372. PMID: 27217249.

Huck, S., N. Szech, and L. M. Wenner (2018). More effort with less pay: On information avoidance, belief design and performance. *Mimeo.*

Jones, R. A. and J. M. Ostroy (1984). Flexibility and uncertainty. *The Review of Economic Studies 51*(1), 13–32.

Kahan, D. M. (2013). Ideology, motivated reasoning, and cognitive reflection. *Judgment and Decision Making 8*(4), 407–424.

Kandul, S. and I. Ritov (2017). Close your eyes and be nice: Deliberate ignorance behind pro-social choices. *Economics Letters 153*, 54–56.

Köszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association 4*(4), 673–707.

Köszegi, B. and M. Rabin (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics 121*(4), 1133–1165.

Köszegi, B. and M. Rabin (2007). Reference-dependent risk attitudes. *The American Economic Review 97*(4), 1047–1073.

Lipnowski, E. and L. Mathevet (2018). Disclosure to a psychological audience. *American Economic Journal: Microeconomics 10*(4), 67–93.

Loomes, G. and R. Sugden (1986). Disappointment and dynamic consistency in choice under uncertainty. *The Review of Economic Studies 53*(2), 271–282.

Lu, J. (2016). Random choice and private information. *Econometrica 84*(6), 1983–2027.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006, November). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica 74*(6), 1447–1498.

Matějka, F. and A. McKay (2014). Rational inattention to discrete choices: a new foundation for the multinomial logit model. *The American Economic Review 105*(1), 272–298.

Niederle, M. and L. Vesterlund (2007). Do women shy away from competition? Do men compete too much? *The Quarterly Journal of Economics 122*(3), 1067–1101.

Niehaus, P. (2014). A theory of good intentions. *Mimeo.*

Oliveira, H., T. Denti, M. Mihm, and K. Ozbek (2017). Rationally inattentive preferences and hidden information costs. *Theoretical Economics 12*(2), 621–654.

Oster, E., I. Shoulson, and E. Dorsey (2013). Optimal expectations and limited medical testing: evidence from Huntington disease. *American Economic Review 103*(2), 804–30.

Pagel, M. (2018). A news-utility theory for inattention and delegation in portfolio choice. *Econometrica 86*(2), 491–592.

Pennesi, D. (2015). Costly information acquisition and the temporal resolution of uncertainty. *Journal of Mathematical Economics 60*, 115–122.

Piermont, E. (2019). Images and norms. *Mimeo*.

Prentice, D. A. and D. T. Miller (1993). Pluralistic ignorance and alcohol use on campus: some consequences of misperceiving the social norm. *Journal of Personality and SocialPsychology 64*(2), 243.

Rudman, L. A. and K. Mescher (2013). Penalizing men who request a family leave: Is flexibility stigma a femininity stigma? *Journal of Social Issues 69*(2), 322–340.

Sarver, T. (2008). Anticipating regret: Why fewer options may be better. *Econometrica 76*(2), 263–305.

Schweizer, N. and N. Szech (2018). Optimal revelation of life-changing information. *Management Science 64*(11), 5250–5262.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics 50*(3), 665–690.

Spence, A. M. (1977). Entry, capacity, investment and oligopolistic pricing. *The Bell Journal of Economics 8*(2), 534–544.

Stets, J. E. and P. J. Burke (2000). Identity theory and social identity theory. *Social Psychology Quarterly*, 224–237.

Sunstein, C. R. (2020). *Too Much Information: Understanding What You Don't Want to Know.* MIT Press.

Tajfel, H. and J. C. Turner (1979). An integrative theory of intergroup conflict. *Organizational Identity: A Reader 56*, 65.

Vandello, J. A., V. E. Hettinger, J. K. Bosson, and J. Siddiqi (2013). When equal isn't really equal: The masculine dilemma of seeking work flexibility. *Journal of Social Issues 69*(2), 303–321.

Viscusi, W. K. (1990). Do smokers underestimate risks? *Journal of Political Economy 98*(6), 1253–1269.

Williams, J. (2006). One sick child away from being fired: When 'opting out' is not an option. *Available at SSRN 2126303*.

Williams, J. C., M. Blair-Loy, and J. L. Berdahl (2013). Cultural schemas, social class, and the flexibility stigma. *Journal of Social Issues 69*(2), 209–234.