ISSN 2279-9362

Collegio Carlo Alberto

Identity and information acquisition

Daniele Pennesi

No. 610
June 2020 (Revised, June 2021)

Carlo Alberto Notebooks
www.carloalberto.org/research/working-papers

Identity and information acquisition*

Daniele Pennesi†

This version June 2021

Abstract

This paper proposes a theory of identity protection and studies its implications for information acquisition. Identity prescribes behaviors, and violating the prescriptions in favor of actions with higher material value is cognitively costly. Information has an ambiguous value because it can reduce or exacerbate the trade-off between material benefits and identity protection. The theory rationalizes information avoidance observed in various domains, including willful ignorance in social dilemmas. The cost of information, however, is not necessarily monotone in the quality of information, so that inferior information can threaten the identity more–hence it is more likely to be avoided–than better information. Although beliefs do not have hedonic value, the theory is consistent with a demand for beliefs that leads to acquire or reject non-instrumental information. 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 destructive behavior, in which the individual commits to payoff-inferior actions.

Keywords: Identity, Self-image, Information Aversion, Willful Ignorance, Gender Identity
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

*A previous version of this paper circulated under the title "Identity and the cost of information".

†ESOMAS Department and Collegio Carlo Alberto, University of Torino. Email: daniele.pennesi@unito.it

in relation to “who they are". The simplest way to protect identity is to follow the *prescriptions*: appropriate behaviors given the identity.[1] However, prescriptions are often different from payoff-maximizing choices so that an individual who cares about her identity faces a trade-off between material benefits and identity protection. A crucial aspect of such a trade-off is its interaction with uncertainty and with information reducing part or all of the uncertainty. Large evidence suggests that information is not always valuable for identity-concerned individuals. In the now-classic moral wiggle room experiment, individuals avoid information about the effects of their choice on the welfare of others, so as to maintain their self-image of being fair while acting selfishly (Dana, Weber, and Kuang, 2007; Grossman and van der Weele, 2016). This is because information that resolves the uncertainty forces the individual to face the trade-off between identity protection (fairly splitting the endowment) and own payoff-maximization, while under ignorance such a trade-off is eschewed. At the same time, not all information is costly for an identity-concerned individual: learning that the only available meal is meat-free makes a starving vegetarian better off.

Existing approaches to model the interaction between information and identity are mostly focused on moral behavior and assume that individuals are unsure about their “deep values”, such as being altruist, and take costly actions to self-signal identity (Bénabou and Tirole, 2011; Grossman and van der Weele, 2016). However, identity-driven information preferences are also observed when there is little to no uncertainty about identity (e.g. gender or political identity). For example, to protect the image of being “strong, robust and tough”, males are less likely to get routine medical tests (Courtenay, 2000), but most males are sure about their gender identity, hence making the self-signalling interpretation of behavior problematic. A second critical aspect of identity-uncertainty concerns the cost of signalling identity. In these approaches, following the prescrip-

[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 terms of physical characteristics and other attributes....agents follow prescriptions, for the most part, to maintain their self-concepts..." 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.”

tions to self-signal identity, for example by taking a prosocial action, is *always* (directly) costly. However, not all prescriptions are costly, as often prescriptions coincide with the payoff-maximizing actions. For a vegetarian, eating vegetables may simultaneously be the prescription and the payoff-maximizing action. Instead, eating meat is the (cognitively and probably physically) costly behavior. In fact, the theory of identity[2] is based on the assumption that *violating* rather than obeying a prescription is costly.

In this paper, I develop a new theory of identity protection in which the individual is sure about her identity and bears a cognitive cost when she *violates* the prescriptions. I then apply the theory to study information acquisition. The theory is "rational", since the individual processes information in a Bayesian way, and evaluates all actions according to their material value.[3] It is driven by behavior, in that identity protection and the value of information are determined by actual behavior rather than beliefs. However, the theory is consistent with the avoidance of instrumental information, commitment to (payoff) destructive behaviors, and a "demand for beliefs" that leads to acquire or reject non-instrumental information.

I consider a two-period choice problem where an individual selects her future opportunities (a menu of actions whose payoffs depend on an unknown state of nature), acquires information or remains ignorant, and lastly selects an action from the menu. Informed by the theory of (social) identity (e.g. Tajfel and Turner, 1979; Stets and Burke, 2000; Akerlof and Kranton, 2000), the model has a unique substantive assumption: acting against the prescriptions is *costly*. For example, due to the "cognitive dissonance"[4] or the anxiety evoked by violating a norm. The cost of violating the prescription depends directly on the divergence between the material value of obeying the prescriptions and the material value of the payoff-maximizing behavior. The choice to acquire information

[2]For example, in Akerlof and Kranton (2000) pp. 716, "violating the prescriptions evokes anxiety and discomfort in oneself".

[3]Empirical evidence that partly supports the rationality 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-protection requires cognitive abilities.

[4]Festinger (1962) "if a person knows various things that are not psychologically consistent with one another, he will, in a variety of ways, try to make them more consistent....The items of information may be about behavior, feelings, opinions, things in the environment and so on....A person can change his opinion; he can change his behavior, thereby changing the information he has about it."

balances contrasting forces: on one hand, information has a positive value, as it helps finding the payoff-maximizing action. Information can also change the prescription, potentially aligning the "informed" prescription to the payoff-maximizing behavior. For the starving vegetarian, the prescription under ignorance would be to refrain from eating the only available meal. Learning that it is actually meat-free is valuable because now the prescriptive behavior (consume the meal) and the payoff-maximizing behavior coincide. On the other hand, information may be "costly" because it can generate or exacerbate the trade-off between payoff-maximization and identity protection, as in the moral wiggle room example.[5]

The first prediction of the model is that the cost of information can be too high compared to its value, so that information avoidance is optimal. The result applies to willful ignorance in social dilemmas (Dana et al., 2007; Grossman and van der Weele, 2016), ignorance concerning the effects of harmful consumption (Viscusi, 1990), avoidance of routine health screenings (Courtenay, 2000) and pluralistic ignorance (Bursztyn, González, and Yanagizawa-Drott, 2020). Avoidance of information arises in the model in absence of cognitive constraints, motivated beliefs, or any other nonstandard assumption on the evaluation of actions (see Section 7). Thus, the present theory provides a rational explanation for identity-driven information avoidance.

Second, the cost and value of information are not necessarily monotone in the quality of information. Better information (in the sense of Blackwell) may be less threatening to identity, so that it can be avoided more often than "poorer" information. Since the cost of information depends on the discrepancy between the payoff-maximizing action and the prescription, when the information reduces or even aligns the payoff-maximizing and the prescriptive actions, its cost diminishes or vanishes. Consider a male worker whose masculine identity prescribes to go to work despite having mild symptoms of a disease (Glick, Berdahl, and Alonso, 2018). There is a costly discrepancy between the payoff-maximizing action (staying home) and the prescription. If, however, he receives a

[5] Konow (2000) proposed a model of fairness based on the idea that individuals seek to reduce cognitive dissonance, either by changing behavior or by adjusting beliefs. In the approach of the present paper, acquiring (avoiding) information, when it reduces (exacerbates) the identity trade-off, makes information an instrument to reduce cognitive dissonance, in the same spirit of Konow (2000).

precise signal about his health status, such as a diagnosis of a dangerous disease, staying home represents a less costly violation of the prescription. In this case, more information reduces if not eliminates the cost of violating the prescription.

Third, the model is consistent with a demand for "beliefs" that leads to reject or seek non-instrumental information. This is because prescriptions vary with beliefs through the acquired information. So, even information that does not affect actual behavior (non-instrumental), may be still valuable or costly because it modifies the beliefs, hence the associated prescriptions. A religious person may prefer not knowing if a life-saving drug contains prohibited substances (e.g. alcohol or parts of pigs or bovines). Information is non-instrumental in the sense that the optimal behavior is to take the drug in any case, but learning that the drug contains prohibited substances is extremely costly, hence ignorance is optimal.

A last prediction of the theory is that an identity-caring individual can display a strict preference for having fewer options. In larger menus, the discrepancy between the prescribed behavior and the payoff-maximizing action is more likely to occur, hence the option value of a larger menu can be outweighed by the larger cost of violating the prescriptions. A preference for commitment can arise both in the absence or in the presence of information arrival (see Section 5). In the latter case, if the information does not add prescriptive value in a larger menu, commitment before information acquisition is optimal. For instance, the protection of an altruistic identity is consistent with escaping a situation where it is possible to act prosocially (e.g. Dana, Cain, and Dawes, 2006; DellaVigna, List, and Malmendier, 2012; Andreoni, Rao, and Trachtman, 2017). Entering a social dilemma forces the individual to face the trade-off between material benefits and identity protection. Anticipating the cost of potentially violating the prescription, the individual may prefer to shy away from the dilemma. If opting in also entails information acquisition, the effect can be stronger. For instance, if meeting the fundraiser implies costly information arrival, for example learning that the charity has low quality, it could be optimal to eschew the meeting (see Corollary 1). Lastly, commitment to payoff-inferior actions becomes optimal when committing to superior action is unfeasible, hence generating an identity-driven form of "destructive" behavior. 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 himself as *bold*, for example, can rationally decide to overinvest in a market as a means to protect his identity, even if he knows that the ex ante optimal behavior is to not enter the market. 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 can rationalize why males select into competitive environments more often than females, but only in patriarchal societies (Niederle and Vesterlund, 2007; Gneezy, Leonard, and List, 2009), why fathers under-use 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). In the last case, the choice to underinvest in human capital can occur even without an explicit role for signalling to the reference group (e.g. to avoid "acting white"). The potential cognitive dissonance of choosing a "white" job–a violation of the prescription–represents a sufficient condition for underinvesting in human capital.

I also study welfare-maximizing information disclosure to an identity-caring individual. 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).

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 and motivates the approach through an example. Section 3 introduces the model. Section 4 contains its application to information acquisition and Section 5 to preference for commitment. Section 6 studies optimal disclosure to an identity-caring individual. 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 will develop the model. 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 actions is denoted by $\mathcal{F}$. Constant actions, $f(\omega) = x$ for all $\omega \in \Omega$, are identified with the element of X they deliver. In the model, the individual evaluates menus: nonempty and compact 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 dynamics of the decision process is the following (see Fig. 1): the individual selects a menu, acquires information (if any), forms a posterior, and then selects an action from the menu before uncertainty resolves. Lastly, uncertainty resolves and the payoff materializes.

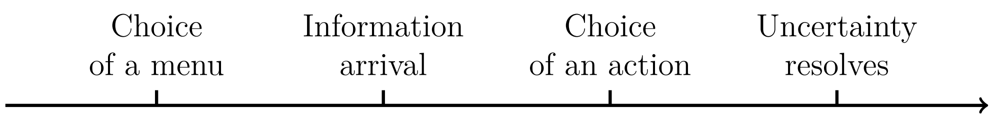

Figure 1: Dynamic of the choice process

Information. I model information as a "Bayesian experiment", a distribution over posteriors 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 $\hat{p} = \int_{\Delta\Omega} q d\mu(q)$ (see, for example, Gollier, 2004). The experiment μ specifies the probability of deriving each posterior q by Bayesian updating the prior $\hat{p}$.[6] The family of all experiments μ consistent with a prior $\hat{p}$ is denoted by $\Gamma(\hat{p})$. The next definition introduces the well-known Blackwell's informativeness order on $\Gamma(\hat{p})$ (see Blackwell, 1953):

Definition 1. *Given $\mu, \nu \in \Gamma(\hat{p})$ for some $\hat{p} \in \Delta\Omega$, ν is Blackwell more informative than μ, written $\nu \trianglerighteq \mu$, if*

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

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

An experiment ν is Blackwell more informative than μ, if any experimenter (associated to a payoff function ϕ) will prefer ν to μ. Lastly, 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 $\bar{\mu}(q) = 0$ if $q \in \Delta\Omega$ and $q \neq \delta_\omega$.

2.1 A motivating example: Moral wiggle room

Growing empirical evidence shows that some individuals exploit uncertainty in social dilemmas to maximize their payoff without compromising their self-image of being altruist (e.g. Dana et al., 2007; Grossman and van der Weele, 2016). Since uncertainty is functional to payoff maximization, the individual can reject information resolving the uncertainty, so as to exploit a "moral wiggle room". Theoretical explanations based on identity or self-image preservation require a combination of imperfect memory and self-control (Bénabou and Tirole, 2011) or an explicit role for self-signalling (Grossman and

[6] 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.

van der Weele, 2016). In both cases, the substantive assumption is that individuals are unsure of their deep values, such as being altruist, and take costly actions to remind themselves who they are. In this paper, I argue that an altruistic individual who is *sure* about her identity, but pays a cost when deviating from the prescription, can rationally decide to avoid information in a social dilemma. This is because information creates a trade-off, absent under ignorance, between the payoff-maximizing action and the protection of identity. Regardless of how it will be resolved, the trade-off is costly: taking the payoff-maximizing action implies a costly deviation from the prescription; by choosing the prescription, the cognitive cost disappears, but the material payoff is lower with respect to ignorance (details below).

Consider the choice that subjects face in the "moral wiggle room" experiment of Dana et al. (2007) and Grossman and van der Weele (2016). There are two actions f, g, two equally probable states of the world $\Omega = \{\omega_1, \omega_2\}$ and two individuals. The payoff of an action f represents an allocation of monetary amounts $f(\omega) = (x, y)$, where x is the amount for the decision-maker (the dictator) and y for the other individual: In their

	ω_1	ω_2
a	$(6, 5)$	$(6, 1)$
b	$(5, 1)$	$(5, 5)$

Table 1: Actions and payoffs in Dana et al. (2007)

experiment, Dana et al. (2007) observed that b was selected by 74% of the subjects,[7] when they knew that the state of the world was ω_2. 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. Among those who preferred to remain ignorant, nobody selected action b. The results support the intuition that uncertainty is exploited as a moral wiggle room. The present theory rationalizes willful ignorance in the following way: under ignorance, the prescriptive behavior and the (selfish) payoff-maximizing action coincide and are both equal to a. Clearly, a is the

[7] A similar proportion is found in Grossman and van der Weele (2016).

choice that maximizes the selfish payoff of the dictator. Moreover, a is prescriptively as good as b, since the payoffs for the other individual are, ex ante, the same under a and b. Therefore, choosing a under ignorance does not threaten the dictator's image of being fair (at least ex ante fair). Consider what happens when information arrives: if the state is ω_1, then a is again the optimal choice from both the payoff-maximizing and prescriptive points of view. However, in state ω_2 the prescription for an altruistic identity is b and it is different from the own payoff-maximizing action a. Hence, in state ω_2, the trade-off between payoff maximization a and the prescriptive behavior b materializes. Once in state ω_2, choosing a implies a costly deviation from the prescription b. By choosing b, the cost disappears, but the payoff is lower than under ignorance (since $(5, 5)$ is prescriptively superior but has a lower material value than $(6, 1)$). In both cases (selecting a or b in state ω_2) ignorance can be preferred to information acquisition (See Fact 1), which is consistent with the modal choice in Dana et al. (2007). An important aspect of the previous example is that the prescriptions vary with information. The prescription under ignorance (e.g. a) may be different from the prescriptions after information arrives (e.g. a or b).

A second and less intuitive consequence of a desire to protect identity is a potential preference for having fewer options. 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 better. Indeed, an individual expecting to receive information prefers larger menus because she will exploit 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 a, hence eliminating b. In this case, facing a can be strictly more valuable than acquiring information and select an action from $a \cup b$. This is because by eliminating b, information acquisition does not create a trade-off.

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

3 The model

Material preferences. The material preference of the individual is represented by an expected utility with respect to a prior $\hat{p} \in \Delta\Omega$ and utility $u : X \to \mathbb{R}$. The value of an action $f \in \mathcal{F}$ is $\mathbb{E}_{\hat{p}}[u(f)]$ and the action $f \in \text{argmax}_{g \in F} \mathbb{E}_{\hat{p}}[u(g)]$ is (one of) the payoff-maximizing action in a menu F.

Identity and prescriptions. Individuals self-categorize or *identify* themselves as members of a social category or occupant of a role.[9] Once established, identity determines prescriptions or norms: appropriate behaviors 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."

Prescriptions are *internalized* (see Akerlof and Kranton, 2000, p. 728) and they determine a ranking of actions.[10] In a menu F, the prescribed action is

$$f_{F,p} = \underset{g \in F}{\text{argmax}}\, U(g|p, F)$$

where $U : \mathcal{F} \times \Delta\Omega \times \mathcal{A} \to \mathbb{R}$ represents the prescriptive ranking.[11] The function U

[9] As in Bénabou and Tirole (2011), I do not distinguish between (personal) identity and social identity. The former is based on the concept of "role" that an individual (believes) occupies in a society (Stets and Burke, 2000), while the latter focuses on "belonging" to a social category or a group (Tajfel and Turner, 1979; Abrams and Hogg, 2006). Though distinct, the two notions differ more in the terminology than in the substance, as argued by Stets and Burke (2000).

[10] In social identity, this assumption is referred to as *depersonalization*, "or seeing the self as an embodiment of the in-group prototype". In personal identity, the process is called *self-verification*, or "seeing the self in terms of the role as embodied in the identity standard" (Stets and Burke, 2000).

[11] Allowing for a unique prescription is a rather weak requirement. If there are multiple prescriptions, I consider a simple tie-breaking rule: for example, the individual considers the prescription that maximizes her material preference. That is, let denote by $P(F, p) = \text{argmax}_{f \in F} U(f|p, F)$, then the unique (up to indifference) prescription is $f = \text{argmax}_{f \in P(F,p)} \mathbb{E}_p[u(f)]$. For example, in the moral wiggle room, it is

depends on:

Beliefs. I consider the prescription to be (potentially) a function of the beliefs. As illustrated in the motivating example, the prescription could vary with information. In particular, if the dictator learns that the true state of the world is ω_2, the prescription becomes b. The same conclusion can be drawn if the dictator does not receive full information, but the realized posterior places a (very) high probability on state ω_2. Therefore, I consider prescriptions that are functions of the belief held by the individual at the moment of choice, and I call $f_{F,p}$ the p-belief prescription in a menu F. I call an action $f \in F$ an "absolute prescription for F" if it does not depend on the beliefs, and I denote it by f_F. Absolute prescriptions derive from extreme forms of identification, for instance, political or religious extremeness. The prescriptions of a committed climate change denier facing a choice between buying or not a hybrid car will hardly adapt to new beliefs.

Situation. Identity is multidimensional and different identities of the same individual may be more or less relevant in a given situation. Situational factors activate one identity and, once activated, "the person behaves so as to maintain consistency with the identity standard" (Stets and Burke, 2000). Here, a situation coincides with the menu F. More broadly, identity "salience" depends on personal factors, for example, the *accessibility* of an identity: the readiness of a role or category to become activated (see, Akerlof and Kranton, 2000; Gennaioli and Tabellini, 2019).[12]

Information, prescriptions and material preferences. I assume that the individual is of type $\gamma \in [0, 1]$, where γ represents the probability that the choice *from* the menu coincides with the payoff-maximizing action and $1 - \gamma$ is the probability that the choice coincides with the prescription. Uncertainty may reflect personal characteristics and can be correlated with situational factors, such as the observability of the second-period choice. An interpretation of γ comes from the Self-Discrepancy theory of Higgins

possible that f and g are prescriptively equivalent under uncertainty, i.e. $f, g \in \text{argmax}_{h \in f \cup g} U(h|\hat{p}, f \cup g)$. In this case, the individual breaks the tie by considering, as the prescription, the materially superior action f.

[12]Other factors affects identity salience, for example, the perceived social distance from a reference-group as well as the status of the group. In strategic situations, the behavior and payoffs of the other players also modify identity salience. However, to discipline the model and given the focus of the model on information acquisition, I limit identity to vary with beliefs and the situation.

(1989), where the individual conceives multiple selves: the actual self and the ideal/ought selves. In this interpretation, γ represents the probability that the "actual self" will decide in the second period, and $1-\gamma$ the probability that the "ideal/ought" selves will decide. Privately or publicly choosing an action can also affect γ. For example, Bursztyn, Callen, Ferman, Gulzar, Hasanain, and Yuchtman (2019) found that costly avoiding an identity-threatening action for Pakistani men (expressing gratitude to the US government) was more likely when the action was private, rather than public. This means that the propensity to follow the prescription $(1-\gamma)$ decreased when the choice was public. An equivalent assumption is the existence of an identity, activated with probability γ, for which the prescription always coincides with the payoff-maximizing action. For the Pakistani men, expressing an anti-American sentiment signals belonging to a minority, so that two competing identities coexist. When the choice was private, the anti-American identity prevailed, whereas the "conformist" identity is activated when the choice is public. For the latter identity, the prescription coincides with payoff-maximization.

For an individual of type γ, I assume that the ex ante material value of learning μ when facing a menu F is:

$$W_\gamma(\mu, F) = \gamma \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)]d\mu(q) + (1-\gamma)\int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\mu(q)$$

With probability γ, she will select the payoff-maximizing action from the menu given her posterior. The expression $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)]d\mu(q)$ represents the ex ante average expected value that can be obtained from F after acquiring μ (see, for example, Arrow and Fisher, 1974; Gollier, 2004; Dillenberger et al., 2014). With probability $1-\gamma$, she will follow the q-belief prescription conditional on a posterior q. The second expression, $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\mu(q)$, therefore represents the (average) expected utility of following the prescriptions in F. In the Example 2.1, b is the δ_{ω_2}-belief prescription and a is the δ_{ω_1}-belief prescription. Therefore, $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{a\cup b,q})]d\bar{\mu}(q) = \frac{1}{2}u(a) + \frac{1}{2}u(b)$. If there is an absolute prescription f_F for F, it follows from the martingale property of Bayesian updating that $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\mu(q) = \int_{\Delta\Omega} \mathbb{E}_q[u(f_F)]d\mu(q) = \mathbb{E}_{\hat{p}}[u(f_F)]$. Therefore, information does not affect the ex ante value of following the prescription in the presence of an absolute prescription. Note also that if $F = f$ then $W_\gamma(\mu, f) = \mathbb{E}_{\hat{p}}[u(f)]$, for all $\gamma \in [0,1]$. In a singleton, there is no second-period choice, hence information and identity play no role. In the absence of information arrival $W_\gamma(\mu_0, F) = \gamma \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] + (1-\gamma)\mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})]$.

If $\gamma = 1$ information is fully exploited to tailor the second-period choice from the menu. It follows that $W_1(\mu, F) \geq W_1(\mu_0, F)$ for all $F \in \mathcal{A}$, thus information always increases the material value of a menu.[13] If $\gamma < 1$, the situation is different. Indeed, information can also modify the prescriptions in a way that the overall variation of material value is negative. This occurs if the "loss" due to the average variation of the prescriptions $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\mu(q) - \mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})]$, overcomes the gain from learning $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)]d\mu(q) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$. For instance, consider the actions a, b with the payoffs given in Table 1. If $u(x,y) = x$, $\hat{p}(\omega) = \frac{1}{2}$, $a = f_{a\cup b,\delta_{\omega_1}}$, $b = f_{a\cup b,\delta_{\omega_2}}$ and $\mu = \bar{\mu}$ the perfectly revealing information, the material value of $\bar{\mu}$ when facing $F = a \cup b$ is $W_\gamma(\mu, a\cup b) = \gamma \int_\Omega \max_{f\in a\cup b} f(\omega)d\hat{p}(\omega) + (1-\gamma)\int_\Omega f_{a\cup b,\delta_\omega}(\omega)d\hat{p}(\omega) = \gamma\left(\frac{1}{2}\cdot a(\omega_1) + \frac{1}{2}\cdot a(\omega_2)\right) + (1-\gamma)\left(\frac{1}{2}\cdot a(\omega_1) + \frac{1}{2}\cdot b(\omega_2)\right) = \gamma 6 + (1-\gamma)\frac{11}{2}$. Depending on γ, the material value of perfect information for $a \cup b$ is between 6 and $\frac{11}{2}$. Under ignorance, $\mu_0 = \delta_{\hat{p}}$, $W_\gamma(\mu_0, a\cup b) = \mathbb{E}_{\hat{p}}[u(a)] = 6$. Since the ex ante payoff-maximizing action f is also the prescription, information changing the prescription (in state ω_2) has a negative material value. A negative material value of information is excluded when there is an absolute prescription f_F. In that case, $W_\gamma(\mu, F) \geq W_\gamma(\mu_0, F)$ for all $\gamma \in [0,1]$. Indeed, $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)]d\mu(q) \geq \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$ and the ex ante value of following the prescription is the independent of information $\int_{\Delta\Omega} \mathbb{E}_q[u(f_F)]d\mu(q) = \mathbb{E}_{\hat{p}}[u(f_F)]$.

The dual role of information. The main assumption of the present model is that the potential discrepancy between the second-period choice (the actual behavior) and the prescription generates a cognitive cost. Indeed, "violating the prescriptions evokes anxiety and discomfort in oneself" (Akerlof and Kranton, 2000). For the personal interpretation of identity, both cognitive dissonance (Festinger, 1962) and the Self-Discrepancy theory of Higgins (1989) rationalize costly deviations from the identity norms.[14] For the social

[13]This is a standard result and follows from the convexity of $q \mapsto \max_{f\in F} \mathbb{E}_q[u(f)]$ that implies $W_1(\mu, F) = \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)]d\mu(q) \geq \max_{f\in F} \int_{\Delta\Omega} \mathbb{E}_q[u(f)]d\mu(q)$, and from the martingale property of Bayesian updating that implies $\max_{f\in F} \int_{\Delta\Omega} \mathbb{E}_q[u(f)]d\mu(q) = \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] = W_1(\mu_0, F)$.

[14]Indeed, discrepancy between the selves cause negative emotions (see Stets and Serpe, 2016, p. 301).

interpretation of identity, deviations from group norms are costly also because of the fear of group rejection or retaliation (Akerlof and Kranton, 2000). I model the cost of an identity threat as an increasing function of the *average cost of deviating from the prescriptions in* F as measured by expected utility. Conditional on a posterior q, the material cost of following the q-belief prescription is $\max_{f\in F}\mathbb{E}_q[u(f)]-\mathbb{E}_q[u(f_{F,q})]$. Since information affects both the payoff-maximizing choice and the prescription, the average cost of deviating from the prescriptions in F when learning μ is

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

and I assume that the overall cost of information μ is $\phi(I(\mu,F))$, for some increasing, proper and convex function ϕ with $\phi(0)=0$. Examples of ϕ are $\phi(x)=\kappa x^2$ for some $\kappa\geq 0$, or $\phi(x)=0$ for $x\leq x^*$ and $\phi(x)=\theta x$ for $x\geq x^*$ and $\theta\geq 0$. This functional form for the cost of information is motivated by the fact that a menu is evaluated before information arrives. However, costly deviations from the prescriptions occur only with probability γ, namely if the second-period choice coincides with the payoff-maximizing action. To conclude, in the present model the value of a menu F when information is acquired and given $\gamma\in[0,1]$ is

$$V_\gamma(F|\mu)=\int_{\Delta\Omega}\gamma\max_{f\in F}\mathbb{E}_q[u(f)]+(1-\gamma)\mathbb{E}_p[u(f_{F,q})]d\mu(q)-\gamma\phi\left(I(\mu,F)\right). \tag{I}$$

The model balances the material value of information and its cognitive cost. Under ignorance the value of a menu F is given by:

$$V_\gamma(F|\mu_0)=\gamma\max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)]+(1-\gamma)\mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})]-\gamma\phi\left(I(\mu_0,F)\right) \tag{U}$$

where

$$I(\mu_0,F)=\max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)]-\mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})].$$

Since information is costless for singletons, by the Bayesian consistency condition, it follows that $V_\gamma(f|\mu)=V_\gamma(f|\mu_0)=\mathbb{E}_{\hat{p}}[u(f)]$ for all $\gamma\in[0,1]$ and I denote this common value by $V_\gamma(f) = \mathbb{E}_{\hat{p}}[u(f)]$, for all $f \in \mathcal{F}$.

Properties of $\phi(I(\cdot, F))$. Note that $I(\mu, F)$ is always greater or equal than zero and that information is costless if the second-period choice is trivial,[15] $\phi(I(\mu, f)) = 0$ for all $f \in \mathcal{F}$. In my interpretation, information cannot create a trade-off in a singleton. The function $\phi(I(\cdot, F))$ has the following properties:

Proposition 1.

1. *For all $F \in \mathcal{A}$, $\phi(I(\cdot, F)) : \Gamma(\hat{p}) \to (-\infty, \infty]$ is positive, convex, and there exists $k_F \geq 0$ such that $\phi(I(\cdot, F)) - k_F$ is normalized. Moreover, if $q \mapsto \mathbb{E}_p[u(f_{F,q})]$ is concave $\phi(I(\cdot, F))$ is (Blackwell's) monotone.*

2. *If there is an absolute prescription for F, $\phi(I(\cdot, F))$ is positive, convex and (Blackwell's) monotone. Moreover, there is $\eta_F \geq 0$ such that $\phi(I(\cdot, F)) - \eta_F$ is normalized.*

Normalization means that the uninformative experiment is costless (i.e. $\phi(I(\mu_0, F)) = 0$). Monotonicity means that more informative experiments, in the sense of Blackwell (Def. 1) are more costly (i.e. $\phi(I(\nu, F)) \geq \phi(I(\mu, F))$ if $\nu \trianglerighteq \mu$), and convexity means that the marginal cost of information is increasing (see Appendix A). An important property of $\phi(I(\cdot, F))$ is its potential non-monotonicity with respect to the Blackwell's informativeness order.[16] This means that more information is not necessarily more threatening to the identity (unless there is an absolute prescription), as the following example illustrates:

Example 1 (Non-monotone cost of information)*.* Consider the actions and payoffs of Table 1. Assume again that $u(x, y) = x$ and $\mu = \bar{\mu}$, the perfectly revealing information. Suppose now that the prescription under ignorance is b, i.e. $f_{a\cup b,\hat{p}} = b$. The posterior-conditional prescriptions are, as before, $f_{a\cup b,\delta_{\omega_1}} = a$ and $f_{a\cup b,\delta_{\omega_2}} = b$. Then $I(\bar{\mu}, a \cup b) = \frac{1}{2}$, whereas

[15] Indeed, $\phi(I(\mu, f)) = \phi\left(\int_{\Delta\Omega} \mathbb{E}_q[u(f)] - \mathbb{E}_q[u(f)] d\mu(q)\right) = \phi(0) = 0$, where the last equality from the property of ϕ.

[16] While the general cost of information $\phi(I(\cdot, F))$ is not necessarily Blackwell's monotone, under special conditions it is related to a well-known information costs. As shown in Fact 1 in Appendix A, under special conditions $I(\mu, F) = \int_{\Delta\Omega} H(\hat{p}) - H(q) d\mu(q) = \int_{\Delta\Omega} D(q\|\hat{p}) d\mu(q)$ where $H(q) = -\sum_{\omega\in\Omega} q(\omega) \ln q(\omega)$ is the entropy of $q \in \Delta\Omega$ and $D(q\|\hat{p}) = \sum_{\omega\in\Omega} q(\omega) \ln\left(\frac{q(\omega)}{\hat{p}(\omega)}\right)$ is the Kullback-Leibler divergence between q and $\hat{p}$. This corresponds to the typical cost of information used in the rational inattention theory (Sims, 2003; Matějka and McKay, 2014).

$I(\mu_0, a \cup b) = \max_{h \in a \cup b} \mathbb{E}_{\hat{p}}[u(h)] - \mathbb{E}_{\hat{p}}[u(f_{a\cup b,\hat{p}})] = \mathbb{E}_{\hat{p}}[u(a)] - \mathbb{E}_{\hat{p}}[u(b)] = 6 - 5 = 1$. Hence, $I(\bar{\mu}, a \cup b) < I(\mu_0, a \cup b)$. Non-monotonicity arises from the fact that the information $\bar{\mu}$ affects, through the posteriors, the prescriptions.

In Example 1, information increases the material value of the prescription $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\mu(q)$ and hence it compensates the cost of deviating from the prescription. It is also clear that, in the Example 1, the function $q \mapsto \mathbb{E}_q[u(f_{a\cup b,q})]$ is not concave.[17] Non-monotonicity of the cost of information has important consequences. For example, for a masculine worker whose identity prescribes to "put work first" (see for example, Glick et al., 2018), leaving work because of a mild sign of disease is more identity damaging than leaving because of a precise diagnosis of a serious illness (see Section 4.3). This result can be used to inform the design of incentive compatible health screenings, for example, in the presence of identity-concerned individuals. Information that aligns the prescriptions with payoff-maximizing actions is valuable even in the presence of identity concerns.

A more extreme case of non-monotonicity occurs when the q-belief prescription is equal to the payoff-maximizing action for all the possible posteriors (up to μ measure zero) i.e. $\max_{f \in F} \mathbb{E}_q[u(f)] = \mathbb{E}_q[u(f_{F,q})]$ μ-a.s. In this case, information completely *eliminates* the cost of deviating from the prescription, implying $I(\mu, F) = 0$. It turns out that this is the unique case in which $I(\mu, F) = 0$, as the next result shows:

Proposition 2.

1. *$I(\mu, F) = 0$ if and only if there exist actions $f_{F,q} \in F$ such that $f_{F,q} \in \operatorname{argmax}_{g \in F} \mathbb{E}_q[u(g)]$, μ-a.s.*

2. *If $\mathbb{E}_{\hat{p}}[u(h)] \geq \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\mu(q)$ for some $h \in F$ and $I(\mu, F) = 0$, there exists $f \in \operatorname{argmax}_{g \in F} \mathbb{E}_{\hat{p}}[u(g)]$ such that $f \in \operatorname{argmax}_{g \in F} \mathbb{E}_q[u(g)]$, μ-a.s.*

Notice that the condition $\mathbb{E}_{\hat{p}}[u(h)] \geq \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\mu(q)$ in point 2 always holds if there is an absolute prescription for F. In that case, $I(\mu, F) = 0$ implies that the absolute prescription is also the payoff-maximizing action for *all* the posteriors.

[17] Indeed, $\mathbb{E}_{\hat{p}}[u(f_{a\cup b,\hat{p}})] = \mathbb{E}_{\frac{1}{2}\delta_{\omega_1}+\frac{1}{2}\delta_{\omega_2}}[u(b)] = 5 < \frac{1}{2}\mathbb{E}_{\delta_{\omega_1}}[u(f_{a\cup b,\delta_{\omega_1}})] + \frac{1}{2}\mathbb{E}_{\delta_{\omega_2}}[u(f_{a\cup b,\delta_{\omega_2}})] = \frac{1}{2}\mathbb{E}_{\delta_{\omega_1}}[u(a)] + \frac{1}{2}\mathbb{E}_{\delta_{\omega_2}}[u(b)] = \frac{1}{2}6 + \frac{1}{2}5 = 5.5$.

4 Information choice, beliefs and applications

4.1 Information avoidance

The choice to acquire or avoid information balances the material value and the cost of information. Information avoidance is formally defined as:

Definition 2. *There is information avoidance for $F \in \mathcal{A}$ if $V_\gamma(F|\mu_0) > V_\gamma(F|\mu)$.*

The strict inequality represents the fact that avoidance must be an "active" choice, hence subjects to a small but strictly positive cost (see e.g. Golman, Hagmann, and Loewenstein, 2017). Let define the benefit-cost ratio of information for F:

$$\lambda_F = \frac{W_\gamma(\mu, F) - W_\gamma(\mu_0, F)}{I(\mu, F) - I(\mu_0, F)}$$

and $\lambda_F = 0$ if $I(\mu, F) - I(\mu_0, F) = 0$. Notice that $\lambda_F = \gamma$ if there is an absolute prescription in F and $I(\mu, F) - I(\mu_0, F) \neq 0$. Next, we recall the notion of subdifferential of a convex function. The subdifferential of ϕ is at $x \in \mathbb{R}$ is defined as:[18]

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

The next result provides a condition to observe information avoidance:

Theorem 1. *For all $F \in \mathcal{A}$, there is information avoidance for F if $I(\mu, F) > I(\mu_0, F)$ and $\gamma\lambda > \lambda_F$ for some $\lambda \in \partial\phi(I(\mu_0, F))$. If there is an absolute prescription $f_F \in F$, information avoidance for F occurs if $I(\mu, F) > I(\mu_0, F)$ and $\lambda > 1$.*

Information avoidance occurs if the marginal cost of information $\gamma\lambda$ is larger than its (net) marginal value λ_F. When $\gamma = 0$, the conditions $\gamma\lambda = 0 > \lambda_F$ and $I(\mu, F) > I(\mu_0, F)$ imply that $\mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})] > \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\mu(q)$, which represents information avoidance for an individual that surely follows the prescriptions in the second period. If there is an absolute prescription in F, information avoidance occurs if $\gamma\lambda > \gamma$, that is possible only if

[18] The set $\partial\phi(x)$ is convex, closed and non-empty on $relint(dom\, f)$, the relative interior of the domain of f.

$\gamma > 0$. In this case, the sufficient condition becomes $\lambda > 1$. Intuitively, with an absolute prescription, the marginal value of information is equal to γ while the marginal cost is equal to $\gamma\lambda$. Clearly, if $\gamma = 0$, in the presence of an absolute prescription information is irrelevant, since the second period choice under information acquisition or avoidance is the same and coincides with the absolute prescription.

Remark 1. Theorem 1 rationalizes willful ignorance in moral dilemmas (e.g. Dana et al., 2007; Grossman and van der Weele, 2016). Consider the two actions a, b with payoffs in Table 1 with the usual $\hat{p}(\omega) = \frac{1}{2}$, $u(x, y) = x$ and $\mu = \bar{\mu}$. The results above shows that, if $f_{a\cup b,\hat{p}} = a$, $f_{a\cup b,\delta_{\omega_2}} = b$ and $f_{a\cup b,\delta_{\omega_1}} = a$, $V_\gamma(a \cup b|\mu_0) = \mathbb{E}_{\hat{p}}[u(a)] = 6$ and $V_\gamma(a \cup b|\bar{\mu}) = \gamma 6 + (1-\gamma)\frac{11}{2} - \gamma\phi(\frac{1}{2})$. Hence, $V_\gamma(a \cup b|\mu_0) \geq V_\gamma(a \cup b|\mu)$, for any $\gamma \in [0, 1]$. The situation changes if the prescription under ignorance is $f_{a\cup b,\hat{p}} = b$ and the prescription in state ω_2 is $f_{a\cup b,\delta_{\omega_2}} = a$, and u is such that $u(5, 5) > u(6, 1)$. Assuming that $\mathbb{E}_{\hat{p}}[u(a)] \geq \mathbb{E}_{\hat{p}}[u(b)]$, then

$$V_\gamma(a \cup b|\mu_0) = \frac{\gamma}{2}\left(u(6,1) + u(6,5)\right) + \frac{(1-\gamma)}{2}\left(u(5,1) + u(5,5)\right) - \\ - \gamma\phi\left(\frac{1}{2}\left(u(6,1) + u(6,5) - u(5,1) - u(5,5)\right)\right)$$

and

$$V_\gamma(a \cup b|\bar{\mu}) = \gamma\left(\frac{1}{2}u(6,5) + \frac{1}{2}u(5,5)\right) + (1-\gamma)\left(\frac{1}{2}u(6,5) + u(6,1)\right) \\ - \phi\left(\frac{1}{2}(u(5,5) - u(6,1))\right)$$

which is consistent with $V_\gamma(a \cup b|\mu_0) \geq V_\gamma(a \cup b|\bar{\mu})$, a preference for ignorance and a contextual choice of b (0% of the subjects in Dana et al. (2007) and 17% of the subjects in Grossman and van der Weele (2016) selected b under ignorance).

The last result of this section illustrates a single-crossing property of the I model:

Proposition 3. *If there is information avoidance for F at $\gamma' < 1$ and at $\gamma = 1$, then there is information avoidance for F at all $\gamma'' \in [\gamma', 1]$.*

If information is too costly, decreasing the probability of behaving according to the prescription does not make information more valuable.

4.2 Demand for "beliefs"

Differently from alternative models of identity protection (e.g. Bénabou and Tirole, 2011; Grossman and van der Weele, 2016), in the present model the beliefs do not enter the utility directly, hence have no hedonic value. However, the I model accommodates a "demand for beliefs" that implies avoidance or acquisition of *non-instrumental* information. I consider information to be non-instrumental in the following sense: suppose that there is a "payoff-dominant action" in F, i.e. an $f^* \in F$ such that $\mathbb{E}_p[u(f^*)] \geq \mathbb{E}_p[u(g)]$ for all $p \in \hat{p} \cup supp\ \mu$. For example, if $u(f^*(\omega)) \geq u(g(\omega))$ for all $\omega \in \Omega$ and all $g \in F$, f^* is a payoff-dominant action in F. When $\gamma = 1$, so that the second-period choice coincides with the payoff-maximizing action, the individual will select f^* from F for *any* posterior as well as for the prior. Therefore, the material value of μ for F is equal to the one of μ_0, since $W_1(\mu, F) = W_1(\mu_0, F) = \mathbb{E}_{\hat{p}}[u(f^*)]$. However,

$$V_1(F|\mu) > V_1(F|\mu_0) \iff \mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})] < \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\mu(q)$$

Indeed, $V_1(F|\mu) = \mathbb{E}_{\hat{p}}[u(f^*)] - \phi\left(\mathbb{E}_{\hat{p}}[u(f^*)] - \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\mu(q)\right)$ and $V_1(F|\mu_0) = \mathbb{E}_{\hat{p}}[u(f^*)] - \phi\left(\mathbb{E}_{\hat{p}}[u(f^*)] - \mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})]\right)$. The opposite inequality will determine avoidance of "non-instrumental" information. A religious person may prefer not knowing if a life-saving drug contains "forbidden" ingredients, if she expects to learn that it is probably the case. Non-instrumental information is valuable (or costly) because it changes the prescription although the choice from the menu is unchanged. This property is consistent with self-serving information acquisition or rejection, which is often rationalized by motivated beliefs (Bénabou, 2015). Beliefs are costly or valuable because of the actions and prescriptions they induce, but they can be changed only indirectly through information acquisition. Thus, acquiring non-instrumental information represents a demand for beliefs that is functional to reduce the cost of deviating from the prescriptions.

4.3 Applications: information avoidance

Blind altruism. Consider three actions, a large donation d, an average donation m or no donation n to a charity of unknown quality (either high ω_h, average ω_m or low ω_l). The payoffs (in utils) are reported in Table 2. A large (average) donation is optimal if

	ω_h	ω_m	ω_l
d	9	0	0
m	4	4	0
n	2	2	2

Table 2: Actions and payoffs

the charity is of high (average) quality and no-donation is optimal 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 absolutely prescribes to make a large donation d, $f_{d\cup m\cup n} = d$. A large donation d also maximizes the ex ante individual's material preference, a form of "warm-glow" giving (Andreoni, 1990). However, information may reveal that the charity is of low quality, so that the optimal action given information is to not donate n, a violation of the prescription. Given γ, the value of $d \cup m \cup n$ under ignorance is $V_\gamma(d \cup m \cup n|\mu_0) = \gamma\frac{9}{3} + (1-\gamma)\frac{9}{3} - \gamma\phi(\frac{9}{3} - \frac{9}{3}) = 3$. The value of $d \cup m \cup n$ if information is acquired is $V_\gamma(d \cup m \cup n|\mu) = \gamma\left(\frac{15}{3}\right) + (1-\gamma)3 - \gamma\phi\,(2)$ equal to $V_\gamma(d \cup m \cup n|\mu) = \gamma 5 + (1-\gamma)3 - \gamma\phi(2)$, and information avoidance is optimal if $2 < \phi(2)$, whenever $\gamma > 0$. If $\gamma = 0$, the individual is indifferent between acquiring information or not since there is an absolute prescription. 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). Differently from the moral wiggle room example, here the avoidance of information is not *motivated* by a desire to maintain an altruistic self-image while acting selfishly. The individual is truly altruistic, but information can be costly if there is a chance that the prescription will be violated.

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 in Oster et al. (2013) 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 gender identity. Getting a medical test threatens the image that *real men* are "independent, self-reliant, strong, robust and tough" (Courtenay, 2000). The I model offers a precise description of the mechanism linking a (stereotypical) masculine identity to the avoidance of health-related information. According to the I 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 (payoff-maximizing behavior) will be to consume less red meat or take a leave from work, behaviors that threaten the (stereotypical) masculine identity. As discussed after Example 1, however, non-monotonicity in the value and cost of information has important implications. If information aligns the prescription and the payoff-maximizing action, information becomes valuable. Therefore, the appeal of routine health screenings can be increased, for example, by making their ability to diagnose dangerous diseases salient.

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), who model information avoidance 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.[19] This model can elegantly rationalize wilfull ignorance for a smoker, but not for someone who decides whether to start smoking and can acquire information before deciding. For example, adolescents underestimating the risk associated with smoking "light" cigarettes may prefer to avoid learning the actual risk (Kropp and Halpern-Felsher, 2004). The I 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.[20] If the identity prescribes smoking, the I model can rationalize the choice to start smoking and the contextual information avoidance concerning its adverse effect. Learning that the probability of developing smoke-related diseases is larger than expected can suggest that not smoking is superior to the prescriptive behavior of start smoking, hence threatening the identity. If the material value of learning the true probability of developing smoking-related diseases 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 represents 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 I 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 (call it w), will make the prescription to restrict women working outside the home (call it n) suboptimal, hence

[19]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.

[20]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).

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. Learning that there is less support for the prescription than previously thought, will shift the prescription itself and potentially change behavior. That is, under ignorance $f_{w\cup n,\hat{p}} = n$, while $f_{w\cup n,p} = w$ for some posterior p when information is acquired.

5 Commitment with and without information arrival

In the previous section, I show that information avoidance arises as a means to reduce the cost of deviating from the prescriptions. This section considers an alternative strategy: strategic commitment. The individual may decide to eliminate actions that threaten the identity even if these are payoff-maximizing. This happens if an action determines a cost to the identity that is higher than its material value. Commitment to a state-by-state dominated action can also be optimal, thus generating an identity-driven form of "destructive" behavior (e.g. self-mutilation).

5.1 Violating the prescriptions and commitment

The next preliminary result establishes conditions to observe a preference for a prescriptively suboptimal menu. I define the benefit-cost ratio between F and G for a given $\nu \in \{\mu, \mu_0\}$:

$$\lambda^{\nu}_{F,G} = \frac{W_\gamma(\nu, F) - W_\gamma(\nu, G)}{I(\nu, F) - I(\nu, G)}$$

if $I(\nu, F) \neq I(\nu, G)$ and $\lambda^{\nu}_{F,G} = 0$ otherwise. Note that $\lambda^{\nu}_{F,G} = \gamma$ if $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\nu(q) = \int_{\Delta\Omega} \mathbb{E}_q[u(f_{G,q})]d\nu(q)$ and $I(\nu, F) \neq I(\nu, G)$.

Theorem 2. *For all $F, G \in \mathcal{A}$ such that $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\nu(q) \geq \int_{\Delta\Omega} \mathbb{E}_q[u(f_{G,q})]d\nu(q)$:*

1. *If $\gamma > 0$, $I(\nu, F) > I(\nu, G)$ and $\gamma\lambda \geq \lambda^{\nu}_{F,G}$ for some $\lambda \in \partial\phi(I(\nu, G))$ then $V_\gamma(G|\nu) \geq V_\gamma(F|\nu)$.*

2. *If $W_\gamma(\nu, G) > W_\gamma(\nu, F)$ and $\gamma\lambda \leq \lambda^{\nu}_{F,G}$ for some $\lambda \in \partial\phi(I(\nu, G))$ then $V_\gamma(G|\nu) \geq V_\gamma(F|\nu)$ and $\gamma > 0$.*

Consider first the case $\nu = \mu_0$. The initial condition becomes $\mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})] \geq \mathbb{E}_{\hat{p}}[u(f_{G,\hat{p}})]$, so that F is prescriptively superior to G. In case 1., the cost of violating the prescriptions in F is higher than that in G, $I(\mu_0, F) > I(\mu_0, G)$. By the initial assumption, it follows that $\max_{f \in F} \mathbb{E}_{\hat{p}}[u(f)] > \max_{f \in G} \mathbb{E}_{\hat{p}}[u(f)]$. Hence F is superior to G from both the prescriptive and the material viewpoint. However, if the marginal cost of deviating from the prescription at G is high enough, the larger cost of violating the prescriptions in F may overcome its higher material value. So, it is optimal to choose G. In point 2., F is prescriptively superior to G, but G has a strictly greater material value. Therefore, G has a higher cost of identity threats than F, $I(\mu_0, G) > I(\mu_0, F)$. However, if the marginal cost of ν at G is "small enough", G is preferred to F.

For the case $\nu = \mu$, the initial condition tells that F is, on average, prescriptively superior to G. In point 1., the average cost of information in F is higher than that in G, $I(\mu, F) > I(\mu, G)$. By the initial assumption, it follows that $\int_{\Delta\Omega} \max_{f \in F} \mathbb{E}_q[u(f)] d\mu(q) > \int_{\Delta\Omega} \max_{f \in G} \mathbb{E}_q[u(f)] d\mu(q)$. 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 cost of information in F may overcome its higher material value. So, it is optimal to choose G. In point 2., F is prescriptively superior to G, but G has a strictly greater material value. Therefore, G has a higher average cost of information than F, $I(\mu, G) > I(\mu, F)$. However, if the marginal cost of information at G is "small enough", G is preferred to F.

Theorem 2 provides sufficient conditions to observe a preference for commitment. For example, suppose that $F = f \cup g$ and $G = g$ and f, then the condition of Theorem 2 becomes $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{f \cup g, q})] d\nu(q) \geq \mathbb{E}_{\hat{p}}[u(g)]$ and it is possible that $V_\gamma(g) \geq V_\gamma(f \cup g|\nu)$, so that commitment to g is superior to having the flexibility of $f \cup g$.

Commitment to payoff-inferior actions. The next result is derived from Theorem 2 and concerns a preference for commitment to a payoff-inferior action. For example, a purely altruistic individual deciding to eschew a situation where she can behave altruistically (e.g. meeting a fundraiser). Suppose that an action g is state-by-state dominated by an action f. It follows that $V_\gamma(f) \geq V_\gamma(g)$, since $V_\gamma(f) = \mathbb{E}_{\hat{p}}[u(f)]$. However, if commitment to f is not feasible, an identity-caring individual may prefer commitment to g to the flexibility of $f \cup g$.

Corollary 1. *Suppose that $u(f(\omega)) \geq u(g(\omega))$ for all $\omega \in \Omega$. If $\gamma > 0$, $I(f \cup g|\nu) > 0$ and $\gamma\lambda \geq \lambda^{\nu}_{f\cup g,g}$ for some $\lambda \in \phi(0)$, the I model is consistent with $V_\gamma(g) \geq V_\gamma(f \cup g|\nu)$.*

Indeed, the condition $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{f\cup g,q})]d\nu(q) \geq \mathbb{E}_{\hat{p}}[u(g)]$ is always satisfied if $f(\omega) \succcurlyeq g(\omega)$ for all $\omega \in \Omega$. Then, if $I(\nu, f \cup g) > I(\nu, g) = 0$ and $\gamma\lambda \geq \lambda^{\nu}_{f\cup g,g}$ for some $\lambda \in \phi(I(\nu, G)) = \phi(0)$, point 1 of Theorem 2 implies that $V_\gamma(g) \geq V_\gamma(f \cup g|\nu)$. Ideally, the individual would commit to f, but if this is unfeasible, committing to the suboptimal g can be strictly preferred to having the flexibility of $f \cup g$. As a consequence, providing the possibility to commit to the prescriptive optimal actions f potentially improves welfare for an identity-concerned individual. Corollary 1 has applications to excess entry and gender-specific sorting into competitive environments that is often rationalized by overconfidence (see Section 5.2). Corollary 1 also rationalizes the findings of DellaVigna et al. (2012) and Andreoni et al. (2017), in which individuals intentionally avoid meeting a fundraiser, hence eschewing social dilemmas even if they are purely altruistic. Consider the choice faced by the individual: meeting the fundraiser and then decide to make a donation d or not n (e.g. with the payoffs of Table 2). Although the prescription is to donate d, entering the social dilemma creates a trade-off and the cost of deviating from the prescription can be so high that avoiding the fundraiser becomes optimal, leading the individual to commit to not donate. This is true under ignorance $\nu = \mu_0$ and under information arrival $\nu = \mu$.

With the payoffs of Table 2 and assuming that d is the absolute prescription, $V_\gamma(n) = 2$ and $V_\gamma(d \cup n|\bar{\mu}) = \gamma\frac{13}{3} + (1-\gamma)3 - \gamma\phi\left(\frac{4}{3}\right)$. 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 d in states ω_m, 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.

Commitment to prescriptions and information acquisition. In the previous sections, I provided conditions to observe commitment to menus that are not prescriptively optimal. Indeed, in Theorem 2, G is weakly prescriptively inferior to F and in

Corollary 1, g is also weakly prescriptively inferior to f. Here, I focus on the conditions under which commitment to the prescriptively optimal menu occurs (either under ignorance and under information acquisition). Consider the actions and payoffs of Table 2. Suppose that d is again the absolute prescription in $d \cup m \cup n$ and in $d \cup m$. Then, $V_\gamma(d \cup m|\bar{\mu}) = \gamma \frac{13}{3} + (1-\gamma)3 - \gamma\phi\left(\frac{4}{3}\right)$ and $V_\gamma(d \cup m \cup n|\bar{\mu}) = \gamma 5 + (1-\gamma)3 - \gamma\phi(2)$. If $\frac{13}{3} - 5 \geq \phi\left(\frac{4}{3}\right) - \phi(2)$, commitment to $d \cup m$ is preferred to $d \cup m \cup n$. Intuitively, the additional material value of having more options is offsets by a larger cost of information. By eliminating the "costly" action n, the individual is better off. The next proposition provides condition to observe a preference for commitment for $\nu \in \{\mu_0, \mu\}$:

Proposition 4. *If $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})]d\nu(q) \geq \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F\cup G,q})]d\nu(q)$ and $\lambda \geq 1$ for some $\lambda \in \partial\phi(I(\nu, F))$ then, for any $\gamma \in [0,1]$, $V_\gamma(F|\nu) \geq V_\gamma(F \cup G|\nu)$.*

When G does not add prescriptive value to F and the cost of deviating from the prescription is steep enough at F, it is optimal to commit to F rather than having the flexibility of $F \cup G$. If $\nu = \mu_0$, the initial condition becomes $\mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})] \geq \mathbb{E}_{\hat{p}}[u(f_{F\cup G,\hat{p}})]$, meaning that G does not add an ex ante prescriptive value to F.

The next result provides simpler conditions to observe commitment to the prescriptively optimal action when it is also the payoff-maximizing action.

Corollary 2. *If $f_{f\cup g,\hat{p}} = f$ and $\mathbb{E}_{\hat{p}}[u(f)] \geq \mathbb{E}_{\hat{p}}[u(g)]$ then there is information avoidance for $f \cup g$ if and only if $V_\gamma(f) > V_\gamma(f \cup g|\mu)$, for all $\gamma \in (0,1]$.*

If information is not valuable for $f \cup g$, commitment to f is strictly preferred to having $f \cup g$ because g can only increase the cost of receiving information. The results in Proposition 4 and Corollary 2 provide a unified rationalization to behaviors observed across choice domains in which identity matters (see Section 5.2). Moreover, Proposition 4 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 may be the unique way to reduce the cost of information.

5.2 Applications: commitment

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).[21] 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.[22] In this case, information about the true state of the world will arrive (inexorably) after the delivery. 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 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[23] (Rudman and Mescher, 2013; Vandello, Hettinger, Bosson, and Siddiqi, 2013). The present model provides a simple explanation. Suppose that the worker wants to preserve a (traditional) masculine and breadwinner identity that absolutely prescribes to delegate d (see, for example, Williams et al., 2013; Rudman and Mescher, 2013; Vandello et al., 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 Proposition 4 or Corollary 2).

Married women and labor market participation. Married women tend to par-

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

[22] 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.

[23] 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"

ticipate to the labor market less than unmarried women. The explanation provided in Bertrand, Kamenica, and Pan (2015) relies on (traditional) gender identity, which prescribes that wives should earn less than their husbands. Indeed, they find that the likelihood that a wife participates in the labor force is inversely proportional to the probability that she will earn more than her husband. The I model can accommodate the empirical evidence in Bertrand et al. (2015) and provides a precise channel through which gender identity affects women's participation in the labor force. Suppose that a married woman has to decide between a high-profile job h and a low-profile job l. Uncertainty pertains, for example, to monthly (relative) income. By selecting the low-profile job l, her earnings are inferior to those of the husband, while the high-profile job h could generate earnings higher than those of the husband. Participating in the labor force allows her to decide between the high-profile job h and the low-profile job l, $E = h \cup l$. Stay at home is represented by the occupation m (e.g. homemaker). A wife unwilling to violate the prescription of earning less than her husband can strictly prefer staying at home m rather than joining the labor force, $V_\gamma(m) > V_\gamma(h \cup l|\mu)$ if, for example, $V_\gamma(m) = V_\gamma(l) \geq V_\gamma(h)$. The mechanism predicted by the I model is the following: if the absolute prescription is to earn less than her husband (hence the prescribed action is to stay home $f_{m \cup l \cup h} = m$), accepting a high-profile job (e.g. if $\gamma = 1$) and reverse the breadwinner role within the household, is costly. Therefore, she may strictly prefer to commit. The empirical fact in Bertrand et al. (2015) that a higher probability of earning more than their husbands implies lower participation in the labor force for wives supports the assumption that information is processed rationally. Wives respond to the information provided by the labor market and act consequentially, some of them by committing.

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 that 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 an 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 absolutely prescribes to enter the market. The I model is consistent with $V_\gamma(e) > V_\gamma(e \cup n|\mu)$, if the conditions of Proposition 4 or Corollary 2 are satisfied. Having n available makes information too costly with respect to 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. Moreover, the model predicts a preference for committing to e even if the entrepreneur knows ex ante that n would be a superior choice. Indeed, by Corollary 1, when commitment to n is not feasible, it may be optimal to commit to an inferior action (in this case e).

On the other side, if boldness does not absolutely prescribe to enter the market, but only conditionally on good news about future demand, information is less costly. Suppose that the ex ante optimal action is n and e is the prescription for $\hat{p}$. Suppose that e becomes the optimal action when a posterior places a sufficiently high probability to state ω_h. In this case, information acquisition becomes valuable as it aligns the prescription and the payoff-maximizing action.

Gender competitiveness. Women are more likely than 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 key role in shaping sorting into competitive environments. The I model provides an identity-based explanation. Suppose that a male faces a real effort task where he can choose to exert either high effort h or low effort l. Uncertainty pertains to his ability (relative to his competitors). Information about ability is often acquired through trial sessions before the task 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 4 or Corollary 2 are consistent with $V_\gamma(h) > V_\gamma(h \cup l|\mu)$. 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 I model. Even if the man knows ex ante that l is superior to h, but commitment to l is not feasible, Corollary 1 implies that he can strictly prefer to enter the competitive environment as a way to avoid the cost of identity threats.

Identity and human capital accumulation. The potential preference for commitment predicted by the I 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 F of potential future occupations, while a high level of human capital corresponds to a menu $G = F \cup w$ (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 $b \in F$, for example, the high-wage occupation w is perceived as a "white job".[24] The result of Proposition 4 is consistent with $V_\gamma(F|\mu) > V_\gamma(F \cup w|\mu)$. Identity protection implies that the cost of w, in terms of identity threats, 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). Note that this conclusion

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

does not require a public signalling role of behavior. Even if $\gamma = 1$, so that choices in the second period coincide with the payoff-maximizing actions, the cognitive dissonance of possibly choosing a white job can be too costly.

Commitment without information acquisition: costly exit in dictator games

The articular case of Theorem 2 in which $\nu = \mu_0$ provides sufficient conditions to observe a preference for commitment under ignorance. For example, suppose that $F = f \cup g$ and $G = g$ and f, then the condition of Theorem 2 with $\nu = \mu_0$ becomes $\mathbb{E}_{\hat{p}}[u(f_{f\cup g,\hat{p}})] \geq \mathbb{E}_{\hat{p}}[u(g)]$ and it is possible that $V_\gamma(g) \geq V_\gamma(f \cup g|\mu_0)$, so that commitment to g is superior to having the flexibility of $f \cup g$.

Such a preference can rationalize the experimental evidence of Dana et al. (2006). In a laboratory experiment, they modified the classical dictator game by allowing the dictators to exit the game before it is played, and without informing the recipient. In case she opts out, the dictator receives \$9 and the recipient receives \$0. A third of the experimental subjects opted out. Such behavior is inconsistent with both purely altruistic and purely selfish preferences. In the former case, the available allocation $(9, 1)$ should be strictly preferred. It the latter case the allocation $(10, 0)$ should be preferred to opting out. Rather, people seem to have a preference for avoiding a moral trade-off. The U model rationalizes the preference for opting out. Suppose that $F = \{(x, 10-x), x \in \{0, \ldots, 10\}\}$ and the prescription is $(x, y)_F = (5, 5)$. Then, $V_\gamma(F|\mu_0) = \gamma \max_{(x,y)\in F} u(x, y) + (1-\gamma)u(5, 5) - \gamma\phi\left(\max_{(x,y)\in F} u(x, y) - u(5, 5)\right)$. If the payoff-maximizing choice is $(10, 0)$, then $V_\gamma(F|\mu_0) = \gamma u(10, 0) + (1-\gamma)u(5, 5) - \gamma\phi(u(10, 0) - u(5, 5))$. Opting out implies commitment to $V_\gamma((9, 0)) = u(9, 0)$. In the extreme case of $u(x, y) = x$, $V_\gamma(F|\mu_0) = \gamma 10 + (1-\gamma)5 - \gamma\phi(10-5)$ and $u(9, 0) = 9$. Therefore, $V_\gamma((9, 0)) \geq V_\gamma(F|\mu_0)$ whenever $9 \geq \gamma 10 + (1-\gamma)5 - \gamma\phi(10-5)$. This is true if $4 \geq \gamma(5 - \phi(5))$ which is always satisfied if $5 \leq \phi(5)$. If $5 > \phi(5)$, the inequality is satisfied if $\gamma \leq \frac{4}{5-\phi(5)}$. Consider the extreme case of $\gamma = 1$, commitment to $u(9, 0)$ is preferred to F if $10 - \phi(10-5) \leq 9$, equivalently if $1 \leq \phi(5)$. If the individual surely follows the prescription, opting out is valuable since the final payoff 9 is larger than the payoff 5 of opting in and following the prescription. If she surely takes the payoff-maximizing action with payoff 10, opting out is valuable if the cost of deviating from the prescription $\phi(5)$ is sufficiently high. Differently from Dillenberger and Sadowski (2012), where resisting the temptation to behave selfishly is costly and leads to opts out, here it is either the costly deviation from the prescriptions or the lower payoff deriving from the prescription that drives a desire for commitment.

6 Optimal disclosure to an identity-caring individual

In this section, I investigate the welfare-maximizing provision of information to an identity-caring individual. I consider a sender, a receiver and an experiment $\mu \in \Gamma(\hat{p})$ called information policy. The sender, who decides the information policy, 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. In this section, *I assume that there is an absolute prescription*. Therefore, the cost of information is increasing in the Blackwell's order (see Proposition 1). A more informative experiment has a dual effect on the receiver's welfare: it has a higher material value, but it is more costly in terms of identity threats. Formally:

Proposition 5. *For all $\mu, \nu \in \Gamma(\hat{p})$:*

(a). For all $F \in \mathcal{A}$ and all $\gamma \in [0,1]$: If $\nu \trianglerighteq \mu$ and if $\lambda \geq 1$ for some $\lambda \in \partial\phi(I(\mu, F))$, then $V_\gamma(F|\mu) \geq V_\gamma(F|\nu)$.

(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 I model predicts that information is not always acquired (Th. 1) and that more information can be welfare decreasing (Prop. 5).

I will consider the following problem faced by the sender: to find the information policy that maximizes the welfare of the receiver. For fixed[25] $\gamma \in (0, 1]$ and fixed $F \in \mathcal{A}$:

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

The following result characterizes the solutions to Problem P:

Theorem 3. *For any $\gamma > 0$, 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, then the information policy $\mu^*_F \in \Gamma(\hat{p})$ solves Problem P.*

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.

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

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

The optimal attention strategy when $\gamma > 0$ and there is an absolute prescription $f_F \in F$ is the solution to Problem P and it is characterized in Theorem 3.

7 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 that the value of an action only depends on its

[25] For $\gamma = 0$ the problem is trivial due to the existence of an absolute prescription in F.

material payoffs, and it is not a function of non-standard elements. Identity here enters the utility of an individual only through the prescriptions. Given the focus of the paper on information acquisition given identity, I do not consider how identity is chosen. After all, information acquisition has consequences in the short-term while choosing an identity is more of a long-term process. 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 costly 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 complements Bénabou and Tirole (2011), in that there is no identity uncertainty. As argued above, imperfect memory of certain "identities" such as gender, 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 I model values actions according to their expected utility and does not suffer from temptation. A related model is that of Konow (2000), where the individual experiences cognitive dissonance by holding beliefs and taking actions different from the ideal. The cognitive dissonance is reduced by directly modifying beliefs and actions. In the present model, beliefs and behavior can be modified only through information acquisition.

The rationality assumptions that characterize the present models also differentiate it from the numerous alternative rationalizations of information aversion or avoidance (see Hertwig and Engel, 2016; Golman et al., 2017; Sunstein, 2020, for recent treatments). The proposed "behavioral" rationalizations of information avoidance or aversion include: 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 et al. (2017) list 13 non-strategic motivations for information avoidance. 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 behaviors. An interesting rationalization of information avoidance comes from motivated beliefs. 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). The I model assumes rationality in information processing and no role for motivated beliefs. Therefore, the current model provides an alternative rationalization of information avoidance with Bayesian information processing and belief-independent utility. This means that further empirical analysis might be required to better understand the source of information avoidance. 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 its 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 spirit of the rational inattention of Sims (2003), where the individual decides what information to pay attention to subject to a cognitive cost. Both Oliveira et al. (2017) and Pennesi (2015) assume a preference for flexibility, because the cost of information is menu-independent. The present model belongs to the ex ante approach since the objects of choice are menus of actions.

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 the present model difficult. In any case, both models (Epstein, 2006; Epstein et al., 2008) assume Set-Betweenness, but the present model can violate it. Most importantly, in the present model information processing is Bayesian, and the preference for commitment derives from a desire to protect identity. In a 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 an individual who wants to preserve her self-image may restrict her future options so as to remain “ignorant".

8 Conclusion and future research

In this paper, I propose a theory of choice for identity concerned individuals, and I applied it to information acquisition. Three main results are predicted by the theory: information avoidance can be optimal, the cost and value of information is not necessarily monotone in the quality of information, the individual can display a demand for beliefs. Lastly, a potential preference for commitment predicted by the theory, has applications in various domains.

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). By endogenizing the possibility to manipulate “identity-salience" or the prior $\hat{p}$, the model can describe self-serving choices of "identity". First, if there are two competing identities with different prescriptions, the individual can invoke the most convenient identity. As argued in Atkin et al. (2021), identity salience responds to group status and group salience, as well as to the material cost of adopting an identity. If salience can be manipulated, the "optimal" identity minimizes the cost of deviating from the prescriptions and contextually maximizes the material value of the prescriptions. Suppose that $U(\cdot|q, F, s)$, where s is the salience of an identity. If identity-salience can be manipulated after information arrives and the realized posterior is q, the individual can invoke the most convenient identity associated with the prescription $f_{F,q} = \text{argmax}_{g \in F} \max_{s \in S} U(g|q, F, s)$. If identity-salience can only be manipulated before information arrives, the "optimal" identity is selected to maximize $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})] d\mu(q)$. Similarly, by managing the prior $\hat{p}$, the prescription also changes so that conveniently "selecting" a prior results in a higher utility (see, for example, Di Tella, Perez-Truglia, Babino, and Sigman, 2015).

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 appendix, I introduce canonical costs of information (Oliveira et al., 2017) and prove Proposition 1. Moreover, I prove that a common specification of the cost of information such as average entropy reduction is a particular case of $\phi(I)$.

Definition 3. *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 $\nu \trianglerighteq \mu$ imply $c(\nu) \geq c(\mu)$.*

(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.

Proof of Proposition 1. For a given $F \in \mathcal{A}$, $\phi(I(\cdot, F))$ is always greater than zero. To see this, take any F, then $\max_{g\in F} \mathbb{E}_q[u(g)] - \mathbb{E}_q[u(f_{F,q})] \geq 0$ for any $q \in \Delta\Omega$. Integrating with respect to μ preserves non-negativity, and the monotonicity of ϕ imply $\phi(I(\mu, F)) \geq 0$. By convexity of ϕ:

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

proving Convexity. If $\mathbb{E}_q[u(f_{F,q})]$ is concave in q, $-\mathbb{E}_q[u(f_{F,q})]$ is convex in q. Hence, $q \mapsto \max_{f\in F} \mathbb{E}_q[u(f)] - \mathbb{E}_q[u(f_{F,q})]$ is convex in q. If $\nu \trianglerighteq \mu$,

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

proving Blackwell's Monotonicity.

Let define $k_F = \phi\left(\max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] - \mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})]\right) \geq 0$, then

$$
\phi(I(\mu_0, F)) - k_F = \phi\left(\max_{g\in F} \mathbb{E}_{\hat{p}}[u(g)] - \mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})]\right) - k_F = 0
$$

showing that $\phi(I(\cdot,F)) - k_F$ is normalized. If there exists an absolute prescription in F, and $\nu \trianglerighteq \mu$,

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

where the first equality follows from $\hat{p} = \int_{\Delta\Omega} q d\mu(q)$, hence proving Blackwell's Monotonicity. By convexity of ϕ:

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

proving Convexity. Let define $\eta_F = \phi\left(\max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)] - \mathbb{E}_{\hat{p}}[u(f_F)]\right) \geq 0$. Then

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

proving Normalization.

□

As argued in Section 3, the entropic costs of information arises in the model under special conditions on the decision problem and the utility u.

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

$$F = \left\{c \in X^\Omega : c > 0, \sum_{\omega\in\Omega}\pi_\omega c_\omega = W\right\}, \quad \text{for some } W > 0 \tag{2}$$

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

Fact 1. If F in Eq. (2), $u(x) = \ln(x)$ and there is an absolute prescription c_F such that $c_F = \frac{\hat{p}(\omega)}{\pi_\omega}$ then

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

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

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

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

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

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

□

B Proofs

I will occasionally use the following notation $\sigma_F(q) \triangleq \max_{f \in F} \mathbb{E}_q[u(f)]$.

Proof of Proposition 2. For point 1, the "only if" direction straightforward. For the "if" direction, suppose that for an $F \in \mathcal{A}$, $I(\mu, F) = 0$. Since $\max_{f \in F} \mathbb{E}_q[u(f)] - \mathbb{E}_q[u(f_{F,q})] \geq 0$ for all $q \in \Delta\Omega$, the equality $I(\mu, F) = \int_{\Delta\Omega} \max_{f \in F} \mathbb{E}_q[u(f)] - \mathbb{E}_q[u(f_{F,q})] d\mu(q) = 0$ implies $\max_{f \in F} \mathbb{E}_q[u(f)] - \mathbb{E}_q[u(f_{F,q})] = 0$ μ-a.s. Therefore, there is $f_{F,q}$ such that $f_{F,q} \in \operatorname{argmax}_{g \in F} \mathbb{E}_q[u(g)]$ μ-a.s..

For point 2, if $\mathbb{E}_{\hat{p}}[u(f)] \geq \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})] d\mu(q)$ for some $f \in F$ and $I(\mu, F) = 0$, then

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

But, in general, $\max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)] \leq \int_{\Delta\Omega}\max_{f\in F}\mathbb{E}_q[u(f)]d\mu(q)$, hence

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

The last equality allows to use the same argument of De Lara and Gossner (2020), therefore, there is $f \in F$ such that $f \in \operatorname{argmax}_{g\in F}\mathbb{E}_{\hat{p}}[u(g)]$ and $f \in \operatorname{argmax}_{g\in F}\mathbb{E}_q[u(g)]$ μ-a.s. □

Proof of Theorem 1. Suppose that $\gamma = 0$, then the second condition implies $0 > \lambda_F$. That is $0 > \frac{\int_{\Delta\Omega}\mathbb{E}_q[u(f_{F,q})]d\mu(q)-\mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})]}{I(\mu,F)-I(\mu_0,F)}$. Since the denominator is strictly positive, the inequality holds only if the numerator is strictly negative, hence $V_0(F|\mu_0) = \mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})] > \int_{\Delta\Omega}\mathbb{E}_q[u(f_{F,q})]d\mu(q) = V_0(F|\mu)$, that is information avoidance for F. Consider $\gamma > 0$, for some $\lambda \in \partial\phi(I(\mu_0, F))$, $\gamma\lambda > \lambda_F$. Then:

$$\begin{aligned}
\phi(I(\mu,F)) - \phi(I(\mu_0,F)) \geq& \lambda(I(\mu,F) - I(\mu_0,F)) \\
>& \frac{\lambda_F}{\gamma}\left(I(\mu,F) - I(\mu_0,F)\right) \\
=& \frac{1}{\gamma}\left(W_\gamma(F,\mu) - W_\gamma(F,\mu_0)\right)
\end{aligned}$$

where the first inequality follows from the definition of subdifferential. The second inequality from the fact that $\gamma\lambda > \lambda_F$ and $I(\mu,F) > I(\mu,F)$. The equality from the definition of λ_F. Rearranging, implies information avoidance for F.

□

Proof of Proposition 3. If there is information avoidance for F at $\gamma' < 1$, it means that the function $T(\gamma) = \gamma I(\mu,F) + \int_{\Delta\Omega}\mathbb{E}_q[u(f_{F,q})]d\mu(q) - \gamma\phi\left(I(\mu,F)\right) - \gamma(\max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)] - \mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})]) + \mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})] + \gamma\phi(\max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)] - \mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})])$ is such that $T(\gamma') < 0$. Deriving T with respect to γ' implies $I(\mu,F) - \phi\left(I(\mu,F)\right) - \max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)] + \mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})] + \phi(\max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)] - \mathbb{E}_{\hat{p}}[u(f_{F,\hat{p}})])$ which is strictly smaller than zero by the fact that information is avoided for F at $\gamma = 1$. Therefore, the function T is strictly negative at γ' and it is strictly decreasing, therefore, it remains strictly negative as γ'' varies in $[\gamma', 1]$. □

Proof of Theorem 2. For point 1., by the second assumption, F and G cannot be singletons. Consider $\lambda^{\nu}_{F,G} = \frac{W_\gamma(\nu,F)-W_\gamma(\nu,G)}{I(\nu,F)-I(\nu,G)}$. Since $\gamma > 0$, note that

$$\begin{aligned}
&0 < I(\nu, F) - I(\nu, G) = \\
&\gamma \left(\int_{\Delta\Omega} \sigma_F(q) d\nu(q) - \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})] d\nu(q) \right) - \gamma \left(\int_{\Delta\Omega} \sigma_G(q) d\nu(q) + \int_{\Delta\Omega} \mathbb{E}_p[u(f_{G,q})] d\nu(q) \right) \\
\leq &\gamma \left(\int_{\Delta\Omega} \sigma_F(q) d\nu(q) - \int_{\Delta\Omega} \mathbb{E}_q[u(f_{G,q})] d\nu(q) \right) - \gamma \left(\int_{\Delta\Omega} \sigma_G(q) d\nu(q) + \int_{\Delta\Omega} \mathbb{E}_q[u(f_{G,q})] d\nu(q) \right) \\
= &\gamma \left(\int_{\Delta\Omega} \sigma_F(q) d\nu(q) - \int_{\Delta\Omega} \sigma_G(q) d\nu(q) \right)
\end{aligned}$$

Therefore, $\int_{\Delta\Omega} \sigma_F(q) d\nu(q) > \int_{\Delta\Omega} \sigma_G(q) d\nu(q)$, that implies $\lambda^{\nu}_{F,G} > 0$. If $\lambda \geq \lambda^{\nu}_{F,G}/\gamma$ for some $\lambda \in \partial\phi(I(\nu, G))$, then

$$\begin{aligned}
\phi(I(\nu, F)) - \phi(I(\nu, G)) \geq &\lambda \left(I(\nu, F) - I(\nu, G) \right) \\
\geq &\frac{\lambda^{\nu}_{F,G}}{\gamma} \left(I(\nu, F) - I(\nu, G) \right) \\
= &\frac{1}{\gamma} \left(W_\gamma(\nu, F) - W_\gamma(\nu, G) \right)
\end{aligned}$$

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

For point 2., by the two initial assumptions F and G cannot both singletons. Moreover, $W_\gamma(\nu, G) - W_\gamma(\nu, F) > 0$ and $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})] d\nu(q) \geq \int_{\Delta\Omega} \mathbb{E}_q[u(f_{G,q})] d\nu(q)$ imply $\gamma > 0$. Notice that $0 < W_\gamma(\nu, G) - W_\gamma(\nu, F) = I(\nu, G) + \int_{\Delta\Omega} \mathbb{E}_q[u(f_{G,q})] d\nu(q) - I(\nu, F) - \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})] d\nu(p) \leq I(\nu, G) - I(\nu, F)$, hence $I(\nu, G) > I(\nu, F)$. Hence $\lambda^{\nu}_{F,G} > 0$ If $\gamma\lambda \leq \lambda^{\nu}_{F,G}$ for some $\lambda \in \partial\phi(I(\nu, G))$, then

$$\begin{aligned}
\phi(I(\nu, F)) - \phi(I(\nu, G)) \geq &\lambda \left(I(\nu, F) - I(\nu, G) \right) \\
\geq &\frac{\lambda^{\nu}_{F,G}}{\gamma} \left(I(\nu, F) - I(\nu, G) \right) \\
= &\frac{1}{\gamma} \left(W_\gamma(\nu, F) - W_\gamma(\nu, G) \right)
\end{aligned}$$

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

Proof of Proposition 4. If $\gamma = 0$, the initial condition $V_0(F|\nu) = \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})] d\nu(q) \geq \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F\cup G,q})] d\nu(q) = V_0(F \cup G|\nu)$ implies the result. Consider $\gamma > 0$. If $\lambda \geq 1$ for some $\lambda \in \partial\phi(I(\nu, F))$, then

$$
\begin{aligned}
&\phi(I(\nu, F \cup G)) - \phi(I(\nu, F)) \geq \lambda(I(\nu, F \cup G) - I(\nu, F)) \\
=&\lambda \left(\int_{\Delta\Omega} \sigma_{F\cup G}(q) d\nu(q) - \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F\cup G,q})] d\nu(q) - \int_{\Delta\Omega} \sigma_F(q) d\nu(q) + \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})] d\nu(q) \right) \\
\geq&\lambda \left(\int_{\Delta\Omega} \sigma_{F\cup G}(q) d\nu(q) - \int_{\Delta\Omega} \sigma_F(q) d\nu(q) \right) \\
\geq& \int_{\Delta\Omega} \max_{f \in F\cup G} \mathbb{E}_q[u(f)] d\nu(q) - \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)] d\nu(q)
\end{aligned}
$$

where the first inequality follows from the definition of subgradient, the second equality by the condition $\int_{\Delta\Omega} \mathbb{E}_q[u(f_{F\cup G,q})] d\nu(q) \leq \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})] d\nu(q)$, and the last inequality from the fact that $\lambda \geq 1$ and the term in the parenthesis is positive. Rearranging and multiplying both sides by $\gamma > 0$ gives $\gamma \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)] d\nu(q) - \gamma\phi(I(\nu, F)) \geq \gamma \int_{\Delta\Omega} \max_{f\in F\cup G} \mathbb{E}_q[u(f)] d\nu(q) - \gamma\phi(I(\nu, F \cup G))$. Lastly, adding $(1-\gamma) \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F\cup G,q})] d\nu(q)$ to the right-hand side and $(1-\gamma) \int_{\Delta\Omega} \mathbb{E}_q[u(f_{F,q})] d\nu(q)$ to the left-hand side preserves the inequality and gives $V_\gamma(F|\nu) \geq V_\gamma(F \cup G|\nu)$.

□

Proof of Corollary 2. If information is not acquired a $f \cup g$, then $V_\gamma(f \cup g|\mu) < V_\gamma(f \cup g|\mu_0)$, meaning $V_\gamma(f \cup g|\mu) = \gamma \int_{\Delta\Omega} \max_{h\in f\cup g} \mathbb{E}_q[u(h)] + (1-\gamma)\mathbb{E}_q[u(f_{f\cup g,q})] d\mu(q) - \gamma\phi(I(\mu, f \cup g)) < \gamma \max_{h\in f\cup g} \mathbb{E}_{\hat{p}}[u(h)] + (1-\gamma)\mathbb{E}_{\hat{p}}[u(f_{f\cup g,\hat{p}})] - \gamma\phi(\max_{h\in f\cup g} \mathbb{E}_{\hat{p}}[u(h)] - \mathbb{E}_{\hat{p}}[u(f_{f\cup g,\hat{p}})]) = \mathbb{E}_{\hat{p}}[u(f)] = V_\gamma(f)$, where the last equality follows from the assumptions that $f_{f\cup g,\hat{p}} = f$ and $\mathbb{E}_{\hat{p}}[u(f)] \geq \mathbb{E}_{\hat{p}}[u(g)]$. Being strict, the inequality forces $\gamma \in (0, 1]$. For the opposite implication, the conditions imply that $V_\gamma(f \cup g|\mu_0) = \gamma \max_{h\in f\cup g} \mathbb{E}_{\hat{p}}[u(h)] + (1-\gamma)\mathbb{E}_{\hat{p}}[u(f_{f\cup g,\hat{p}})] - \gamma\phi(\max_{h\in f\cup g} \mathbb{E}_{\hat{p}}[u(h)] - \mathbb{E}_{\hat{p}}[u(f_{f\cup g,\hat{p}})]) = V_\gamma(f) = \mathbb{E}_{\hat{p}}[u(f)] > V_\gamma(f \cup g|\mu)$, hence, information is not acquired at $f \cup g$ for all $\gamma \in (0, 1]$. □

Proof of Proposition 5. For a., since $\nu \trianglerighteq \mu$ either $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)] d\nu(q) = \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)] d\mu(q)$ or $\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)] d\nu(q) > \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)] d\mu(q)$. In the first case, $V_\gamma(F|\nu) = V_\gamma(F|\mu)$ for all $\gamma \in [0, 1]$. Consider the strict inequality. If $\gamma = 0$, then again $V_0(F|\nu) = V_0(F|\mu)$ and the result follows. Take $\gamma > 0$, if $\lambda \in \partial\phi(I(\mu, F))$ and $\lambda \geq 1$,

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

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, multiplying by γ both sides and adding $(1-\gamma)\mathbb{E}_{\hat{p}}[u(f_F)]$ to both sides gives $V_\gamma(F|\nu) \geq V_\gamma(F|\mu)$.

For b., since ϕ is convex, the ratio $\frac{\phi(x)-\phi(y)}{x-y}$ is weakly increasing in x for fixed y. By the condition $\nu \trianglerighteq \mu$, it follows that $I(\nu, F) \geq I(\mu, F)$. If information μ is not acquired for F, $W_\gamma(\mu, F) - \gamma\phi(I(F|\mu)) < W_\gamma(\mu_0, F) - \gamma\phi(I(\mu_0, F))$ (which implies $\gamma > 0$ and

$\int_{\Delta\Omega} \sigma_F(q)d\mu(q) > \sigma_F(\hat{p})$), that is

$$\begin{aligned}\gamma \int_{\Delta\Omega} \sigma_F(q)d\mu(q) + (1-\gamma)\mathbb{E}_{\hat{p}}[u(f_F)] - \gamma\phi\left(\int_{\Delta\Omega} \sigma_F(q)d\mu(q) - \mathbb{E}_{\hat{p}}[u(f_F)]\right) < \\ <\gamma\sigma_F(\hat{p}) + (1-\gamma)\mathbb{E}_{\hat{p}}[u(f_F)] - \gamma\phi\left(\sigma_F(\hat{p}) - \mathbb{E}_{\hat{p}}[u(f_F)]\right)\end{aligned} \tag{3}$$

that is equivalent to

$$\begin{aligned}1 <& \frac{\phi\left(\int_{\Delta\Omega} \sigma_F(q)d\mu(q) - \mathbb{E}_{\hat{p}}[u(f_F)]\right) - \phi\left(\sigma_F(\hat{p}) - \mathbb{E}_{\hat{p}}[u(f_F)]\right)}{\int_{\Delta\Omega} \sigma_F(q)d\mu(q) - \mathbb{E}_{\hat{p}}[u(f_F)] - \sigma_F(\hat{p}) + \mathbb{E}_{\hat{p}}[u(f_F)]} \\ \leq& \frac{\phi\left(\int_{\Delta\Omega} \sigma_F(q)d\nu(q) - \mathbb{E}_{\hat{p}}[u(f_F)]\right) - \phi\left(\sigma_F(\hat{p}) - \mathbb{E}_{\hat{p}}[u(f_F)]\right)}{\int_{\Delta\Omega} \sigma_F(q)d\nu(q) - \mathbb{E}_{\hat{p}}[u(f_F)] - \sigma_F(\hat{p}) + \mathbb{E}_{\hat{p}}[u(f_F)]}\end{aligned}$$

hence $V_\gamma(F|\nu) < V_\gamma(F|\mu_0)$, so information ν is not acquired for F as well. □

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:

$$\begin{aligned}V_\gamma(\mu_1, F) - V_\gamma(\mu, F) =& W_\gamma(\mu_1, F) - \gamma\phi(I(\mu_1, F)) - W_\gamma(\mu, F) + \gamma\phi(I(\mu, F)) \\ \geq& W_\gamma(\mu_1, F) - W_\gamma(\mu, F) + \gamma(I(\mu, F) - I(\mu_1, F)) = 0\end{aligned}$$

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

Fact 2. *For each $t \in [I(\mu_0, F), I(\bar{\mu}, F)]$, there is $\mu_t \in \Gamma(\hat{p})$ such that $I(\mu_t, F) = t$.*

Proof of Fact 2. Let $\mu_t = \frac{t-I(\mu_0,F)}{I(\bar{\mu},F)-I(\mu_0,F)}\bar{\mu}+(1-\frac{t-I(\mu_0,F)}{I(\bar{\mu},F)-I(\mu_0,F)})\mu_0$. Then $I(\mu_t, F) = \int_{\Delta\Omega} \sigma_F(q)d\mu_t(q) - \mathbb{E}_{\hat{p}}[u(f_F)] = t$. □

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

□

References

Abrams, D. and M. A. Hogg (2006). *Social identifications: A social psychology of intergroup relations and group processes.* Routledge.

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. (1990). Impure altruism and donations to public goods: A theory of warm-glow giving. *The Economic Journal 100*(401), 464–477.

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.

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.

Bertrand, M., E. Kamenica, and J. Pan (2015). Gender identity and relative income within households. *The Quarterly Journal of Economics 130*(2), 571–614.

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., M. Callen, B. Ferman, S. Gulzar, A. Hasanain, and N. Yuchtman (2019). Political Identity: Experimental Evidence on Anti-Americanism in Pakistan. *Journal of the European Economic Association 18*(5), 2532–2560.

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.

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.

Di Tella, R., R. Perez-Truglia, A. Babino, and M. Sigman (2015, November). Conveniently upset: Avoiding altruism by distorting beliefs about others' altruism. *American Economic Review 105*(11), 3416–42.

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.

Festinger, L. (1962). Cognitive dissonance. *Scientific American 207*(4), 93–106.

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

Glick, P., J. L. Berdahl, and N. M. Alonso (2018). Development and validation of the masculinity contest culture scale. *Journal of Social Issues 74*(3), 449–476.

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.

Higgins, E. T. (1989). Self-discrepancy theory: What patterns of self-beliefs cause people to suffer? In *Advances in experimental social psychology*, Volume 22, pp. 93–136. Elsevier.

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.

Konow, J. (2000). Fair shares: Accountability and cognitive dissonance in allocation decisions. *American economic review 90*(4), 1072–1091.

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.

Kropp, R. Y. and B. L. Halpern-Felsher (2004). Adolescents' beliefs about the risks involved in smoking "light" cigarettes. *Pediatrics 114*(4), e445–e451.

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

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.

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.

Stets, J. E. and R. T. Serpe (2016). *New Directions in Identity Theory and Research.* Oxford University Press.

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.