ISSN 2279-9362

Collegio Carlo Alberto

Identity and the cost of information

Daniele Pennesi

No. 610
June 2020

Carlo Alberto Notebooks

www.carloalberto.org/research/working-papers

Identity and the cost of information

Daniele Pennesi*

February 2020

Abstract

This paper proposes a theory of rational information avoidance based solely on the desire to protect identity, a person's sense of self. Identity-caring individuals follow the prescriptions: behaviors that are appropriate given the identity. Information is costly because it can recommend a behavior that violates the prescriptions. The theory rationalizes information avoidance observed in different domains without resorting to cognitive biases, motivated reasoning or time-inconsistency. By analyzing revealed preferences over menus of actions, it is possible to disentangle identity protection from the many alternative rationalizations of information avoidance. The key behavioral restriction, a desire to have fewer options, supplies an identity-based explanation for excess entry in competitive environments, low participation of married women to the labor force, under-use of flexible work arrangements, and ethnic-specific differences in human capital accumulation.

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

1 Introduction

In a series of influential works, Akerlof and Kranton (2000, 2005, 2010) introduced identity, a person's sense of self, into economic analysis. When deciding, identity-caring individuals take into account the material consequences of an action as well as the image, public or personal, that such action convey. The most direct and unequivocal way to (self-) signal identity is to follow the *prescriptions*: behaviors that are appropriate given the identity.[1] For instance, a right-wing politician can follow the prescription to vote against softening immigration laws as a means to protect her identity. A man willing to protect his self-image of masculine,

*ESOMAS Department, University of Torino. Email: daniele.pennesi@unito.it

[1] According to Akerlof and Kranton (2000) "Prescriptions indicate the behavior appropriate for people in different social categories in different situations. The prescriptions may also describe an ideal for each category in terms of physical characteristics and other attributes....agents follow prescriptions, for the most part, to maintain their self-concepts...violating the prescriptions evokes anxiety and discomfort in oneself...".

breadwinner and committed worker, can follow the prescription to "put work first" (Glick et al., 2018). In a situation where learning is possible, the desire to protect identity generates a tension between information and ignorance. On the one hand, information is valuable because it is instrumental: it can be used to make better decisions. On the other hand, information is cognitively "costly": it can recommend a behavior that violates the prescriptions.[2] If an individual wants to protect her identity by following the prescriptions, information loses its instrumental value, but its cost persists, hence leading to information avoidance. A masculine, breadwinner and committed worker eager to follow the prescription of "putting work first", will deem detrimental any information on his health. Learning that he is sick and needs a break is costly because it contradicts the prescription and has no instrumental value because he will continue to work. As a consequence, he may prefer to avoid health-related information altogether.[3]

In this paper, I consider the rich implications of the above intuition and develop a theory of information avoidance based solely on the desire to protect identity. I show that identity protection offers a unified explanation to many instances of information avoidance in absence of cognitive constraints, motivated beliefs, anticipatory feelings, disappointment aversion or sophisticated time-inconsistency.[4] In my theory the individual is fully rational and sophisticated: she is able to anticipate her desire to preserve identity, she processes information in a Bayesian way and she values actions according to expected utility.[5] Identity protection rationally explains willful ignorance in social dilemmas (Dana et al., 2007; Grossman and van der Weele, 2016; Kandul and Ritov, 2017), avoidance of information concerning the effects of harmful consumption (Viscusi, 1990), willful ignorance of work incentives (Huck et al., 2018) and avoidance

[2]From a psychological perspective, the theory of cognitive dissonance developed in (Festinger, 1962) offers a foundation for the cost of information in this paper. In his words "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." In this sense, information is costly because it induces dissonance (see also Golman et al., 2017).

[3]There is large empirical evidence showing that men are less likely than women to get routine health checks (e.g. Courtenay, 2000), see also the discussion in Section 4.1.

[4]These are all proposed rationalizations of information avoidance or aversion. In particular, Sims (2003) for cognitive constraints, Köszegi (2006) for motivated beliefs, Caplin and Leahy (2001) and Brunnermeier and Parker (2005) for anticipatory feelings, Andries and Haddad (2019) for disappointment aversion and Carrillo and Mariotti (2000) and Bénabou and Tirole (2011) for sophisticated time-inconsistency. In their review, Golman et al. (2017) list 13 non-strategic motivations for information avoidance.

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

of routine health screenings (Courtenay, 2000). For instance, an individual who thinks of herself as an altruist can avoid freely available information about the efficacy of her donations because her identity prescribes to donate. Information is costly as it may suggest to violate the prescription, for example, if she learns that her favorite charity is less reliable than previously thought. The theory identifies a precise mechanism through which identity protection leads to information avoidance: the desire to *behave* in conformity with the prescriptions. Information is costly only when it recommends an *action* that violates the prescriptions. If the information modifies the beliefs, but the recommended action is still consistent with the prescriptions, information is costless. Implying that identity is protected through *actions* rather than beliefs and "dangerous" thoughts may or may not become taboos. A feature that discriminates the current approach from alternative approaches such as motivated beliefs or anticipatory feelings. Indeed, in the latter cases, even the information changing beliefs but not the recommended action can be costly.

A further contribution of the paper, that leads to a different set of results, is the identification of the behavioral restrictions that characterize the desire to protect identity. Hence allowing to discriminate identity protection from the numerous alternative rationalizations of information avoidance. I consider a dynamic choice problem where an individual first decides whether to become informed or not (I will allow for multiple sources of information in Section 8). Conditional on remaining ignorant, the individual selects a menu of actions whose payoffs depend on an unknown state of nature and then an action from the menu. Conditional on becoming informed, the individual selects a menu, receives information and selects an action from the menu. In the latter case, a key prediction of the present theory is that an individual anticipating a desire to preserve identity can strictly prefer to have fewer options. Indeed, the likelihood that information makes the prescribed behavior suboptimal is lower in smaller menus. Therefore, beyond avoidance, an identity-caring individual uses commitment as a means to reduce the cost of information, a prediction relevant in situations where information is physically unavoidable (see Section 5). In these cases, restricting future options represents the only way to reduce the cost of information. As a consequence, the theory provides an identity-based rationalization for why some entrepreneurs overinvest in certain markets, why married women participate less than unmarried women to the labor force , why males select into competition more often than females (but only in patriarchal societies), why fathers under-use flexible work arrangements, and why different ethnic groups make specific choices concerning human capital accumulation. An entrepreneur who sees herself as *bold*, for example, can rationally decide to overinvest in a market as a means to protect her identity, even knowing that information about the level of future demand will arrive. This prediction does not require overconfidence (Camerer and Lovallo, 1999) or time-inconsistency (Brocas and Carrillo, 2004) and contradicts the classical result linking delayed information arrival to underinvestment (Arrow and Fisher, 1974). A married woman unwilling to violate the prescription of earning less than her husband, can strictly prefer to commit herself to stay at home, rather than participating to the labor force (Bertrand et al., 2015). To protect his self-image of bold and skilled, a man will embrace competition regardless of his true ability (Niederle and Vesterlund, 2007; Gneezy et al., 2009). A man about to become a father and willing to protect his image of breadwinner and committed worker will avoid asking a flexible work arrangement, hence committing to delegate childcare (Williams et al., 2013). Similarly, a black student willing to protect her identity may prefer to accumulate a low human capital, hence committing to occupations that are consistent with her identity and avoiding "white jobs" (Austen-Smith and Fryer Jr, 2005).

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 (henceforth, I will refer to this class of models as the *standard models*). 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 et al., 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 value commitment: the desire to preserve her identity.

The structure of the model allows a sharp comparative static analysis. I introduce and characterize through restrictions on the parameters of the model, a comparative notion of having a "greater ability to follow the prescription" (GAFP). Behaviorally, an individual 1 has a GAFP than 2 if the choices of 1 are always more aligned with the prescriptions of her identity than the choices of 2. Parametrically, it means that they share the same identity, but the cost of information for 1 is always smaller than the cost of information for 2.

Lastly, I extend the model to allow for flexible information acquisition. In this extension, there are multiple sources of information that can be selected by the individual. I show that an information source is optimal if and only if it is consistent with the prescriptions, hence inducing an extreme form of selective attention.

The paper is structured as follows. Section 2 introduces the framework. Section 2.1 motivates the approach through examples arguing that both Preference for Flexibility and Set-Betweenness are unable to capture a preference for identity protection. Section 3 introduces the model and Sections 4 and 5 contain its applications. Section 6 provides the behavioral characterization of the model and Section 6.1 develops the comparative analysis. Section 7 studies the design of incentive-compatible information policies. Lastly, Section 8 provides an extension of the model allowing for flexible information acquisition. The Appendix A contains all the proofs.

2 Framework

There is a finite set Ω of states of the world representing uncertainty. An act is a function from states to consequences $f : \Omega \to X$, where X is a compact, convex subset of a normed vector space. The set of all acts is denoted by $\mathcal{F}$. As usual, constant acts $f(\omega) = x$ for all $\omega \in \Omega$, are identified with the element of X they deliver. I denote by $\mathcal{A}$ the set of all nonempty and finite subsets of $\mathcal{F}$, the menus. The mixture $+$ of two acts, denoted by $\alpha f + (1-\alpha)g$ for all $\alpha \in [0, 1]$, is performed state-wise, $(\alpha f + (1-\alpha)g)(\omega) = \alpha f(\omega) + (1-\alpha)g(\omega)$. The mixture of two menus is the menu containing all mixtures of the elements in F and G, $\alpha F + (1-\alpha)G = \{\alpha f + (1-\alpha)g : f \in F,\ g \in G\}$, for all $\alpha \in [0, 1]$. To simplify notation, I will write f for the singletons $\{f\}$ and $f \cup g$ for the doubletons $\{f, g\}$. Given a topological space M, I denote by ΔM the set of probabilities defined on the Borel σ-algebra of M and I endow ΔM with the weak* topology. Given a probability distribution $\hat{p} \in \Delta\Omega$, a probability $\nu \in \Delta\Delta\Omega$ satisfying the martingale property of Bayesian updating $\hat{p} = \int_{\Delta\Omega} p d\nu(p)$ is interpreted as an *experiment* consistent with $\hat{p}$ (see, for example, Gollier, 2004). The experiment μ specifies the probability of deriving a posterior by Bayesian updating of the prior $\hat{p}$.[6] The set of all experiments consistent with a prior $\hat{p}$ is denoted by $\Gamma(\hat{p})$, i.e. $\Gamma(\hat{p}) = \{\nu \in \Delta\Delta\Omega : \int_{\Delta\Omega} p d\nu(p) = \hat{p}\}$.

For most part of the paper, the primitives are the two binary relations $(\succsim_\upsilon, \succsim_\iota)$ defined on $\mathcal{A}$ and representing preferences conditional on avoiding information and conditional on acquiring information from the "source" ι, respectively.[7] The interpretation of $F \succsim_\iota G$ is "*the menu F is preferred to G conditional on acquiring information from the source ι, but prior to information arrival*". Whereas $F \succsim_\upsilon G$ means "*the menu F is preferred to G conditional on avoiding information*". The precise definition of ι is, in a sense, irrelevant. Indeed, even if ι represents an experiment and its informational content is known to the external observer, what matters in determining the cost of information is the *perceived* informational content of ι. In the representation, it corresponds to a *subjective* experiment consistent with the prior, that may or may not coincide with the information content of ι. I show that the behavioral conditions allow to uniquely identify the subjective experiment.

The choice process evolves as follows (see Fig. 1): conditional on avoiding information, the individual selects a menu and an then action from the menu. Conditional on acquiring information, the individual selects a menu, observes a signal, forms a posterior and then selects an action from the menu. Lastly, in both cases, uncertainty resolves and the payoff is delivered.

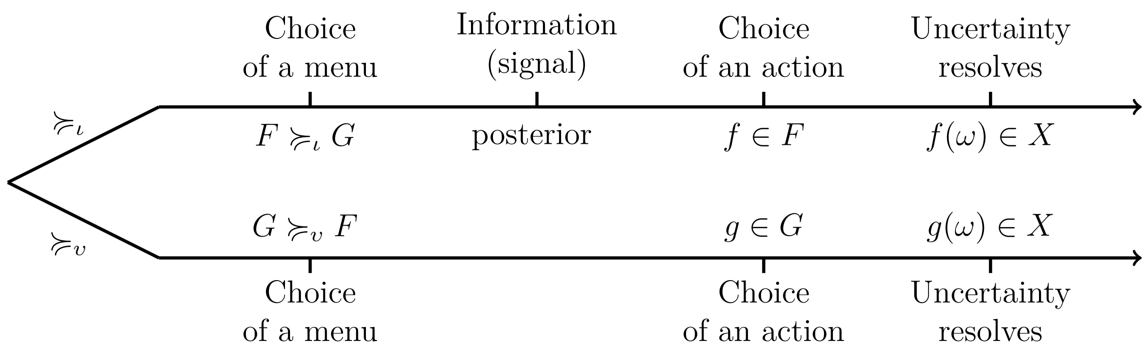

Figure 1: Dynamics of the choice process

[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.

[7] I will consider flexible information acquisition in Section 8, where the primitive is a family $\{\succsim_\theta\}_{\Theta\cup\upsilon}$ of conditional preferences, where $\theta \in \Theta \cup \upsilon$ represents different sources of information.

2.1 A motivating example

In this section, I provide an example motivating why the desire to protect identity, in the presence of arriving information, induces violations of two well-known axioms: Preference for Flexibility and Set-Betweenness. The next axiom is Preference for Flexibility:

Axiom (Preference for Flexibility.). *For all menus $F, G \in \mathcal{A}$, $F \succcurlyeq_\iota G$ whenever $G \subseteq F$*

The natural interpretation is the following: an individual expecting to receive information prefers larger menus because she will use the information to optimize her choice *from* the menu. The Preference for Flexibility axiom is satisfied in the standard models (Arrow and Fisher, 1974; Dillenberger et al., 2014) as well as in models of costly information acquisition (Sims, 2003; Oliveira et al., 2017; Pennesi, 2015).

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).[8] 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 pertains how stressful is to look after his newborn, it can be either high ω_h or low ω_l.[9] Information about the true state of the world will arrive (inexorably) after the delivery. Models satisfying the Preference for Flexibility axiom predict that the worker will *never* strictly prefer to maintain the current work schedule G prior to information arrival. Since information will be used to optimize the choice between c and d, flexibility is always weakly valuable, hence $F = c \cup d$ is preferred (according to $\succcurlyeq_\iota$) to $G = d$.

However, suppose that the worker wants to preserve a (traditional) masculine and breadwinner identity. In this case, child-caring represents a threat to the identity, hence the prescribed behavior is to delegate d (see, for example, Williams et al., 2013; Rudman and Mescher, 2013; Vandello et al., 2013). Anticipating that he will protect his self-image by delegating childcare (to follow the prescription), the flexibility entailed by the flexible work arrangement $c \cup d$ is not more valuable than the current work schedule d. However, maintaining the possibility to look

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

[9]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.

after his newborn c can generate a costly threat to his identity. This happens if the information makes the worker realize that child caring c is superior to delegating d, hence suggesting a violation of the prescription. To avoid such a situation, commitment to $G = d$ can be strictly preferred to $F = c \cup d$, a violation of the Preference for Flexibility axiom:

$$d \succ_\iota c \cup d \tag{1}$$

Empirical evidence consistent with the preference in (1) 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 flexibility, especially for family caregiving, is seen as an impermissible lack of commitment, if not a feminine behavior[10] (Rudman and Mescher, 2013; Vandello et al., 2013). Note that a symmetric argument would apply for an *involved father* (Marks and Palkovitz, 2004), whose prescription is c. The involved father can strictly prefer to commit to c, for example, by changing job.

A possible weakening of the Preference for Flexibility axiom, consistent with the preference $d \succ_\iota c \cup d$, is the Set-Betweenness axiom of Gul and Pesendorfer (2001):

Axiom (Set-Betweenness.)**.** *For all $F, G \in \mathcal{A}$, if $F \succcurlyeq_\iota G$ then $F \succcurlyeq_\iota F \cup G \succcurlyeq_\iota G$.*

In a setting with objective uncertainty and information acquisition, Set-Betweenness is assumed in Epstein (2006). In Epstein's interpretation, a sophisticated individual who anticipates her future under or overreaction to information (non-Bayesian updating) can prefer to commit. However, assume the existence of a third action l available to the worker (e.g. take a leave). Identity prescribes d better than l, for example, men that take a leave for childcare are "not serious about heir work" (Vandello et al., 2013). The following preferences violate Set-Betweenness:

$$d \succ_\iota c \cup d \sim_\iota d \cup l \succ_\iota c \cup d \cup l \tag{2}$$

The preferences $d \succ_\iota c \cup d$ and $d \succ_\iota d \cup l$ can be rationalized as above. The last preference can also be rationalized by the desire to preserve self-image: the menu $c \cup d \cup l$ is worse than $c \cup d$

[10] 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"

and $d \cup l$ because in $c \cup d \cup l$ it is more likely that the information will contradict the prescribed behavior d. Information is costly if a posterior ranks looking after his newborn c better than delegating d, but also if a different posterior ranks taking a leave l better than delegating d. Hence, in the larger menu $c \cup d \cup l$, there are "more threats" to the identity than in $c \cup d$ or in $d \cup l$, leading to a strict preference for the smaller menus.

Note that the previous rationalization of excess commitment (to work) *does not* require violations of rationality. A fully rational individual can strictly prefer to commit because of a desire to protect his identity. The model introduced in the next section is consistent with the preferences in (1) and (2) and also predicts a preference for information avoidance. Namely, the worker would prefer not to know how stressful is to look after his newborn (a remote possibility). However, in the model information processing is Bayesian and actions are evaluated according to expected utility.

3 Modeling information avoidance

In this section, I introduce the Information Avoidance model for the binary relations $\succsim_\upsilon$ and $\succsim_\iota$.

Definition 1. *An Information Avoidance (IA) representation of a pair of binary relations* $(\succsim_\upsilon, \succsim_\iota)$ *is a tuple* $(u, \hat{p}, \mu, \kappa)$ *where* $u : X \to \mathbb{R}$ *is non-constant and affine,* $\hat{p} \in \Delta\Omega$*,* $\mu \in \Gamma(\hat{p})$*,* $\kappa \geq 0$*, and such that* $\succsim_\upsilon$ *has a representation* $V(\cdot|\upsilon) : \mathcal{A} \to \mathbb{R}$ *given by:*

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

and $\succsim_\iota$ *has a representation* $V(\cdot|\iota) : \mathcal{A} \to \mathbb{R}$ *given by:*

$$V(F|\iota) = \max_{f \in F} \mathbb{E}_{\hat{p}}[u(f)] - \kappa C(\mu, F) \tag{I}$$

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

In the IA model, I assume that the prescriptions are *internalized* (see Akerlof and Kranton, 2000, p. 728). The identity determines a preference among actions represented by the prior $\hat{p}$ and the utility u. In particular, the identity prescribes the action f over g if and only if

$\mathbb{E}_{\hat{p}}[u(f)] \geq \mathbb{E}_{\hat{p}}[u(g)]$. This assumption makes identity identifiable by observing revealed preferences among singleton menus. Moreover, the assumption allows interpreting $V(F|\upsilon)$ in Eq. U as the value of "following the prescription" when facing a menu F: the (expected) utility of selecting from F the best action according to the prescription. Conditional on expecting (but not yet receiving) information, Eq. I implies that the value of a menu F is given by the difference between its "follow-the-prescription value" $\max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$ and the cost of information $\kappa C(\mu, F)$. To clarify why $C(\mu, F)$ represents the cost of information, suppose that the identity prescribes $f \in F$. Since u and $\hat{p}$ represent the the prescription, $\mathbb{E}_{\hat{p}}[u(f)] \geq \max_{g\in F} \mathbb{E}_{\hat{p}}[u(g)]$. Suppose that, for a posterior q with $\mu(q) > 0$, it is true that $\mathbb{E}_q[u(g)] > \mathbb{E}_q[u(f)]$, for some $g \in F$, then $C(\mu, F) > 0$.[11] Information is costly because, with strictly positive probability $\mu(q) > 0$, the individual may learn that the prescribed action f is inferior (in expectation) to an alternative action g. Since different posteriors in the support of μ can rank g better than f, the cost of information $C(\mu, F)$ represents the average cost of identity threats in F when expecting to learn μ and it is always greater or equal than zero.[12]

The IA model rationalizes the preference for commitment displayed in Equation (1) of Section 2.1. If the identity prescribes the action d over c, it is true that $\mathbb{E}_{\hat{p}}[u(d)] = \max_{h\in c\cup d} \mathbb{E}_{\hat{p}}[u(h)]$. It follows that (when $\kappa > 0$) $V(d|\iota) > V(c \cup d|\iota)$, the preference in (1). For all singletons $g \in \mathcal{A}$, the condition $\hat{p} = \int_{\Delta\Omega} q d\mu(q)$ implies $V(g|\iota) = \mathbb{E}_{\hat{p}}[u(g)]$. Therefore, $V(d|\iota) = \mathbb{E}_{\hat{p}}[u(d)] = \max_{h\in c\cup d} \mathbb{E}_{\hat{p}}[u(h)] > \max_{h\in c\cup d} \mathbb{E}_{\hat{p}}[u(h)] - \kappa C(\mu, c \cup d) = V(c \cup d|\iota)$. Adding an option that is not optimal according to the prescription can only decrease the value of the flexible work arrangement $c \cup d$, in case a posterior ranks to look after the baby c better than to delegate d. Moreover, if the worker anticipates that he will follow the prescription d regardless of the arriving information, the latter becomes non-instrumental. Being the information costly, the worker (weakly) prefers to remain ignorant: $V(c\cup d|\upsilon) \geq V(c\cup d|\iota)$, and the inequality is strict if $\kappa > 0$.[13] For the preference in (2), if $d \succ_\upsilon c$ and $d \succ_\upsilon l$, adding the option l to $c \cup d$ does not

[11]Indeed, by the condition $\hat{p} = \int_{\Delta\Omega} q d\mu(q)$, $C(\mu, F) = \int_{\Delta\Omega} [\max_{h\in F} \mathbb{E}_q[u(h)] - \mathbb{E}_q[u(f)]]\, d\mu(q)$. Since for a q with $\mu(q) > 0$ and a $g \in F$, $\mathbb{E}_q[u(g)] > \mathbb{E}_q[u(f)]$, the cost $C(\mu, F) > 0$.

[12]For a given $F \in \mathcal{A}$ and by the condition $\hat{p} = \int_{\Delta\Omega} q d\mu(q)$, the convexity of $q \mapsto \max_{g\in F} \mathbb{E}_q[u(g)]$ implies that $C(\mu, F) = \int_{\Delta\Omega} \max_{g\in F} \mathbb{E}_q[u(g)] d\mu(q) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] \geq \max_{g\in F} \int_{\Delta\Omega} \mathbb{E}_q[u(g)] d\mu(q) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] = 0$.

[13]In the IA model information avoidance derives from the anticipating that the choice from menu will conform the prescription, hence information is non-instrumental. Assuming that information is costly because of identity protection, but information is used instrumentally, is consistent with the following model: $V'(F) = \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_q[u(f)] d\mu(q) - \kappa C(\mu, F)$. This model is observationally equivalent to the Subjective Learning model of Dillenberger et al. (2014), hence it satisfies the Preference for Flexibility axiom. In the model V' information is always weakly preferred to ignorance.

increase the prescription value of $c \cup d$, but it can increase the cost of information.[14]

The term $C(\mu, F)$ corresponds to the *instrumental value of receiving information* μ in the choice problem F (see De Lara and Gossner, 2020; Frankel and Kamenica, 2019). Namely, the difference between the value of F for an individual who anticipates the use of information μ to select the best action from F, and the value of F in the absence of information arrival. Since the instrumental value of information is null if the choice from the menu is trivial, in the IA model information is not costly when facing a singleton, $C(\mu, f) = 0$ for all $f \in \mathcal{F}$. In my interpretation, the information cannot challenge identity in a singleton. The next result follows from simple duality arguments (see De Lara and Gossner, 2020) and characterizes all the cases in which the cost of information is null in the IA model:

Proposition 1 (Actions speak louder than words)**.** *For all $F \in \mathcal{A}$, $C(\mu, F) = 0$ if and only if there exists $f \in \operatorname{argmax}_{g \in F} \mathbb{E}_{\hat{p}}[u(g)]$ such that $f \in \operatorname{argmax}_{g \in F} \mathbb{E}_p[u(g)]$, μ-a.s.*

Actions speak louder than words in protecting identity. Indeed, information is costly only when it recommends, with strictly positive probability, an action different from the prescribed behavior. Information that modifies the posterior beliefs, but not the recommended action, is costless. Therefore, in the IA model prohibited thoughts induced by information are not necessarily *taboos*. For the masculine worker, learning that childcare is indeed less stressful than expected. is not necessarily costly. It becomes a cost when, given information, looking after his newborn becomes the optimal action, hence superior to the prescription.

The result of Proposition 1 has two consequences: first, it allows determining when an individual has a strictly positive willingness to pay to avoid information (e.g. 46% of the subjects in Grossman and van der Weele, 2016). Second, it allows discriminating between the IA and models with belief-dependent utilities (such as Köszegi, 2006). Information changing the posterior but not the recommended action can be costly when the beliefs enter the utility directly. In this case, the mere learning that looking after his newborn is better than expected may be costly for the worker, even if the prescription to delegate remains optimal given the information. In Section 7, I show that a policy-maker can exploit the result in Proposition 1 to design an information policy that is, at the same time, costless and informative. A similar

[14]Indeed, $V(d \cup l|\iota) = \mathbb{E}_{\hat{p}}[u(d)] - \kappa C(\mu, d \cup l) \geq \mathbb{E}_{\hat{p}}[u(d)] - \kappa C(\mu, c \cup d \cup l) = V(c \cup d \cup l|\iota)$, since the cost function satisfies $C(\mu, c \cup d \cup l) = \int_{\Delta\Omega} \max_{h \in c \cup d \cup l} \mathbb{E}_p[u(h)] d\mu(p) - \mathbb{E}_{\hat{p}}[u(d)] \geq \int_{\Delta\Omega} \max_{h \in d \cup l} \mathbb{E}_p[u(h)] d\mu(p) - \mathbb{E}_{\hat{p}}[u(d)] = C(\mu, d \cup l)$ and the inequality is strict if, for some q with $\mu(q) > 0$, $\mathbb{E}_q[u(l)] > \max_{h \in c \cup d} \mathbb{E}_q[u(h)]$.

result may not hold in a model with a belief-dependent utility.

I conclude the section with two remarks:

Remark 1. Note that the cost of information occurs *before* uncertainty resolves. It may be the case that the prescribed action remains *eventually* superior to the action suggested by the information. For example, suppose that $\mathbb{E}_{\hat{p}}[u(g)] > \mathbb{E}_{\hat{p}}[u(g)]$ and once information is revealed, the realized posterior q is such that $\mathbb{E}_q[u(g)] > \mathbb{E}_q[u(f)]$. However, after uncertainty resolves $u(f(\omega^*)) > u(g(\omega^*))$ where ω^* is the true state of the world. In this case, information is costly even if the prescribed action is eventually optimal.

Remark 2. Consider a menu F containing only constant acts, i.e. $F \subseteq X$: it easy to see that $V(F|\iota) = \max_{x \in F} u(x)$, the pure preference for flexibility of Kreps (1979). This property is motivated by the fact that the desire for commitment can solely derive from a desire to protect identity and not from tastes uncertainty: tastes are fixed and known to the individual. Therefore, when evaluating menus of state-independent actions, flexibility is desirable since information cannot threaten the identity. This feature differentiates my model from models of choice over menus (of lotteries) that allows for tastes uncertainty with or without a preference for flexibility (Dekel et al., 2001; Gul and Pesendorfer, 2001; Sarver, 2008; Dekel et al., 2009).

4 Applications: information avoidance

In this section, I apply the IA model to rationalize various instances of information avoidance observed empirically. The main tool employed in the applications is the next result:

Proposition 2. *If the pair $(\succsim_\upsilon, \succsim_\iota)$ has an IA representation then:*

1. *For all $F \in \mathcal{A}$, $V(F|\iota) \leq V(F|\upsilon)$.*

2. *If $V(F|\upsilon) \geq V(G|\upsilon)$ then $V(F|\upsilon) \geq V(F \cup G|\iota)$*

Point 1 shows that information avoidance is, *ceteris paribus*, always weakly preferred to information acquisition. Moreover, committing to the prescription and remaining ignorant are weakly preferred to become informed and maintain flexibility (point 2). This last property, relating identity preservation, preference for commitment and information avoidance, has a direct application in the choice problems of the following Sections.

A combination of 1 and 2 in Proposition 2 implies the following property: $V(F|\upsilon) \geq V(G|\upsilon)$ then $V(F|\iota) \geq V(F \cup G|\iota)$. The extra-flexibility gained by adding options that are not optimal according to the prescription can be detrimental. In case a posterior ranks an action in G better than the prescription, identity is threatened. This form of preference for commitment represents a means to reduce the cost of information alternative to (physical) avoidance and it is particularly relevant for choice situations in which the physical avoidance of information is unfeasible, as I show in Section 5.

4.1 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 such 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 et al. (2013) show that at-risk individuals, those having one parent with the Huntington disease, avoid undergoing a perfectly informative test. A suitable rationalization of such behavior is optimism maintenance (Brunnermeier and Parker, 2005). However, the same rationalization seems less compelling to explain why men avoid routine medical tests. At-risk individuals in the study of Oster et al. (2013) are likely to know their condition, while it is less plausible to assume a similar awareness for the average man. An alternative explanation relies on identity protection. Getting medical tests threaten the image that *real men* are “independent, self-reliant, strong, robust and tough" (Courtenay, 2000). The IA model offers a precise description of the mechanism linking masculine identity and the avoidance of health-related information. According to the IA model, it is not the mere act of getting tested to threaten masculinity, but the impact that health-related information will have on identity. After a diagnosis of high blood pressure, the recommendation will be to consume less red meat or take a leave from work, a threat to the (stereotypical) masculine identity. In Section 7, I address the problem of how to provide information to an identity-caring individual respecting incentive compatibility.

4.2 Willful ignorance in social dilemmas

Growing empirical evidence shows that individuals often avoid information concerning the effect that their choices have on others (e.g. Dana et al., 2007; Grossman and van der Weele, 2016; Niehaus, 2014; Kandul and Ritov, 2017). Theoretical explanations are based on a combination of imperfect memory and self-control problems (Bénabou and Tirole, 2011) or explicit signaling (Grossman and van der Weele, 2016). The IA model rationalizes willful ignorance in social dilemmas only assuming the desire to preserve self-image. Consider a choice between two actions: donate d or not n to a charity of unknown quality. Suppose that the individual sees herself as an altruist, Proposition 2 shows that she may strictly prefer commitment to donate, rather than receiving information and having the flexibility to donate or not, $V(d|\upsilon) = V(d \cup n|\upsilon) > V(d \cup n|\iota)$. Indeed, if the information makes not-donating optimal, an altruistic individual caring about her identity will donate anyway while knowing it is suboptimal. Anticipating such a situation, she will prefer remaining ignorant (see for example Kandul and Ritov, 2017). Consistent with this conclusion, Niehaus (2014) found that only 3 percent of donors acquire information prior to making a donation. The above inequality also shows that the provision of information about a charity can be costly. Individuals that donate in any case because of their identity can experience a welfare loss if information contradicts the prescription. Along similar lines, the IA model predicts that a consumer identifying himself as *ethical* can prefer to ignore information about the attributes of a product (Ehrich and Irwin, 2005).

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

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

In their experiment, Dana et al. (2007) observed that g is the modal choice when knowing that the state of the world was ω_2. Namely $u(5,5) \geq u(6,1)$ for 74% of the subjects. In their main treatment, the subjects did not know the state of the world but can learn it at no cost. In this case, only 56% of their subjects decide to learn the state of the world. The IA model is consistent with willful ignorance, indeed $V(f \cup g|\upsilon) \geq V(f \cup g|\iota)$. Moreover, *conditional on remaining ignorant* (the moral wiggle room), the modal choice for the subjects in Dana et al. (2007) was f. The IA model can account for an even stronger form of moral wiggle room: namely, a preference for commitment to f, $V(f|\upsilon) > V(f \cup g|\iota)$. Suppose that $\hat{p}(\omega_1) = \frac{1}{2}$ as in Dana et al. (2007), then:

Fact 1. If (expected) information is perfectly revealing, $u(5,5) > u(6,1)$ and $u(6,5) + u(6,1) > u(5,1) + u(5,5)$, then the individual will prefer to commit to f and remain ignorant, $V(f|\upsilon) > V(f \cup g|\iota)$.

4.3 Information avoidance and "self-control"

Information avoidance is also observed in individuals experiencing self-control problems, such as smokers (Viscusi, 1990). In particular, Viscusi (1990) found that smokers overestimated the risk of developing lung cancer even if the available information is virtually free and abundant. In their theory, Carrillo and Mariotti (2000) argue that information avoidance acts as a commitment device in a intra-personal game for a sophisticated individual experiencing self-control problems. 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. 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. The IA model rationalizes information avoidance without requiring self-control problems/time-inconsistency. Suppose that a smoker sees herself as a *light* smoker, for example, she is a school teacher or a physician. There are two actions, continue to smoke moderately (e.g. 5 cigarettes a day) m or smoke more m^+ (e.g. 10 cigarettes a day). Her self-image prescribes to keep smoking moderately. The IA model predicts $V(m|\upsilon) = V(m \cup m^+|\upsilon) \geq V(m \cup m^+|\iota)$. Learning that the probability of developing lung cancer is smaller than expected can suggest that smoking moderately m (the prescribed action) is inferior to smoking more m^+, hence threatening the identity. Anticipating her decision to continue to smoke moderately, the individual will prefer to remain ignorant. Note that if $V(m^+|\iota) \geq V(m|\iota)$ the inequality becomes $V(m^+|\upsilon) = V(m \cup m^+|\upsilon) \geq V(m \cup m^+|\iota)$. So, even an individual who's self-image prescribes to smoke a pack of cigarettes a day (e.g. a rock star) can also prefer to avoid information. A similar argument is consistent with information avoidance of the calorie content of foods (Thunström et al., 2016). For instance, an individual who sees herself as an athlete can prefer to avoid calorie-related information in absence of self-control problems.

4.4 Identity and cultural cognition

Identity preservation as a rationale for information avoidance is particularly suitable to explain the persistence of ideological polarization. Issues such as the nature of climate change, the effect of gun control, the effectiveness of the death penalty, the role of the state in the economy, the economic impact of migrations are well-known to polarize opinions as a function of the political identity (the ideology) despite the abundance of factual information. The *cultural cognition theory* developed by Kahan and Braman (2006) suggests that, for ideology-sensitive issues such as those listed above, the ideology affects cognition directly. Factual information is distorted in a self-serving manner as to maintain consistency between the interpretation of facts and the ideological prescriptions. It is worth noting that the IA model allows for subjective cognition: the information μ is individual-specific and may be a function of the ideology. However, the subjectivity of μ is not the channel through which polarization occurs. Even if factual information is not distorted in a self-serving manner, the IA model predicts a *polarization in actions*. Two individuals with diverse ideologies (e.g. a Democrat and a Republican) will behave differently even after observing the same information ι and even if they do not distort factual information. In the IA model polarization in actions is not determined by distorted beliefs, since posteriors are generated from the prior by Bayesian updating, but by the prescriptions. It follows that without eliciting posterior beliefs and by only observing choices *from* menus, it is not possible to understand the source of polarization. At the same time, if an analyst elicits the prior of the individual and observes choices from menus, an identity-caring individual will appear insensitive to information. However, information is received and processed rationally although it is not employed in the second-period choice. Section 6 shows that, by observing revealed preference among menus, it is possible to identify the IA model, hence solving the mentioned identification issue.

4.5 Work incentives and ignorance

In a laboratory experiment, Huck et al. (2018) study subjects' performance in a real-effort task in which the performance pay is uncertain (either a high or a low piece-rate). They found that one-third of their subjects, given the choice, prefer to remain ignorant about the piece-rate. Huck et al. (2018) rationalize their findings with the optimal expectation model of Brunnermeier and Parker (2005), where ignorance helps workers to maintain motivation. The IA model provides an alternative rationalization: suppose that the subject can choose to exert either high effort h or low effort l. She sees herself as a *tireless worker*, hence her identity prescribes to exert high effort. In this case, $V(h \cup l|\upsilon) \geq V(h \cup l|\iota)$. If information about the piece-rate recommends low effort l, the prescribed action h becomes suboptimal, hence identity is threatened. Anticipating that she will exert high effort to preserve her identity, she may prefer to avoid information. A symmetric argument will predict information avoidance for an individual whose identity prescribes l.

5 Applications: identity, information and commitment

In this section, I consider choice situations in which information is hardly avoidable. As predicted by the IA model (Proposition 2, point 2), the only strategy to reduce the cost of information in these cases is to commit. The following sections contain applications of the previous property to economically relevant decisions.

5.1 Excess entry without overconfidence

Consider the following decision problem: an entrepreneur has to decide whether to enter e a new market or not n. Uncertainty concerns the level of future demand, it can be either high ω_h or low ω_l. Before deciding, she can either make an irreversible investment H or not N. By making a irreversible investment H, she commits to enter the market in the second period $H = e$. By making no initial investment N, the entrepreneur maintains flexibility $N = e \cup n$. Information concerning the level of future demand is *de facto* unavoidable. Models satisfying the Preference for Flexibility axiom predict that the entrepreneur will *never* strictly prefer to make an irreversible investment prior to information arrival. Since information will be used to optimize the choice between e and n, flexibility is weakly valuable. This is the classical result linking under-investment and delayed information arrival due to Arrow and Fisher (1974) (see also Gollier, 2004).

However, suppose that the entrepreneur sees herself as a *bold* (for example, Brocas and Carrillo, 2004) and wants to preserve her self-image. The latter prescribes to enter the market, hence choosing e in the second period. The IA model is consistent with $V(e|\iota) > V(e \cup n|\iota)$. Anticipating that she will follow the prescription anyway, the flexibility entailed by $N = e \cup n$ is not more valuable than committing to $H = e$ (since information is not used instrumentally), but having the action n available increases the cost of information. Note that the previous rationalization of excess entry/overinvestment *does not* require overconfidence (Camerer and Lovallo, 1999) or time-inconsistency (Brocas and Carrillo, 2004). A fully rational entrepreneur can strictly prefer entering a new market because of a desire to protect her identity.

Additional empirical evidence consistent with the IA model is the observed gender differences in competitive inclinations. In general, women are more likely to shy away from competition while men embrace it, even if there is no gender differences in performance (Niederle and Vesterlund, 2007). Interestingly, the opposite result holds in *matrilineal* societies (Gneezy et al., 2009), indicating that also cultural aspects drive sorting into competitive environments. The IA model provides an identity-based explanation for the gender-specific attitude to compete. Men (resp. women) wanting to preserve their image may prefer to compete. Suppose that the subject faces a real task. 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 trials sessions before the real experiment starts. Entering the competitive environment (e.g. a tournament) implies exerting high effort, $F = h$. Entering the non-competitive environment (e.g. a flat-rate payment scheme) leaves the flexibility to exert either high effort or low effort, $G = h \cup l$. Suppose that his self-image prescribes to exert high effort. Therefore, Proposition 2 is consistent with $V(h|\upsilon) > V(h \cup l|\iota)$. If information about the ability recommends a low effort l, the prescribed action h becomes suboptimal, hence identity is threatened. By anticipating that he will exert a high effort to preserve identity, entering the competitive environment can be strictly optimal. An equivalent argument holds for women in a matrilineal society. This type of sorting is often ascribed to overconfidence of men (Niederle and Vesterlund, 2007). However, differently from overconfidence, information processing is rational in the IA model. Although the prescription may be determined by a wrong assumption that men performs better than women, the posterior is formed by Bayes' rule. Hence, a man rationally anticipating that low effort will be optimal can still enter the competitive environment.

5.2 Identity and labor force participation

Married women tend to participate to the labor market less than unmarried women. The explanation provided in Bertrand et al. (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 to the labor force is inversely proportional to the probability that she will earn more than her husband. The IA model can accommodate the empirical evidence in Bertrand et al. (2015) and provides a precise channel through which identity affects the 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 (e.g. a part-time job). Uncertainty pertains, for example, to monthly (relative) income. Information concerning the features of the labor market is difficult to avoid. Moreover, as I argue at the end of the section, the result of Bertrand et al. (2015) supports the conclusion that married women do acquire and process information about the labor market rationally. 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). According to the IA model, 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, $V(m|\upsilon) > V(h \cup l|\iota)$ if, for example, $V(m|\upsilon) = V(l|\upsilon) \geq V(h|\upsilon)$.[15] The mechanism predicted by the IA model is the following: if the prescription is to earn less than her husband (hence the prescribe action is to stay home or to embrace a low-profile job), the possibility to accept

[15]Indeed, $V(l|\upsilon) = V(h|\upsilon)$ implies $V(l|\upsilon) = V(l|\iota) \geq V(l \cup h|\iota)$, hence $V(m|\upsilon) \geq V(l \cup h|\iota)$.

a high-profile job and reverse the breadwinner role within the household, can be detrimental. Although she will protect her identity by earning less than her husband, learning that the high-profile job is superior to the low-profile job is costly. Therefore she may strictly prefer to commit. The fact 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 (by committing).

5.3 Identity and human capital accumulation

The preference for commitment in the IA model can rationalize suboptimal human capital investment decisions in minorities. Suppose that a student belonging to a minority (e.g. a black student) has to decide how much human capital she wants to accumulate, for simplicity, either high or low. A low accumulation of human capital corresponds to a menu G of potential future occupations, while a high level of human capital corresponds to a menu $F = G \cup m$ (where m represents a high-wage occupation such as a lawyer). Uncertainty pertains to the features of the labor market and it is difficult, if not impossible, to avoid receiving information. Suppose that the identity prescribes an occupation $g \in G$, for example, the high-wage occupation m is perceived as a "white job".[16] Then the IA model is consistent with $V(G|\iota) > V(G \cup m|\iota)$. Identity-protection implies that the student will select the occupation g regardless of the features of the labor market, therefore maintaining flexibility may be costly. If after information arrives, the prescribed occupation g is inferior to m, identity is threatened. Hence, the student may strictly prefer commitment to a low level of human capital G (e.g. Austen-Smith and Fryer Jr, 2005). A symmetric argument applies to a student whose identity prescribes to aim at high wage occupations. In this case, committing to m, for example by attending an elite college, may be strictly preferred to having flexibility.

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

6 Behavioral characterization

In this section, I introduce the axioms characterizing the Information Avoidance model (Def. 1). The primitives are the binary relations $\succsim_\upsilon$ and $\succsim_\iota$ defined on $\mathcal{A}$.

Axiom (Consistent Weak Order). *For $\theta \in \{\upsilon, \iota\}$, the binary relation $\succsim_\theta$ is a weak order and there exist $f, g \in \mathcal{F}$ such that $f \succ_\theta g$. Moreover, for all $f, g \in \mathcal{F}$, $f \succsim_\upsilon g \iff f \succsim_\iota g$.*

The only non-standard part of Consistent Weak Order is the ordinal equivalence between $\succsim_\upsilon$ and $\succsim_\iota$ when restricted to singletons. The next continuity axiom requires the introduction of the Hausdorff metric between menus:

$$d_H(F, G) = \max \left\{ \sup_{f \in F} \inf_{g \in G} d(f, g), \sup_{g \in G} \inf_{f \in F} d(f, g) \right\}$$

where $d(f, g)$ is the Euclidean distance between $f \in \mathcal{F}$ and $g \in \mathcal{F}$, when $\mathcal{F}$ is identified with X^Ω.

Axiom (Continuity).

a. *For $\theta \in \{\upsilon, \iota\}$, if $F \succ_\theta G \succ_\theta H$, then there are $\gamma, \gamma' \in (0, 1)$ such that $\gamma F + (1 - \gamma)H \succ_\theta G \succ_\theta \gamma' F + (1 - \gamma')H$.*

b. *For $\theta \in \{\upsilon, \iota\}$, there exist menus $F^*, F_* \in \mathcal{A}$ and $k > 0$ such that for any $F, G \in \mathcal{A}$ and $\gamma \in (0, 1)$ with $d_H(F, G) \leq \gamma/k$, $\gamma F^* + (1 - \gamma)F \succsim_\theta \gamma F_* + (1 - \gamma)G$.*

The first condition is an Archimedean continuity. The second implies, together with the other axioms, that the representation is Lipschitz continuous (see Sarver, 2008). The third axiom is the classical independence axiom: preferences are not reversed after mixing with a common menu.

Axiom (Independence). *For $\theta \in \{\upsilon, \iota\}$ and for all menus $F, G, H \in \mathcal{A}$ and $\alpha \in (0, 1)$,*

$$F \succsim_\theta G \iff \alpha F + (1 - \alpha)H \succsim_\theta \alpha G + (1 - \alpha)H$$

The independence axiom is incompatible with models of costly information acquisition as Oliveira et al. (2017) and Pennesi (2015). Hence it excludes confounding factors in the identification of the cost of information.

I now state a behavioral conditions that is peculiar to the binary relation $\succsim_{\upsilon}$.

Axiom (U-Strategic Rationality)**.** *If $f \succsim_{\upsilon} g$ and $f \in F$, then $F \sim_{\upsilon} F \cup g$.*

The above axiom is similar to Kreps (1979)'s Strategic Rationality. In the absence of information arrival, adding an action that is not optimal according to prescription leaves the individual indifferent, since the choice from the menu F cannot be better (or worse) than the choice from $F \cup g$.

The next axiom bundles the Anscombe-Aumann monotonicity axiom with a weaker form of strategic rationality. If a menu contains an act f that dominates *in any state* an act g, then adding g to the menu leaves the individual indifferent.[17]

Axiom (Dominance)**.** *If $f(\omega) \succsim_{\upsilon} g(\omega)$ for all $\omega \in \Omega$ and $f \in F$, then $f \succsim_{\upsilon} g$ and $F \sim_{\iota} F \cup g$.*

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

The last axiom is the main behavioral restriction characterizing the IA model and it is motivated by the discussion in Section 2.1. It is called Identity Preservation:

Axiom (Identity Preservation)**.** *If $f \succsim_{\upsilon} g$ and $f \in F$, then $F \succsim_{\iota} F \cup g$.*

If an action g is not inferior to an f according to the prescription, adding g to a menu containing f can make the decision-maker strictly worse off. Indeed, g does not add a material value to a menu containing f, but one or more posteriors can rank g better than f, threatening the identity. Anticipating such a situation, the decision-maker can strictly prefer to commit to F.[18] Note that the condition $f \succsim_{\upsilon} g$ and $f \in F$ has distinct implications for the two preferences $\succsim_{\upsilon}$ and $\succsim_{\iota}$: in one case $F \sim_{\upsilon} F \cup g$ and $F \succsim_{\iota} F \cup g$ in the other. Lastly, note

[17] A similar, though weaker condition, is assumed in Dillenberger et al. (2014); Oliveira et al. (2017). Their condition combined with the Preference for Flexibility axiom implies Dominance.

[18] An axiom with a similar structure appears in Sarver (2008) even if his setting and interpretation of the model are different from the current setting. In section 9, I spell out the key differences between the two approaches.

that Dominance restricts the scope of Identity Preservation. Suppose that $f(\omega) \succcurlyeq_\upsilon g(\omega)$ for all $\omega \in \Omega$ and $f \in F$. Then by Dominance, $f \succcurlyeq_\upsilon g$ and Identity Preservation implies $F \succcurlyeq_\iota F \cup g$. However, Dominance again forces $F \sim_\iota F \cup g$. This property implies that a strict preference for commitment can only occur if information is strictly more costly in the larger menu $F \cup g$.

Then next theorem is the main result of this section:

Theorem 1. *A pair of binary relations $(\succcurlyeq_\upsilon, \succcurlyeq_\iota)$ satisfies Weak Order, Continuity, Independence, Dominance, Identity Preservation and $\succcurlyeq_\upsilon$ satisfies U-Strategic Rationality if and only if it has an IA representation $(u, \hat{p}, \mu, \kappa)$.*

6.1 Comparative analysis

In this section, I perform a comparative static analysis aimed at capturing a comparative notion of greater ability to follow the prescriptions. I consider two individuals 1 and 2, each of which has preferences represented by an Information Avoidance model.

Definition 2. *Suppose that $(\succcurlyeq^1_\upsilon, \succcurlyeq^1_\iota)$ and $(\succcurlyeq^2_\upsilon, \succcurlyeq^2_\iota)$ are represented by two IA $(u_j, \hat{p}_j, \mu_j, \kappa_j)$, for $j = 1, 2$. I say that $(\succcurlyeq^2_\upsilon, \succcurlyeq^2_\iota)$ has a* greater ability to follow the prescriptions *(GAFP) than $(\succcurlyeq^1_\upsilon, \succcurlyeq^1_\iota)$ if and only if, for all $F, G \in \mathcal{A}$,*

$$F \succcurlyeq^1_\upsilon G \quad \text{and} \quad F \succcurlyeq^1_\iota G \quad \Longrightarrow \quad F \succcurlyeq^2_\iota G$$

Note that the condition implies that $u_1 \approx u_2$ and $\hat{p}_1 = \hat{p}_2$ that is, $\succcurlyeq^1_\upsilon$ and $\succcurlyeq^2_\upsilon$ are equivalent, meaning that 1 and 2 follows the same prescriptions. Indeed, take $F = f$ and $G = g$, then $f \succcurlyeq^1_\upsilon g$ implies directly $f \succcurlyeq^1_\iota g$ and the comparative definition over singletons reduces to: $f \succcurlyeq^1_\upsilon g$ implies $f \succcurlyeq^2_\upsilon g$. By the uniqueness of the expected utility representation over singletons, $u_1 \approx u_2$ and $\hat{p}_1 = \hat{p}_2$. Therefore, the intuition of 2 having GAFP than 1 is that whenever the preference $\succcurlyeq^1_\iota$ is aligned with the prescription $\succcurlyeq^1_\upsilon$, the preference $\succcurlyeq^2_\iota$ is also aligned with (the common) prescription $\succcurlyeq^2_\upsilon$.

The following result characterizes the parametric restrictions of having a greater ability to follow the prescriptions:

Theorem 2. *Given binary relations $(\succcurlyeq^j_\upsilon, \succcurlyeq^j_\iota)$, $j = 1, 2$ with IA representations $(u_j, \hat{p}_j, \mu_j, \kappa_j)$. Then $(\succcurlyeq^2_\upsilon, \succcurlyeq^2_\iota)$ has a GAFP than $(\succcurlyeq^1_\upsilon, \succcurlyeq^1_\iota)$ if and only if $u_1 \approx u_2$, $\hat{p}_1 = \hat{p}_2$, hence $V^1(\cdot|\upsilon) =$*

$V^2(\cdot|\upsilon)$*, and there exists* $\gamma \in [0,1]$ *such that, for all* $F \in \mathcal{A}$*,*

$$V^2(F|\iota) = \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] - (1-\gamma)\kappa_1 \left[\int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)] d\mu_1(p) - \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] \right]$$

Individual 2 has a GAFP than 1 if they share the same identity (the prior $\hat{p}$, the utility over payoffs u), expect to receive the same information μ, but the cost of information for 2 is lower than that for 1. This holds because $\gamma \in [0,1]$, hence $\kappa_2 = (1-\gamma)\kappa_1 \le \kappa_1$. The parameter γ represents a normalized measure of the comparative ability to follow the prescriptions. When $\gamma = 0$, the two individuals have identical preferences. When $\gamma = 1$, the preference $\succsim^2_\iota$ coincides with $\succsim^2_\upsilon$, hence individual 2 has the highest ability to follow the prescription (since information is not costly for her).

7 Information and welfare

In this section, I investigate the possibility to provide information to an identity-caring individual without lowering her welfare, hence maintaining incentive compatibility. I consider a sender, a receiver and an experiment $\mu \in \Delta\Delta\Omega$ called an information policy. A sender could be, for example, a regulator imposing information disclosure to sellers in a given market: in this case, the receiver is the consumer. Alternatively, the sender could be an authority that designs medical tests. The IA model predicts that information is always weakly welfare decreasing (Proposition 2). *Ceteris paribus*, an individual with a desire to protect identity weakly prefers to remain ignorant.

Moreover, the interpretation of the IA model makes information non-instrumental, hence the Sender cannot expect to persuade the receiver. However, the sender may be interested in providing information in itself, regardless of the receiver's response to information.[19] The problem faced by the sender is to design an information policy that is informative but welfare equivalent to ignorance. The next definition introduces the well-known Blackwell's informativeness order (Blackwell, 1953):

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

[19] Alternatively, the sender can exploit naivety. If second-period choice partially (and unanticipatedly) reacts to information, the sender will like to provide information.

informative than ν, written $\mu \trianglerighteq \nu$, if

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

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

An experiment μ is Blackwell more informative than ν, if any experimenter (associated to a payoff function ϕ) will prefer μ to ν.

Formally, the information policy designed by the Sender has to solve the following problem, for a fixed $F \in \mathcal{A}$:

$$\begin{aligned} \mu_F \in \underset{\mu \in \Gamma(\hat{p})}{\operatorname{argmax}} \trianglerighteq \\ \text{s.t. } C(\mu_F, F) = 0 \end{aligned} \tag{P}$$

The solution μ^*_F to P is the most informative experiment consistent with the prior $\hat{p}$ that is also costless. The problem P has always a trivial solution $\mu^*_F = \delta_{\hat{p}}$, consisting in no-information provision. However, other solutions are possible. For an $F \in \mathcal{A}$ and $f \in \operatorname{argmax}_{g \in F} \mathbb{E}_{\hat{p}}[u(g)]$, I define $\Delta^*_F(f) = \left\{p \in \Delta\Omega : f \in \operatorname{argmax}_{g \in F} \mathbb{E}_p[u(g)]\right\}$, the set of beliefs that recommend to choose f from F. Note that $\Delta^*_F(f)$ is a convex set, it is also compact since it is the intersection of a normal cone and $\Delta\Omega$. Moreover, $\hat{p} \in \Delta^*_F(f)$. The next result provides a sufficient condition for a non-trivial solution of problem (P) to exist:

Proposition 3. *If $\hat{p} \subsetneq \Delta^*_F(f)$ there exists $\mu_F \neq \delta_{\hat{p}}$ that solves Problem* (P)*.*

It is essential to note that the result in Proposition 3 crucially depends on the structure of the IA model, in particular, the cost of information $C(\mu, F)$. Suppose that the individual has a belief-dependent utility. In that case, information changing the posterior can equally affect welfare even if the recommended action is unchanged. That is, an information policy solving problem (P) can be strictly welfare decreasing. Therefore, there might be no solutions to problem (P), except for the trivial one.

The result in Proposition 3 is qualitatively similar to the result of Schweizer and Szech (2018). They consider the optimal design of payoff-relevant (medical) tests when individuals have anticipatory utility. They show that the optimal test should be perfectly informative if the true state is the preferred state of the agent and noisy otherwise. In the IA model, rather than preferred state, the agent has a preferred action, but the intuition of Schweizer and Szech

(2018) is still valid. The key difference is that, in the IA model, there is no role for anticipatory utility. Alternative results concerning the provision of information to behavioral individuals highlight the importance of recommendations (e.g. Lipnowski and Mathevet, 2018). Given an information policy μ, a recommendation is an information policy ν that induces the same choice (ex-post), but it is less informative than μ (i.e. $\mu \trianglerighteq \nu$). In the IA model recommendations cannot be strictly optimal because they may recommend actions that are different from the prescription (to the same extent of the original signal).

8 Extension: flexible information acquisition

The IA model can be easily extended to allow for flexible information acquisition. To characterize such an extension, I consider a set of information structures (experiments) parameterized by $\theta \in \Theta \cup \{\upsilon\}$. The primitive is a family of preferences $\{\succsim_\theta\}_{\theta \in \Theta \cup \upsilon}$ defined on $\mathcal{A}$ and conditional on receiving information from the experiment $\theta \in \Theta \cup \upsilon$.

Definition 4. *A family of preferences $\{\succsim_\theta\}_{\theta \in \Theta \cup \upsilon}$ has a Flexible Information Aversion representation if, for each $\theta \in \Theta$, the pair $(\succsim_\upsilon, \succsim_\theta)$ has a IA representation $(u, \hat{p}, \mu_\theta, \kappa_\theta)$, that is, $\succsim_\upsilon$ has a representation $V(\cdot|\upsilon) : \mathcal{A} \to \mathbb{R}$ given by:*

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

and each $\succsim_\theta$ has a representation $V(\cdot|\theta) : \mathcal{A} \to \mathbb{R}$ given by:

$$V(F|\theta) = \max_{f \in F} \mathbb{E}_{\hat{p}}[u(f)] - \kappa_\theta C(\mu_\theta, F) \tag{I_θ}$$

where

$$C(\mu_\theta, F) = \int_{\Delta\Omega} \max_{g \in F} \mathbb{E}_p[u(g)] d\mu_\theta(p) - \max_{f \in F} \mathbb{E}_{\hat{p}}[u(f)]$$

The behavioral characterization of the FIA representation is a straightforward extension of the one provided for the IA. In the FIA representation, all preferences $\succsim_\theta$ share the same prescriptions (the prior $\hat{p}$ and utility u), hence all experiments μ_θ are consistent with $\hat{p}$, that is $\mu_\theta \in \Gamma(\hat{p})$ for all $\theta \in \Theta$. The FIA model rationalizes an extreme form of *selective attention*. Consider the problem of selecting the optimal information θ for an individual with preferences represented by the FIA model. For a fixed menu F, the individual will select an information source θ as to maximize $V(F|\theta)$. The following result is an immediate consequence of Proposition 2.

Proposition 4. *For a given $F \in \mathcal{A}$:*

$$\theta \in \operatorname*{argmax}_{\theta' \in \Theta} V(F|\theta') \quad \Longleftrightarrow \quad C(\mu_\theta, F) = 0$$

The optimal solution is the set of all experiments θ delivering costless information. By Proposition 1: $C(\mu_\theta, F) = 0$ if and only if there exists $f \in \operatorname{argmax}_{g\in F} \mathbb{E}_{\hat{p}}[u(g)]$ such that $f \in \operatorname{argmax}_{g\in F} \mathbb{E}_p[u(g)]$, μ_θ-a.s. Proposition 4 predicts that the individual is indifferent between remaining ignorant and acquiring the information that confirms the prescription. As shown in Proposition 3, there are solutions to the optimization problem that are non-trivial.

9 Related literature

Identity was introduced to economics by Akerlof and Kranton (2000, 2005, 2010) with the objective of integrating results in social psychology into the traditional analysis of individual behavior based only on material incentives. Identity is a multidimensional and dynamic concept. An individual has multiple coexisting identities depending on gender, occupation, political view, social class etc. that may be more or less salient for a given choice.[20] What matters for the present work is that identity, whether mono or multidimensional, is internalized by the individual and generates a ranking among actions.

There are two papers that are most closely related to the current work: Bénabou and Tirole (2011) in spirit and Sarver (2008) from a technical point of view (the discussion of the differences with Sarver (2008) is at postponed until the end of this Section). In Bénabou and Tirole (2011), the individual has imperfect memory and uses past actions to (self-) signal her identity. Combined with temptation in the form of quasi-hyperbolic discounting their model accommodates a variety of empirical facts including a form of information avoidance. The present work takes an alternative approach and shows that identity-protection does not require

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

behavioral biases, such as imperfect memory and temptation. The individual in the IA model knows who she is, processes information in a Bayesian way, values actions according to their expected utility and does not suffer from temptation. These assumptions also differentiate the present work from the numerous alternative rationalizations of information aversion or avoidance (see Hertwig and Engel, 2016; Golman et al., 2017, for recent reviews). These include cognitive constraints, motivated beliefs, anticipatory feelings, disappointment aversion and sophisticated time-inconsistency. The theory of rational inattention developed in Sims (2003) posits that cognitive limitations forces the individual to optimally acquire information subject to a cost. In this model and its generalizations (Oliveira et al., 2017; Pennesi, 2015), the cost of information depends on the informativeness of the acquired signal, hence independent of the menu faced by the individual. Therefore, flexibility is weakly valuable, an assumption that seems unable to capture identity-protective behavior (see the example in Section 2.1).

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

The paper contributes to the literature studying unobservable information acquisition. There are two main approaches that are different in the type of revealed preferences adopted:

ex-ante (choices among menus) or ex-post (choices from menus). Models focusing on ex-ante choices are Epstein (2006), 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). The IA model is related to Dillenberger et al. (2014). They axiomatize the Arrow and Fisher (1974) model where information arrival leads to a preference for flexibility. Theorem 3 in the present paper generalizes their subjective learning model by relaxing the preference for flexibility axiom. Lu (2016) studied private information when the available data are stochastic choices from menus. Oliveira et al. (2017) and Pennesi (2015) axiomatize a model of costly information acquisition, in the tradition of rational inattention of Sims (2003), where the individual decides what information to pay attention to subject to a cognitive cost. Both models assume a preference for flexibility, because the cost of information is menu-independent.

Concerning the relation between commitment and information processing, few models allow for a preference for commitment in the presence of objective uncertainty and learning. The seminal work of Epstein (2006) builds 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 et al. (2008). In my model, the second-period choice is made from the menu selected initially. This makes the comparability with my model difficult. In any case, both models (Epstein, 2006; Epstein et al., 2008) assume Set-Betweenness. The motivating example in Section 2.1 argues that Set-Betweenness is unable to capture identity protection. Most importantly, in the IA model information processing is Bayesian, and the preference for commitment derives from a desire to protect identity. In a different setting with menus of lotteries, Alauoi (2016) proved that the preservation of self-image may lead to a preference for smaller menus. In his approach, lotteries in the menu convey information about the unobservable ability of the decision-maker and a doubt-prone individual may restrict her options as to remain “ignorant".

Lastly, I discuss the differences between the IA model and the anticipated regret model of Sarver (2008). The function $V(\cdot|\iota)$ has a structure similar to the regret model of Sarver (2008), but the setting and both the interpretation of the model and of the axioms is different.

First, in his model uncertainty is subjective, depends on the possible subjective states (tastes or moods) of the individual and the objects of choice are menus of lotteries. In the IA model uncertainty is objective, there is the possibility to acquire information (this is not the case in Sarver's model) and the choice objects are menus of Anscombe-Aumann acts. Second, information acquisition in the present paper is modeled assuming two (or more) preference relations and the characterization of the IA model requires axioms relating both preferences (the Identity Preservation axiom and the Consistent Weak Order axiom). In Saver's model there is no information acquisition and only one preference is necessary to identify the regret model. Third, the regret interpretation in Sarver's model is based on the assumption that choices from menus happen before the resolution of the subjective uncertainty. In the IA model, the interpretation is that choices from menus occurs after the resolution of (a first layer of) uncertainty, the one concerning the posterior. Forth, consider two acts f, g with $f(\omega) \succcurlyeq_{v} g(\omega)$, then by Dominance and Identity Preservation $f \succcurlyeq_{v} g$ and $f \sim_{\iota} f \cup g$. So a preference for commitment can occur only when information contradicts the prescription. If an action unambiguously dominates another, their ranking will remain the same regardless of the information acquired. This property has no counterpart in the regret representation. Lastly, it is commonly understood that a subject experiences regret *after* deciding and after uncertainty resolves.[21] This is the model of Sarver (2008), where regret occurs after a choice from the menu is made and after the uncertain taste of the individual realizes. Differently, the cost of information in the IA model occurs in an intermediate period: after choosing a menu but *before* selecting an action from the menu. Hence, it is the dissonance between prescriptions and information that generate the cost of information. It is clearly possible to include regret: if the selected action is, after state-uncertainty resolves, inferior to an available alternative, the individual can also experience regret. However, also an individual that is insensitive to regret will bear the cost of information if she cares about identity preservation (as it happens in the IA). Allowing for regret will exacerbate the cost of information. To avoid confounding factors, in the IA model actions are evaluated according their expected utility, hence leaving no room for regret.

[21]According to Loomes and Sugden (1982), an individual experiences regret if "he may reflect on how much better his position would have been, had he chosen differently, and this reflection may reduce the pleasure that he derives.. "

A Proofs

Proof of Proposition 2. For the first inequality, since $\kappa \geq 0$ and $C(\mu, F) \geq 0$, $V(F|\iota) = \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] - \kappa C(\mu, F) \leq \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] = V(F|\upsilon)$. For the second point, $\mathbb{E}_{\hat{p}}[u(f)] = V(f|\upsilon) = V(f \cup g|\upsilon) \geq V(f \cup g|\iota)$, where the last inequality follows form point 1. □

I first prove the following result that generalizes Dillenberger et al. (2014):

Theorem 3. *For $\theta \in \{\upsilon, \iota\}$, a binary relation $\succsim_\theta$ is a complete, transitive and satisfies Continuity, Independence and Dominance if and only if there exist a non-constant and affine $w_\theta : X \to \mathbb{R}$, a $\hat{p}_\theta \in \Delta\Omega$ and $\eta_\theta \in ca(\Delta\Omega)$ with $\int_{\Delta\Omega} p d\eta_\theta(p) = \hat{p}_\theta$, such that $\succsim_\theta$ is represented by $V_\theta : \mathcal{A} \to \mathbb{R}$:*

$$V_\theta(F) = \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[w_\theta(f)] d\eta_\theta(p) \tag{3}$$

The preference $\succsim_\theta$ also satisfies Preference for Flexibility if and only if $\eta_\theta \in \Delta\Delta\Omega$.

Proof of Theorem 3. Since $\mathcal{F}$ is a convex subset of a vector space, by the Anscombe-Aumann theorem, completeness, transitivity, Continuity, Dominance and Independence restricted to singletons hold, if and only if, there exist an affine and cardinally unique function $w_\theta : X \to \mathbb{R}$ and a probability $\hat{p}_\theta \in \Delta\Omega$ such that $V_\theta(f) \equiv \int_\Omega w_\theta(f) d\hat{p}_\theta$ represents the restriction of $\succsim_\theta$ to singleton menus. W.l.o.g. let assume that $[0, 1] \subseteq w_\theta(X)$. By Dominance and Continuity, every menu has a singleton (certainty) equivalent: $x_F^\theta \sim_\theta F$ with $x_F^\theta \in X$. I can extend V_θ from $\mathcal{F}$ to $\mathcal{A}$ by defining $V_\theta(F) \equiv V_\theta(x_F^\theta)$ for each $F \in \mathcal{A}$. Now, for each $F \in \mathcal{A}$, define $\sigma_F : \Delta\Omega \to \mathbb{R}$ as $\sigma_F(p) = \max_{f\in F} \int_\Omega w_\theta(f) dp$ and $S \equiv \{\sigma_F(\cdot) \in C(\Delta\Omega) : F \in \mathcal{A}\}$, where $C(\Delta\Omega)$ is the space of continuous functions defined on $\Delta\Omega$. Define $W_\theta : S \to \mathbb{R}$ as $W_\theta(\sigma_F) = w_\theta(x_F^\theta)$. By completeness, transitivity, Independence and Continuity (part a), $\succsim_\theta$ satisfies indifference to randomization, so $F \sim_\theta co\, F$ for all $F \in \mathcal{A}$. To show W_θ is well defined, consider $\sigma_F = \sigma_G$, then $co\, F = co\, G$, by Indifference to Randomization, $F \sim_\theta co\, F = co\, G \sim_\theta G$, hence $W_\theta(\sigma_F) = w_\theta(x_F) = w_\theta(x_G) = W_\theta(\sigma_G)$, so W_θ is well defined and represents $\succsim_\theta$. By definition $V_\theta(F) = W_\theta(\sigma_F)$. By standard arguments, Independence implies affinity of the functional W_θ, i.e. $W_\theta(\alpha\sigma_F + (1 - \alpha)\sigma_G) = \alpha W_\theta(\sigma_F) + (1 - \alpha)W_\theta(\sigma_G)$ for all $F, G \in \mathcal{A}$ and all $\alpha \in [0, 1]$. By definition, W_θ is also normalized, $W_\theta(\sigma_x) = w_\theta(x)$, then $W_\theta(0) = 0$. By Continuity (part b) and an argument similar to Sarver (2008) W_θ is Lipschitz continuous. Define $rS = \{r\phi : \phi \in S\}$ and let $H = \cup_{r\geq 0} rS$. Lastly, let $H^* = H - H = \{\phi_1 - \phi_2 \in C(\Delta\Omega) : \phi_1, \phi_2 \in H\}$. For all

$\phi \in H$ and $\phi \neq 0$, there is $r > 0$ such that $\frac{1}{r}\phi \in S$. There exists a linear extension of W_θ to H^* denoted by W_θ^*. Such an extension is also Lipschitz continuous.

Since H^* is dense in $C(\Delta\Omega)$ and W_θ^* is linear and Lipschitz continuous from H^* to $\mathbb{R}$, there exists a unique linear extension of W_θ^* to $C(\Delta\Omega)$ denoted by $\hat{W}_\theta$. By the Riesz's Representation Theorem there exists a signed countably additive measure $\eta_\theta \in ca(\Delta\Omega)$ such that $\hat{W}_\theta(\phi) = \int_{\Delta\Omega} \phi(p)\eta_\theta(dp)$ for all $\phi \in C(\Delta\Omega)$. To see that $\int_{\Delta\Omega} p\eta_\theta(dp) = \hat{p}_\theta$, consider a singleton f, then $V_\theta(f) = \int_\Omega w_\theta(f)d\hat{p}_\theta = w_\theta(x_f) = \hat{W}_\theta(\sigma_f) = \int_{\Delta\Omega} \mathbb{E}_p[w_\theta(f)]d\eta(p)$ and consider $w_\theta(f(\omega)) = 1$ and $w_\theta(f(\omega')) = 0$ otherwise, then $\hat{p}_\theta(\omega) = \int_{\Delta\Omega} p(\omega)d\eta_\theta(p)$ for all $\omega \in \Omega$. Normalization implies $\eta_\theta(\Delta\Omega) = 1$, indeed, take $x \in X$ with $w_\theta(x) = 1$, then $1 = \hat{W}_\theta(\sigma_x) = \int_{\Delta\Omega} d\eta_\theta(p) = \eta_\theta(\Delta\Omega)$.

If $\succsim_\theta$ satisfies the Preference for Flexibility axiom, the result follows from Dillenberger et al. (2014).

□

Before proving Theorem 1, I will prove the following lemmas:

Lemma 1. *If $\succsim_v$ satisfies Consistent Weak Order, Continuity, Independence, Dominance and U-Strategic Rationality, and for all $g \in G$ there is $f \in F$ with $f \succsim_v g$, then $F \sim_v F \cup G$.*

Proof of Lemma 1. Take any $G = (g_1, g_2, \ldots, g_n)$ in $\mathcal{A}$. If for all $g \in G$, there is $f \in F$ such that $f \succsim_v g$, then $F \sim_v F \cup G$. Consider g_1, by definition $g_1 \in G$, hence there is $f \in F$ such that $f \succsim_v g_1$. By U-Strategic Rationality, $F \sim_v F \cup g_1$. An identical argument applies to g_2, so that $F \cup g_1 \sim_v F \cup g_1 \cup g_2$. By Consistent Weak Order, $F \sim_v F \cup g_1 \cup g_2$. Repeating the argument n-times implies that $F \sim_v F \cup G$. □

Lemma 2. *The binary relation $\succsim_v$ satisfies Consistent Weak Order, Continuity, Independence, Dominance and U-Strategic Rationality if and only if $\succsim_v$ is represented by $V_v : \mathcal{A} \to \mathbb{R}$ given by $V_v(F) = \max_{f \in F} \mathbb{E}_{\hat{p}_v}[w_v(f)]$, where w_v and $\hat{p}_v$ are defined in Theorem 3.*

Proof of Lemma 2. By Theorem 3, Consistent Weak Order, Continuity, Independence and Dominance implies that there exists an affine representation V_v of $\succsim_v$. When restricted to singletons, the condition $\int_{\Delta\Omega} pd\eta_v(p) = \hat{p}_v$ implies $V_v(f) = \int_{\Delta\Omega} \mathbb{E}_p[w_v(f)]d\eta_v(p) = \mathbb{E}_{\hat{p}_v}[w_v(f)]$. I will prove the Lemma by induction on the size of $F \in \mathcal{A}$. For all $|F| = 1$, the result is trivially true. Suppose it is true for all $F' \in \mathcal{A}$ with $|F'| = n-1$. Take an arbitrary F' with $|F'| = n-1$ and an $f \notin F'$ and define $F = F' \cup f$. Then, there are two cases: $F' \succ_v f$ or $f \succsim_v F'$. In the case of $F' \succ_\upsilon f$, the induction hypothesis implies $\max_{f' \in F'} \mathbb{E}_{\hat{p}}[u(f')] = V_\upsilon(F') > V_\upsilon(f) = \mathbb{E}_{\hat{p}}[w_\upsilon(f)]$. Therefore, there is $g \in F'$ such that $V_\upsilon(F') = \mathbb{E}_{\hat{p}_\upsilon}[w_\upsilon(g)]$ and $g \succ_\upsilon f'$ for all $f' \in F' \cup f = F$. By Lemma 1 $g \sim_\upsilon F' \sim_\upsilon F$. Then, $\max_{f \in F} \mathbb{E}_{\hat{p}_\upsilon}[w_\upsilon(f)] = \max\{\mathbb{E}_{\hat{p}_\upsilon}[w_\upsilon(f)], V_\upsilon(F')\} = V_\upsilon(F') = V_\upsilon(F)$. For the case $f \succsim_\upsilon F'$, the induction hypothesis implies $\mathbb{E}_{\hat{p}_\upsilon}[w_\upsilon(f)] = V_\upsilon(f) \geq V_\upsilon(F') = \max_{f' \in F'} \mathbb{E}_{\hat{p}_\upsilon}[w_\upsilon(f')]$. Lemma 1 implies $f \sim_\upsilon F' \cup f = F$. Therefore $\max_{f \in F} \mathbb{E}_{\hat{p}_\upsilon}[w_\upsilon(f)] = \max\{\mathbb{E}_{\hat{p}_\upsilon}[w_\upsilon(f)], \max_{f' \in F'} \mathbb{E}_{\hat{p}_\upsilon}[w_\upsilon(f')]\} = \mathbb{E}_{\hat{p}}[w_\upsilon(f)] = V_\upsilon(f) = V_\upsilon(F)$. Since the representation holds for all $F \in \mathcal{A}$ with $|F| = n$. By induction it holds for all $F \in \mathcal{A}$.

$\square$

Fact 2. If $\succsim_\iota$ satisfies Consistent Weak Order, Continuity, Independence, Dominance and Identity Preservation, and for all $g \in G$ there is $f \in F$ with $f \succsim_\upsilon g$, then $F \succsim_\iota F \cup G$.

Proof of Fact 2. Take any $G = (g_1, g_2, \ldots, g_n)$ in $\mathcal{A}$. If for all $g \in G$, there is $f \in F$ such that $f \succsim_\upsilon g$, then $F \sim_\iota F \cup G$. Consider g_1, by definition $g_1 \in G$, hence there is $f \in F$ such that $f \succsim_\upsilon g_1$. By Identity Preservation, $F \succsim_\iota F \cup g_1$. An identical argument applies to g_2, so that $F \cup g_1 \succsim_\iota F \cup g_1 \cup g_2$. By Consistent Weak Order, $F \succsim_\iota F \cup g_1 \cup g_2$. Repeating the argument n-times implies that $F \succsim_\iota F \cup G$. $\square$

Proof of Theorem 1. Necessity follows from straightforward arguments. For sufficiency, Theorem 3 implies that both $\succsim_\upsilon$ and $\succsim_\iota$ have an additive representation and, therefore, there exist affine functionals $V_\upsilon, V_\iota : \mathcal{A} \to \mathbb{R}$ such that $F \succsim_\upsilon G$, if and only if, $V_\upsilon(F) \geq V_\upsilon(G)$ and $F \succsim_\iota G$, if and only if, $V_\iota(F) \geq V_\iota(G)$. When restricted to singletons, $V_\upsilon(f) = \mathbb{E}_{\hat{p}_\upsilon}[w_\upsilon(f)]$ for some affine $w_\upsilon : X \to \mathbb{R}$ and some $\hat{p}_\upsilon \in \Delta\Omega$. Similarly, $V_\iota(f) = \mathbb{E}_{\hat{p}_\iota}[w_\iota(f)]$ for some affine $w_\iota : X \to \mathbb{R}$ and some $\hat{p}_\iota \in \Delta\Omega$. By Consistent Weak Order, V_υ and V_ι are ordinally equivalent when restricted to singletons. By the uniqueness property of the Anscombe-Aumann expected utility, w_υ and w_ι are cardinally equivalent and they can be normalized (if necessary) to have $w_\upsilon = w_\iota \equiv u$. Moreover, $\hat{p}_\upsilon = \hat{p}_\iota \equiv \hat{p}$. Since both V_υ and V_ι have an additive representation (3), there exist linear functionals $W_\upsilon, W_\iota : C(\Delta\Omega) \to \mathbb{R}$ such that $F \succsim_\upsilon G$, if and only if, $W_\upsilon(\sigma_F) \geq W_\upsilon(\sigma_G)$ and $F \succsim_\iota G$, if and only if, $W_\iota(\sigma_F) \geq W_\iota(\sigma_G)$. Moreover, by the common normalization implied by Consistent Weak Order, W_υ and W_ι coincides when restricted to singletons: $W_\upsilon(\sigma_f) = \mathbb{E}_{\hat{p}}[u(f)] = W_\iota(\sigma_f)$ for all $f \in \mathcal{F}$. Now, consider $F, G \in \mathcal{A}$ such that $\max_{f \in F} \mathbb{E}_{\hat{p}}[u(f)] = \max_{g \in G} \mathbb{E}_{\hat{p}}[u(g)]$ and $\max_{f \in F} \mathbb{E}_p[u(f)] \leq \max_{g \in G} \mathbb{E}_p[u(g)]$ for all $p \in \Delta\Omega$.

Take a $f' \in \operatorname{argmax}_{f\in F} \mathbb{E}_{\hat{p}}[u(f)]$, then

$$\mathbb{E}_{\hat{p}}[u(g)] \leq \max_{g\in G} \mathbb{E}_{\hat{p}}[u(f)] = \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] = \mathbb{E}_{\hat{p}}[u(f')]$$

hence $f' \succsim_\upsilon g$ and by Fact 2, $F \succsim_\iota F \cup G$. By the condition $\max_{f\in F} \mathbb{E}_p[u(f)] \leq \max_{g\in G} \mathbb{E}_p[u(g)]$ for all $p \in \Delta\Omega$, it holds that $F \subseteq G$, therefore $F \succsim_\iota F \cup G = G$ implying $W_\iota(\sigma_F) \geq W_\iota(\sigma_G)$. A proof similar to that of Sarver (Lemma 12 in 2008) implies that the extension of W_ι to $C(\Delta\Omega)$ defined in the proof of Theorem 3 has the Domination property: for all $\phi, \psi \in C(\Delta\Omega)$, if $\phi(\hat{p}) = \psi(\hat{p})$ and $\phi \leq \psi$, then $\hat{W}_\iota(\phi) \geq \hat{W}_\iota(\psi)$. Since $\hat{W}_\iota$ is a linear functional, it can be rewritten as the difference of two positive linear functionals $\hat{W}_\iota(\phi) = \hat{W}_\iota^+(\phi) - \hat{W}_\iota^-(\phi)$. An argument similar to Lemma 13 in Sarver (2008) can be applied to show that the domination property of $\hat{W}_\iota$ implies that, for any pair of support functions $\sigma_F, \sigma_G \in C(\Delta\Omega)$, $W_\upsilon(\sigma_F) = \sigma_F(\hat{p}) \geq \sigma_G(\hat{p}) = W_\upsilon(\sigma_G)$ implies $W_\iota^+(\sigma_F) \geq W_\iota^+(\sigma_G)$. Since both W_υ and W_ι^+ are affine, standard results (e.g. De Meyer and Mongin, 1995) there are $\alpha > 0$ and $\beta \in \mathbb{R}$ such that $W^+(\sigma_F) = \alpha\sigma_F(\hat{p}) + \beta$. Taking $\sigma_x(p)$ with $u(x) = 0$, one has $\sigma_x(p) = 0$ for all $p \in \Delta\Omega$, hence $\beta = 0$.

By normalization, taking $u(y) = 1$, $1 = W_\iota(\sigma_y) = \alpha\sigma_y(\hat{p}) - W^-(\sigma_y)$, then $W^-(1) = \alpha - 1$. Since W^- is a positive functional $\alpha - 1 \geq 0$. Then W_ι can be rewritten as $W_\iota(\sigma_F) = \alpha\sigma_F(\hat{p}) - (\alpha - 1)\bar{W}_\iota^-(\sigma_F)$, where $\bar{W}_\iota^- \equiv \frac{W_\iota^-}{\alpha - 1}$ is a linear, positive and normalized functional. By Riesz's representation theorem, there exists a probability $\mu \in \Delta\Delta\Omega$ such that $\bar{W}_\iota^-(\sigma_F) = \int_{\Delta\Omega} \max_{f\in f} \mathbb{E}_p[u(f)]d\mu(p)$. By definition, for all singletons $f \in \mathcal{F}$, $\mathbb{E}_{\hat{p}}[u(f)] = W_\iota(\sigma_f) = \alpha\mathbb{E}_{\hat{p}}[u(f)] - (\alpha - 1)\int_{\Delta\Omega} \mathbb{E}_p[u(f)]d\mu(p)$, that implies $\int_{\Delta\Omega} p\mu(p) = \hat{p}$, hence $\mu \in \Gamma(\hat{p})$. Lastly $W_\iota(\sigma_F) = \alpha \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] - (\alpha - 1)\int_{\Delta\Omega} \max_{f\in f} \mathbb{E}_p[u(f)]d\mu(p)$. Since $\alpha \geq 1$, let define $\kappa \equiv 1 - \alpha$ that implies $W_\iota(\sigma_F) = (1 - \kappa)\max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] - \kappa\int_{\Delta\Omega} \max_{f\in f} \mathbb{E}_p[u(f)]d\mu(p)$, the representation in Equation (I). Lemma 2 proves that the binary relation $\succsim_\upsilon$ has the representation in Equation (U). By defining $V(\cdot|\iota) \equiv V_\iota$ and $V(\cdot|\upsilon) \equiv V_\upsilon$ the proof is complete.

□

Proof of Theorem 2. Note that $\succsim_\upsilon^1$ coincides with $\succsim_\iota^1$ of $\mathcal{F}$, hence $f \succsim_\upsilon^1 g$ if and only if $f \succsim_\iota^1 g$. By the condition $f \succsim_\iota^1 g$ implies $f \succsim_\iota^2 g$. By the standard arguments there exists $\alpha > 0$ and $\beta \in \mathbb{R}$ such that $u_1 = \alpha u_2 + \beta$ and $\hat{p}_1 = \hat{p}_2$. So I can normalize $u_1 = u_2$ if necessary. Since $\succsim_\upsilon^1$ and $\succsim_\upsilon^2$ are identical, I denote by $\succsim_\upsilon$ both of them. By the condition of the theorem,

$V(F|\upsilon) \geq V(G|\upsilon)$ and $V^1(F|\iota) \geq V^1(G|\iota)$ imply $V^2(F|\iota) \geq V^2(G|\iota)$. Then, by standard results on aggregation of affine functions (see e.g. De Meyer and Mongin, 1995) there exist $\gamma, \beta \geq 0$ and $c \in \mathbb{R}$ such that:

$$V^2(F|\iota) = \gamma V(F) + \beta V^1(F|\iota) + c$$

for all $F \in \mathcal{A}$. Taking $F = f$, $V^2(f|\iota) = V(f|\upsilon) = \gamma V(f|\upsilon) + \beta V_\upsilon(f|\upsilon) + c$, hence $c = 0$ and $\gamma + \beta = 1$.

$$\begin{aligned}
V^2(F|\iota) =& \gamma \max_{f\in F} \mathbb{E}_{\hat{p}}[u(f)] + (1-\gamma)\left[(1+\kappa_1)\max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)] - \kappa_1 \int_{\Delta\Omega} \max_{f\in F} \mathbb{E}_p[u(f)]d\mu_1(p)\right] \\
=& (\gamma + (1-\gamma)(1+\kappa_1))\max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)] - (1-\gamma)\kappa_1 \int_{\Delta\Omega}\max_{f\in F}\mathbb{E}_p[u(f)]d\mu_1(p) \\
=& \max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)] - (1-\gamma)\kappa_1 \left[\int_{\Delta\Omega}\max_{f\in F}\mathbb{E}_p[u(f)]d\mu_1(p) - \max_{f\in F}\mathbb{E}_{\hat{p}}[u(f)]\right]
\end{aligned}$$

For the reverse implication. By the assumption $u_1 \approx u_2$ and $\hat{p}_1 = \hat{p}_2$, it follows that $V^1(\cdot|\upsilon) = V^2(\cdot|\upsilon)$ and I denote it $V(\cdot|\upsilon)$. By the condition of the theorem, $V^2(F|\iota) = \gamma V(F|\upsilon) + (1-\gamma)V^1(F|\iota)$. Assume that $V(F|\upsilon) \geq V(G|\upsilon)$ and $V^1(F|\iota) \geq V^1(G|\iota)$. Hence, $V^2(F|\iota) = \gamma V(F|\upsilon) + (1-\gamma)V^1(F) \geq \gamma V(G|\upsilon) + (1-\gamma)V^1(G|\iota) = V^2(G|\iota)$. Since F, G were arbitrary, the result holds.

□

Proof of Proposition 3. By the Krein-Milman Theorem, $\Delta^*_F(f) = \overline{co}(ext\Delta^*_F(f))$, where $\overline{co}(A)$ is the closed convex hull of an arbitrary set A and $ext(A)$ is the set of its extreme points. Therefore, $\hat{p}$ cannot be the unique extreme point of $\Delta^*_F(f)$. By the Choquet-Bishop-de Leeuw Theorem (Alfsen, 1971, Theorem I.4.8), since $\hat{p} \in \Delta^*_F(f)$ there exists a maximal (in the order $\trianglerighteq$) measure μ_F on $ext(\Delta^*_F(f))$ such that $\mu_F \neq \delta_{\hat{p}}$. □

References

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

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

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

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

Alfsen, E. M. (1971). *Compact Convex Sets and Boundary Integrals.* Springer-Verlag, New York.

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.

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.

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., 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.*

De Meyer, B. and P. Mongin (1995). A note on affine aggregation. *Economics Letters 47*(2), 177–183.

Dekel, E., B. L. Lipman, and A. Rustichini (2001, July). Representing Preferences with a Unique Subjective State Space. *Econometrica 69*(4), 891–934.

Dekel, E., B. L. Lipman, and A. Rustichini (2009). Temptation-driven preferences. *The Review of Economic Studies 76*(3), 937–971.

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

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

Ehrich, K. R. and J. R. Irwin (2005). Willful ignorance in the request for product attribute information. *Journal of Marketing Research 42*(3), 266–277.

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.

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

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

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

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.

Gul, F. and W. Pesendorfer (2001). Temptation and self-control. *Econometrica 69*(6), 1403–1435.

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

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

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

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

Kahan, D. M. and D. Braman (2006). Cultural cognition and public policy. *Yale L. & Pol'y Rev. 24*, 149.

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

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

Köszegi, B. and M. Rabin (2006, Nov). A model of reference-dependent preferences. *The Quarterly Journal of Economics CXXI*(4).

Kreps, D. M. (1979, May). A Representation Theorem for "Preference for Flexibility". *Econometrica 47*(3), 565–577.

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

Loomes, G. and R. Sugden (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal 92*(368), 805–824.

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

Marks, L. and R. Palkovitz (2004). American fatherhood types: The good, the bad, and the uninterested. *Fathering: A Journal of Theory, Research & Practice about Men as Fathers 2*(2).

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.

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

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

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

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

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

Thunström, L., J. Nordström, J. F. Shogren, M. Ehmke, and K. van't Veld (2016, Apr). Strategic self-ignorance. *Journal of Risk and Uncertainty 52*(2), 117–136.

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.