RATIONALITY
& COMPETITION
CRC TRR 190

Persuasion Against Self-Control Problems

Jonas von Wangenheim (Humboldt University Berlin)

Discussion Paper No. 98

May 29, 2018

Persuasion against Self-Control Problems*

Jonas von Wangenheim

Humboldt-Universität zu Berlin

This Version: May 25, 2018

First Version: March 15, 2018

Abstract

I derive a social planner's optimal information design in an environment with quasi-hyperbolic discounting consumers without commitment. Consumption induces instantaneous utility, but unknown delayed cost. Consumers may or may not acquire additional costless information on the cost parameter. The planner's optimal signal can be interpreted as an incentive compatible consumption recommendation whenever the cost parameter is below some cut-off. Welfare strictly exceeds the one under full information. I characterize distributional conditions under which welfare attains first best.

Keywords: Bayesian Persuasion, Present Bias, Hyperbolic Discounting, Rational Inattention

JEL classification: D01, D18, D62, D82

*I thank Paul Heidhues, Roland Strausz, and Jean Tirole for helpful comments. Financial support by Deutsche Forschungsgemeinschaft through CRC TRR 190 is gratefully acknowledged.

1 Introduction

Evidence suggests that many consumers have self-control problems (see e.g. DellaVigna (2009)): in order to receive instantaneous gratification from consumption, they overconsume products which create long-term health risks. The finding that consumers act against their own interest has led to various suggestions on possible paternalistic policies to help consumers make better choices: nudges in form of default options (Thaler and Sunstein (2003)), optimal "sin-taxes" on unhealthy products (O'Donoghue and Rabin (2006)), or simple prohibition of drugs as implemented in most countries.

I analyze optimal information design by a social planner as an alternative form of paternalism to induce consumption behavior that is (more) aligned with consumers' long run interests.

I draw a connection between the dynamic consumption model with quasi-hyperbolic discounting consumers by Carrillo and Mariotti (2000) and the optimal information design approach by Kamenica and Gentzkow (2011), and derive a consumer-optimal information signal about the consumers' risk type. The optimal signal consists of a simple binary signal, which displays whether or not the risk is below some threshold.

There are numerous examples where institutions implicitly use cut-off signals in form of recommendations, guidelines or definitions, when consumers themselves fail to thoroughly interpret data. For instance, many governments have adopted guidelines that define thresholds for responsible alcohol consumption (Kalinowski and Humphreys (2016)). the world health organization (WHO) defines a body mass index of $25kg/m^2$ as the cut-off point for overweight.[1]

Moreover, labels are a common tool to provide consumption recommendation on the basis of thresholds. For instance, in 2006 the Food Standard Agency (FSA) in the UK introduced a traffic light rating system for food nutrition values. Products with sugar or saturated fats above certain thresholds are highlighted with a red light to display the health risk of the product. Further, the organic food label by the European Union defines minimum requirements for food to be labelled as organic.[2]

My paper provides a rationale for the use of these information policies, and

[1] http://apps.who.int/iris/bitstream/10665/37003/1/WHO_TRS_854.pdf
[2] (EG) Nr. 889/2008 and (EG) Nr. 834/2007

shows how to choose the respective thresholds in order to induce the welfare maximizing consumption decision for consumers. Moreover, I give conditions under which the derived information signal is robust against additional costless learning by the consumer.

I consider an infinite horizon discrete time model, where the consumer may consume one unit of a good each period. Consumption induces instantaneous utility, but gives rise to a negative future externality with unknown probability. For instance, consider food products where the consumption risk of obesity or diabetes are a priori not transparent. Similarly, smoking, lack of sports, or other unhealthy activities feature health risks that may depend on genetic predispositions only testable by medical experts. The consumer has a strong value for instantaneous gratification in form of quasi-hyperbolic preferences, which leads to time-inconsistent preferences (Laibson (1997)).

Carrillo and Mariotti (2000) find that in this model there is value of rational inattention with respect to the risk parameter: Even though information helps a consumer to achieve a better myopic consumption choice, any information is shared with future incarnations and may lead to future overconsumption due to the present bias. Thus, rational inattention may work as a commitment device towards future incarnations to enable more favorable long-term consumption choices.

I follow the normative approach suggested by O'Donoghue and Rabin (2002, 2003) in defining consumer welfare as ex-ante consumer utility before facing consumption decisions.[3] As a social planner therefore takes the perspective of the consumer at time 0, the model can equivalently be regarded as a model of self-persuasion in which not the regulator but the consumer himself strategically acquires information some time before he faces the consumption decision.

For the analysis of the optimal information signal, I distinguish two cases.

First, I analyze the consumer-optimal signal when the consumer has no access to further information before making the consumption decisions. As the incentives of a welfare maximizing planner are aligned with the consumer whenever he does not face an instantaneous consumption decision, this seems appropriate where information is not immediately available at reasonable cost. For instance, it seems very implausible

[3]This approach implies that quasi-hyperbolic discounting is regarded as an "error".

that a consumer would first conduct a medical test on his risk type, whenever he instantaneously faces the opportunity to smoke.

I show that in this case the optimal signal consists of a risk threshold together with the simple information whether the risk type is above or below the threshold. Similar to Carrillo and Mariotti (2000), this signal improves welfare upon full information: while under full information the present bias leads to inefficient consumption of intermediate risk types, this cut-off signal pools (some of) these intermediate types with high risk types, and induces them to abstention. For many distributions, the optimal signal can implement first-best welfare.

There are numerous examples where institutions implicitly use cut-off signals in form of recommendations, guidelines or definitions, when consumers themselves fail to thoroughly interpret data. Definitions of obesity, limits for responsible drinking, classification of medical risk groups are only some examples. Moreover, certification in form of labels—such as the European Union eco label and energy label—are commonly based on predefined thresholds.

While there may also be other motives for the use of simple information, my model provides a novel rationale for the use of these tools in regulation. In particular, it provides an instruction on how such tools may be used to incentivize more preferable consumption decisions.

As a second case, I assume that consumers can acquire costless additional information at any time. This assumption seems pertinent in situations where the consumer can instantaneously access relevant product information online. In the welfare maximizing Markov perfect equilibrium of this game, the provided ex-ante information, again, consists of a simple cut-off signal, and consumers don't acquire further relevant information. Moreover, I show that under some regularity conditions on the risk distribution, this signal coincides with the optimal signal in case of full information control by the planner, and consequently achieves the same welfare. Intuitively, consumers abstain from acquiring more precise information as they fear that they will eventually end up in the full information equilibrium and overconsume forever.

The remainder of this paper is structured as follows: In Section 2, I discuss the related literature. In Section 3, I introduce the model. I start Section 4 with the benchmark of full information, before I analyze the case of full information control by the planner, and the case of costless consumer learning. Section 5 concludes this paper. All proofs are relegated to the appendix.

2 Related Literature

There is substantial evidence that individuals have dynamically inconsistent time preferences (Frederick et al. (2002), DellaVigna (2009)). As a common feature, discount rates increase as the date approaches. The employed (β, δ)-model of quasi-hyperbolic discounting dates back to Phelps and Pollak (1968), who used it to model imperfect inter-generational altruism. Laibson (1997) was first to use it in the context of an intra-personal conflict. It has arguably become the standard model of time-inconsistent preferences.

The existence of self-control problems and time-inconsistent preferences inherently gives value to devices that enable individuals to commit to future actions.

Many papers have analysed how the market can offer such a device by selling adequate goods such as illiquid assets (Laibson (1998), Diamond and Kőszegi (2003)), rationed quantities (Wertenbroch (1998)), or long-term memberships (DellaVigna and Malmendier (2006)).

My paper connects to another strand of the literature where individuals use belief manipulation as an intrapersonal commitment device. Bénabou and Tirole (2002) show how endogenously chosen imperfect recall may lead to overconfidence to overcome motivational problems. Brocas and Carrillo (2000) show how information avoidance can be welfare increasing under time-inconsistent preferences. My paper builds on the dynamic consumption model by Carrillo and Mariotti (2000). Consumption yields instantaneous utility but with unknown risk some delayed cost. They show that individuals with time-inconsistent preferences may prefer to abstain from information acquisition on the risk parameter, even if information is costless. Consumers, who have the ability to sample information according to a Bernoulli process, may fear to be trapped in inefficient consumption forever for intermediate risk estimates. In order to avoid this overconsumption they may stop sampling at beliefs that induce abstention.

My paper is also closely related to the very recent and independently developed work by Mariotti et al. (2018) who also find cut-off signals as the optimal persuasion mechanism in a 4-period model with consumers with self control problems. While they derive optimal mechanisms for different principal's objectives (e.g. consumption maximizing lobbyist, consumption minimizing policy maker), my work focuses on the robustness of the optimal mechanism in an infinite horizon game with additional learning opportunities.

While Carrillo and Mariotti show that there is value of stopping Bernoulli sampling, I derive the optimal information policy for their consumption model, when the consumer (or a regulator on behalf of the consumer) is unconstrained in the ability to design information signals (Kamenica and Gentzkow (2011)), and derive precise conditions when the optimal signal induces first-best consumption utility.

My paper also relates to the literature on paternalistic motives. In the context of quasi-hyperbolic discounting, O'Donoghue and Rabin (2003, 2006) study the optimal paternalistic tax, when consumption exerts a negative utility on future periods. Arguably, the regulation of information is—at least if information is freely available—a much softer form of paternalism than sin taxes. The provided information with its implicit consumption recommendation can be interpreted as a default action, which the consumer may or may not follow. If we think of information being available at some very small cost, my model is more in the spirit of libertarian paternalism, and relates to the example in Thaler and Sunstein (2003), where the planner of a cafeteria may place dessert in a further location to induce small transaction costs for its consumption.

Based on the work of Kamenica and Gentzkow (2011) and Rayo and Segal (2010), there has been a quickly evolving literature on Bayesian persuasion and information design. Yet, very little is known about optimal sequential information design by conflicting parties. Li and Norman (2017) study sequential persuasion with multiple senders. Similar to my model, attention can be restricted to equilibria in which information is only provided in the first period. Terstiege and Wasser (2017) analyze buyer-otpimal information structures in a monopoly that are robust to additional information provision by the seller. To my best knowledge, this paper provides the first dynamic model in which the receiver himself may acquire addition information before choosing an action.

3 The Model

The consumption model with intertemporal preferences closely follows Carrillo and Mariotti (2000). Consider an infinite discrete time model, indexed by $t = 0, 1, 2, ...$ In every period $t \geq 1$, a risk-neutral consumer decides whether he wants to consume one unit of an indivisible good. Consumption induces an instantaneous utility normalized to one. Let therefore $x_t \in \{0, 1\}$ denote the consumption choice at time t.[4]

Consumption exerts a negative externality on the welfare of future periods. More precisely, consumption at time t reduces the consumer's utility in period $t + \tau$ with probability θ by an amount $c_\tau \in [0, \overline{c}]$ for all $\tau \geq 1$. In particular, the magnitude of the externalities is assumed to be independent of past consumption choices.

The probability θ of a realization of the externality is unknown to the consumer. It is distributed according to some prior distribution with cdf F_0 and continuous, positive density $f_0(\theta)$ on support $[0, 1]$.[5] However, at time 0 a regulator on behalf of the consumer may design an information signal about θ, which the consumer observes at no cost. In the beginning of any subsequent period the consumer may acquire further information about θ. Details on the informational process are explained after the characterization of the intrapersonal conflict.

Intertemporal Payoffs and Intrapersonal Conflict

The instantaneous expected utility u_t at each date t consists of the potential consumption utility and the expected externality costs $c_{t-\tau}\theta$, acquired in former periods $\tau \in \{1, ..., t - 1\}$, i.e.

$$u_t = x_t - \sum_{\tau=1}^{t-1} x_\tau c_{t-\tau}\theta.$$

To abstract away from updating the value of θ due to the realization of the externality, I follow Carrillo and Mariotti (2000) in assuming that the consumer does not observe his current utility.

[4] As Carrillo and Mariotti (2000) point out, the restriction to binary decisions is without loss of generality. Indeed, whenever the consumer decides to consume, he wishes to consume the maximal amount.

[5] I follow Carrillo and Mariotti (2000) in assuming that the magnitude of the externality is known to consumers, whereas its probability of occurrence is unknown. It is straightforward to derive an equivalent model where the externality occurs with certainty, but its magnitude is unknown.

As a key assumption, I assume the consumer has present-biased preferences as developed by Phelps and Pollak (1968) and employed by Laibson (1996, 1997): The consumer at time t assigns a discount factor of $\beta\delta^{\tau}$ $(\beta, \delta < 1)$ to the instantaneous utility in period $t+\tau$ $(\tau \geq 1)$. Utility from the consumer's perspective at time t (in the following called "self-t") then reads

$$U_t = u_t + \beta \sum_{\tau=1}^{\infty} \delta^{\tau} u_{t+\tau}.$$

The parameter β can be regarded as the "impatience" or "impulsiveness" (Ainslie (1992)). In contrast to classical exponential discounting, this quasi-hyperbolic discounting leads to decisions that are time inconsistent: The optimal contingent plan of self-t for consumption in some future period $t+\tau$ may not longer be optimal to implement for self-$t{+}\tau$, as self-$t{+}\tau$ has a strong taste for instantaneous gratification.

Consequently, the collection of different selves of the consumer play a non-cooperative game against each other. To focus on this intra-personal conflict, I assume that the consumer has no commitment power towards his future selfs. The main scope of this paper is to analyze to which extend optimal information provision in the sense of Bayesian persuasion (Kamenica and Gentzkow (2011)) can mitigate the time-inconsistency problem and increase consumer's utility.

Besides his self control problem the consumer behaves fully rational: he perfectly anticipates his behavior and has perfect recall.

Information Provision and Learning

Consider now the role of a welfare maximizing regulator. Given the conflict between the desires of different consumer selfs, it is a priori unclear how to define an appropriate welfare function. I follow the approach suggested by O'Donoghue and Rabin (2002, 2003) and pursued by many others in focusing on the welfare of self-0. Thus, the regulator takes the perspective of the consumer before he faces any consumption decision.

Definition 3.1. A signal structure for θ consists of a finite signal space $S = \{s_1, ..., s_n\}$ together with a joint distribution G on the measurable space $([0,1] \times S, \mathcal{B}([0,1] \times S))$, where $\mathcal{B}([0,1] \times S)$ is the induced Borel algebra. Defining the joint distribution the usual way by $G(\tilde{\theta},\tilde{s}) = Pr(\theta \leq \tilde{\theta}, s \leq \tilde{s})$, we say S is consistent with belief F about θ, if for all $\theta \in [0,1]$ the marginal distributions satisfy

$$G(\theta, s_n) = F(\theta).$$

At time 0 the regulator may costlessly provide any F_0-consistent signal structure S_0 about the risk type θ. Observing a signal realization s, the consumer forms posterior belief F_1 according to the conditional distribution $G(\cdot|s)$. We can for example think of a costless health check which reveals some information about individual health risk of smoking, or some legal requirements for product labelling, which provides the consumer with some (incomplete) information about the (un)healthiness of food products.

In many circumstances it seems plausible that the regulator has exclusive control over information in the sense that it is too costly for the agent to acquire additional information himself whenever he faces a consumption decision.[6]

In other environments the consumer may be able to instantaneously find relevant information online at (almost) no cost.

In the following analysis, I therefore consider two cases. First, I analyze the optimal signal structure whenever the regulator (or self-0) has full control over information. Then, I consider the case where the consumer may acquire additional information each period before the consumption decision. Formally, at time $t \geq 1$ the consumer may design any signal structure S_t that is consistent with the belief F_t derived from Bayesian updating in period $t-1$.

Whenever we think of the case of full information control by the regulator we formally restrict consumer's information acquisition at any time $t \geq 1$ to the trivial signal $S = \{s\}$.

Equilibrium Concept

In each period $t \geq 1$, a strategy for the consumer consists of a learning decision in form of a signal S_t, and the consumption decision $x_t \in \{0,1\}$, based on the

[6] As the regulator's interest is aligned with the consumer whenever the consumer does not face an instantaneous consumption decision, an alternative interpretation of the model would be that at time zero—say at home—the consumer himself has costless access to information about his risk type, whereas in the situation of consumption—say in the supermarket—such information would be excessively costly as it is not directly available.

posterior derived from updating with respect to the signal. Since the payoff relevant information at the beginning period t is captured by belief F_t, and is independent of time, it is natural to focus on Markov strategies.

Definition 3.2. A Markov strategy for the consumer prescribes for each belief F an F-consistent signal structure S_F, and for each signal realization $s \in S_F$ a consumption decision $x(F, s) \in \{0, 1\}$.

Certainly, as the consumption decision has no impact on future behavior, the consumer optimally consumes whenever consumption utility exceeds expected externality cost, given the updated belief from signal s. Formally, in any equilibrium we have $x(F, s) = 1$ if and only if

$$1 > \left(\sum_{\tau=1}^{\infty} \beta \delta^{\tau} c_{\tau} \right) \mathbb{E}[\theta | F, s].$$

Definition 3.3. We say a subgame perfect equilibrium of the prescribed game is a Markov perfect equilibrium (MPE), if the consumer's strategy for $t \geq 1$ forms a Markov strategy. We say a Markov perfect equilibrium is preferred, if it maximizes self-0's utility among all Markov perfect equilibria.

The solution concept to the game will be preferred Markov perfect equilibrium (PMPE). To ensure existence, the following assumption is maintained throughout the paper.

Assumption 3.4. The consumer breaks any tie in favor of the regulator: Whenever the consumer at any time is indifferent between preferred signal structures or the two consumption decision, he takes the decision that maximizes utility of self-0.

We will see that the unique PMPE arises naturally in this context: the regulator provides information which can be interpreted as a consumption recommendation. The consumer finds it optimal to follow the consumption recommendation and abstains from further information acquisition.

4 Analysis

Let

$$C = \sum_{\tau=1}^{\infty} \delta^{\tau} c_{\tau}$$

be the present value magnitude of the externality without present bias. The following condition is assumed to hold for the remainder of the paper.

Assumption 4.1. If the consumer knows with certainty that the externality realizes he prefers to abstain, i.e.

$$\beta C > 1.$$

The Full Learning Benchmark

In order to understand the benefit of incomplete information, it is insightful to first analyze benchmark where the consumer has complete information about θ.

From the perspective of self-0, the expected value of consumption at any time $t \geq 1$ is

$$\delta^t(\beta x_t - \beta x_t C\theta),$$

thus consumption is optimal if and only if $\theta \leq \frac{1}{C}$. However, the value of instantaneous consumption for self-t is

$$x_t - x_t \beta C\theta.$$

Consumption therefore is optimal for self-t if and only if $\theta \leq \frac{1}{\beta C}$. Hence, a conflict of interest between self-0 and self-t arises if and only if $\theta \in [\frac{1}{C}, \frac{1}{\beta C}]$, where self-$t$ will consume even though his past selves would have liked to commit to abstention. The loss due to the lack of commitment power is depicted in Figure 3.1.

The dashed line depicts the self-t's utility from his optimal consumption decision in period t to consume whenever $\theta \leq \frac{1}{C}$. The solid line represents self-0's (δ^t-undiscounted) utility of self-t's decision. From the perspective of self-0, the shaded area illustrates the loss of self-t's action compared to self-0's preferred action.

From an ex-ante perspective total expected utility under complete information is

$$\mathbb{E}U_{\text{full info}} = \sum_{t=1}^{\infty} \beta\delta^t \int_0^{\frac{1}{\beta C}} (1 - C\theta) f(\theta) d\theta = \frac{\delta\beta}{1-\delta} \int_0^{\frac{1}{\beta C}} (1 - C\theta) f(\theta) d\theta.$$

Compared to the first-best utility

$$\mathbb{E}U_{\text{first best}} = \sum_{t=1}^{\infty} \beta\delta^t \int_0^{\frac{1}{C}} (1 - C\theta) f(\theta) d\theta = \frac{\delta\beta}{1-\delta} \int_0^{\frac{1}{C}} (1 - C\theta) f(\theta) d\theta,$$

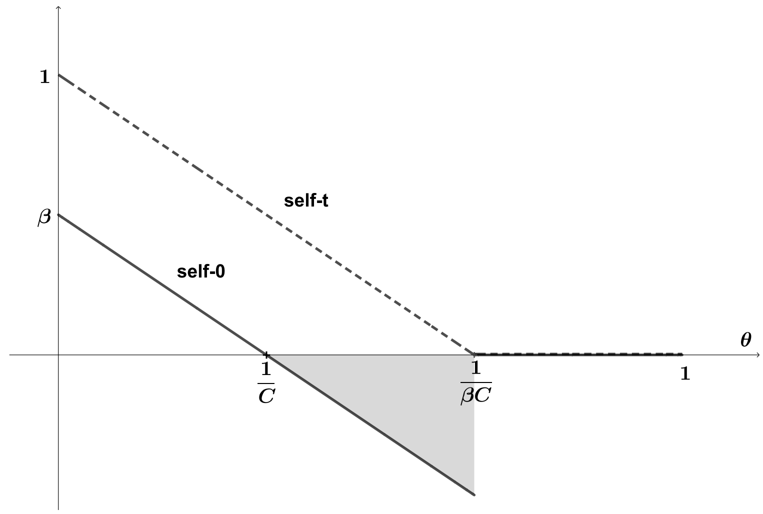

Figure 1: Welfare loss under complete information

the commitment problem induces a welfare loss of

$$\text{Loss} = \frac{\delta\beta}{1-\delta}\int_{\frac{1}{C}}^{\frac{1}{\beta C}}(1 - C\theta)f(\theta)d\theta.$$

In the following we will see how less information about θ may induce self-t to consumption choices that are more aligned with the interest of self-0.

Full Information Control

In this section, I analyze the equilibrium, when the regulator has full control over the information provided to consumers.

Since there is no information acquisition at any $t \geq 1$, and we restrict to Markov strategies, the consumption choices are history-independent and identical at all times. Consequently, the regulator faces a persuasion problem in the sense of Kamenica and Gentzkow (2011). Recall that there is a conflict of interest between the regulator and consumer's self $t \geq 1$ if and only if $\theta \in [\frac{1}{C}, \frac{1}{\beta C}]$. In her position of an information sender, the regulator's objective is to let the consumer receive information, which induces him to abstain in the conflicting interval as much as possible.

The solution to the regulator's problem is to use a cut-off strategy. She will inform the consumer, whether his risk type θ is above or below some threshold y.

Proposition 4.2. *If the regulator has full control over information, the welfare maximizing signal structure is described by a threshold $y \in \left[\frac{1}{C}, \frac{1}{\beta C}\right)$, and the signal*

$$S = \begin{cases} s_1, & \theta < y, \\ s_2, & \theta \geq y. \end{cases}$$

The consumer consumes if and only if $\theta < y$.

1. *If*
$$\mathbb{E}\left[\theta \middle| \theta > \frac{1}{C}\right] \geq \frac{1}{\beta C}$$
then $y = \frac{1}{C}$, and the signal induces first best welfare.

2. *If*
$$\mathbb{E}\left[\theta \middle| \theta > \frac{1}{C}\right] < \frac{1}{\beta C}$$
then y is uniquely determined by the condition that $\mathbb{E}\left[\theta | \theta > y\right] = \frac{1}{\beta C}$. Welfare under S strictly exceeds welfare under full information.

As the consumer abstains if and only if his expected risk θ weakly exceeds $\frac{1}{\beta C}$, the regulator pools as many types as possible from conflicting interval $\theta \in [\frac{1}{C}, \frac{1}{\beta C}]$ with high risk types by maintaining an expected risk weakly above $\frac{1}{\beta C}$. Within this interval the regulator prefers to induce abstention for the high types for two reasons: Firstly, for those types consumption creates the highest disutility to self-0. Secondly, the regulator is able to pool more types while maintaining an expected risk above $\frac{1}{\beta C}$.[7]

The optimal signal can be interpreted as an incentive compatible recommendation by the regulator to the consumer. Given consumer's risk type, the regulator recommends whether to consume or abstain, and the consumer wishes to follow the recommendation. It is therefore consistent with commonly observed health recom-

[7]The reason why it suffices to have only two signal realizations is very similar to the logic revelation principle. Whenever different signal realizations lead to the same action, we can instead use the signal structure, where the consumer cannot distinguish among these, without changing his optimal action. One can therefore assume without loss of generality that any signal structure has (at most) as many states as the action space.

mendations by experts as well as certified product labels which require products to satisfy legally defined limits on harmful substances.

To better understand the conflict of interest between self-0 and self-t, I derive the information signal at time 0 that is not preferred by self-0 but by self-t for $t \geq 1$. Again, as all consumer incarnations have the same information, any Markov strategy prescribes the same consumption decision to all incarnations. The following Lemma derives the preferred consumption decision of self-t if the decision has to be the same at all times t.

Lemma 4.3. *If all consumers have to take the same consumption decision rule $x(\theta) \in \{0,1\}$, then the optimal decision rule from the perspective of self-t for $t \geq 1$ is*

$$x(\theta) = 1 \qquad \text{if and only if} \qquad \theta \leq \delta \frac{1}{C} + (1-\delta)\frac{1}{\beta C}$$

Recall that the first best consumption choice of self-0 was to consume if and only if $\theta \leq \frac{1}{C}$, while the optimal myopic decision under full information was to consume if and only if $\theta \leq \frac{1}{\beta C}$. The optimal consumption threshold is a weighted average between these two objectives as it trades off self-t's instantaneous gain from high consumption with the loss that all future incarnations will consume equally much. The higher the long-run discount factor δ the more weight the consumer puts on the long run utility. The lower δ, the stronger is the consumer's desire for instant gratification.

Following the argument of Proposition 4.2 with this objective, one immediately obtains

Corollary 4.4. *If self-t for $t \geq 1$ has full control over information provided by the regulator, the welfare maximizing signal structure is described by a threshold $y \in \left[\delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}, \frac{1}{\beta C}\right)$, and the signal*

$$S = \begin{cases} s_1, & \theta < y, \\ s_2, & \theta \geq y. \end{cases}$$

The consumer consumes if and only if $\theta < y$.

1. If

$$\mathbb{E}\left[\theta \middle| \theta > \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}\right] \geq \frac{1}{\beta C}$$

then $y = \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}$, and the signal implements self-t's preferred consumption decision rule.

2. If

$$\mathbb{E}\left[\theta \middle| \theta > \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}\right] < \frac{1}{\beta C}$$

then y is uniquely determined by the condition that $\mathbb{E}\left[\theta|\theta > y\right] = \frac{1}{\beta C}$. Welfare under S strictly exceeds welfare under full information.

Costless Learning

I now look at the case, where the consumer may learn additional information costlessly at any time.

I begin the analysis with the observation that the consumer never benefits from postponing the acquisition of relevant information.

Definition 4.5. Consider a Markov perfect equilibrium. An information signal S_F is relevant, if acquiring S_F when the belief is F at any time t induces different expected utility to self-t or self-0 than acquiring no information.

Lemma 4.6. *On equilibrium path in a Markov perfect equilibrium the consumer will not acquire relevant information signals at any time $t > 1$.*

The intuition for this result is straightforward. Any acquired information by self-t ($t \geq 1$) can only help to improve his consumption decision. The only incentive for self-t to remain nevertheless uninformed about θ is to "discipline" future selfs to take a more favorable consumption decision. If self-t anticipates that self-$t+1$ will acquire additional information anyway, self-t (weakly) prefers to acquire that information herself.

Note that the regulator and self-1 are indifferent between information provision in period 0 and information acquisition in period 1. We can therefore in the following restrict without loss of generality to equilibria where (on equilibrium path) information is solely provided in period 0.

Since no incarnation wants to know less than his successor, each incarnation will acquire full information himself whenever the successor would do so. As an immediate consequence, the time-independent strategy of always acquiring full information regardless of the current belief forms a Markov perfect equilibrium.[8]

Corollary 4.7. *For any subgame starting at any time $t \geq 0$ the Markov strategy of a full information signal*

$$S_F = \begin{cases} s_1, & \theta < \frac{1}{\beta C}, \\ s_2, & \theta \geq \frac{1}{\beta C}, \end{cases}$$

for all beliefs F constitutes a Markov perfect equilibrium.

Besides the full information equilibrium, there may be a plethora of other potential equilibria. In the following, I am looking for the regulator-preferred Markov perfect equilibrium, and argue that it arises natural the context of our model.

The ability to support information provision by the regulator without further learning as an equilibrium depends on the ability to punish deviations to this information policy by future incarnations. Since any punishment has to be sequentially rational, it turns out that the maximum punishment to deviations from the equilibrium path is given by the full information subgame.

Lemma 4.8. *For any $t \geq 0$, the Markov perfect equilibrium of the subgame starting time $t+1$ which minimizes the utility of self-t is given by the full information equilibrium, where each self-s for $s \geq t+1$ for every belief chooses a full information signal.*

The equilibrium in Proposition 4.2 where the regulator has full control over information can be sustained with costless consumer learning if the one time gain from acquiring a preferable information structure does not exceed the loss from being stuck in the full information equilibrium afterwards. Proposition 4.9 gives precise conditions when this is the case.

[8]Formally, full information cannot be attained, since the state space is continuous, whereas we restrict the signal space to be finite. However, since the action space is only binary, the full information outcome can be replication with a binary signal (see Kamenica and Gentzkow (2011)). Indeed, take the cut-off signal that displays s_1 if and only if $\theta \leq \frac{1}{\beta C}$. This signal induces the consumer to make the full information consumption choice to consume whenever $\theta > \frac{1}{\beta C}$ — independently of further information realization. In the following, whenever I refer to a full information signal, one may think of this signal.

Proposition 4.9. *1. There is a preferred Markov perfect equilibrium with costless learning in which the only learning takes place at $t = 0$. The regulator provides a cut-off signal*

$$S = \begin{cases} s_1, & \theta < y, \\ s_2, & \theta \geq y. \end{cases}$$

The consumer always consumes if $S = s_1$ and always abstains if $S = s_2$.

2. *The cut-off y and the consumption decisions coincide with those in Proposition 4.2, where the regulator has full control over information, if and only if the optimal cut-off in Proposition 4.2 satisfies*

$$\mathbb{E}\left[\theta \middle| \theta \in \left[y, \frac{1}{\beta C}\right]\right] \geq \delta \frac{1}{C} + (1-\delta)\frac{1}{\beta C}. \qquad (1)$$

3. *Otherwise, if condition (1) is not satisfied then y is uniquely determined by*

$$\mathbb{E}\left[\theta \middle| \theta \in \left[y, \frac{1}{\beta C}\right]\right] = \delta \frac{1}{C} + (1-\delta)\frac{1}{\beta C}.$$

Welfare is strictly higher than under full information and strictly higher than in the case of Corollary 4.4, where the consumer at time $t = 1$ chooses the signal structure.

The intuition for the threshold y is depicted in Figure 3.2.

On equilibrium path, the consumer does not exert learning at any t, and consumes according to signal S whenever $\theta < y$. If some consumer incarnation deviates and acquires additional information at some time t then all subsequent selfs $s > t$ will acquire full information, and accordingly consume whenever $\theta < \frac{1}{\beta C}$. The solid line represents the (δ-undiscounted) per period utility of self-t from such a consumption choice in all periods $s > t$. Compared to abstention for $\theta > y$ on equilibrium path, self-t obtains each period $s > t$ a discounted loss proportional to the light grey shaded area. He will find this deviation profitable only if the sum of these discounted losses are exceeded by the one time gain of the deviation in period t. The most profitable deviation is full information, as it allows the best informed consumption choice. The dashed line depicts self-t's utility from full information in period t. The one time gain for self-t compared to the utility on equilibrium path

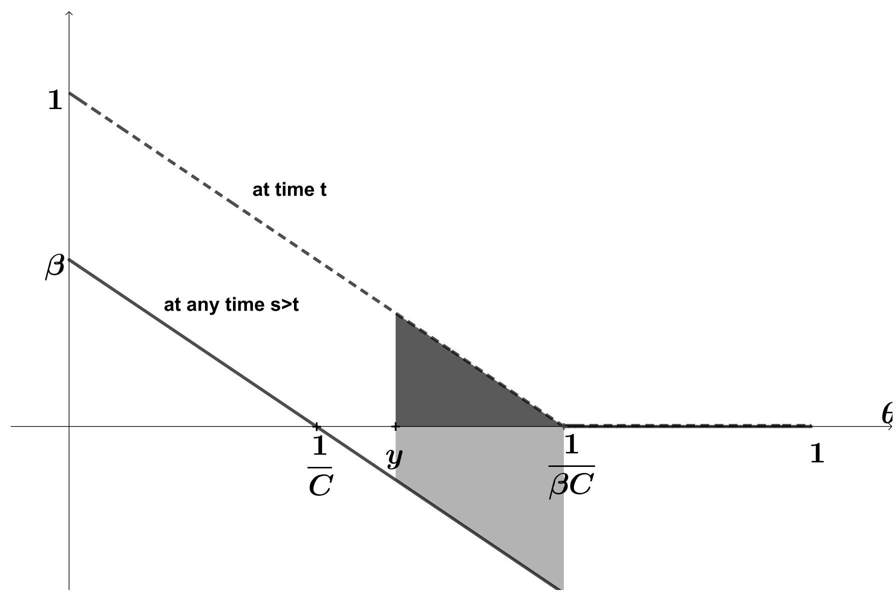

Figure 2: Gain and Loss of a Deviation from Self-t's Perspective

is depicted by the dark grey shaded area, where self-t consumes whenever $\theta < \frac{1}{\beta C}$.

While there are many other Markov perfect equilibria for the regulator's signal in period 0 (including the full information equilibrium), the described equilibrium is not only welfare maximizing, but also arises naturally in this context. The information signal in period 0 can be regarded as an incentive compatible recommendation of a default option by the regulator: All consumer incarnations find it optimal to take the information as given and base their consumption decision on the implied recommendation.[9]

Next, we look at sufficient conditions for the distribution to satisfy condition 1 for the optimal cut-off y, so that the outcome under full information control and in the PMPE under costless learning coincide.

Recall from Proposition 4.2 that the optimal threshold y under full information control satisfies $y \in [\frac{1}{C}, \frac{1}{\beta C}]$. Consequently, $\mathbb{E}[\theta|\theta \in [y, \frac{1}{\beta C}]]$ is in $[\frac{1}{C}, \frac{1}{\beta C}]$ with the exact value depending on where the prior distribution on $[\frac{1}{C}, \frac{1}{\beta C}]$ has most of its mass. Intuitively, in order to satisfy

$$\mathbb{E}[\theta|\theta \in [y, \frac{1}{\beta C}]] \geq \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C},$$

[9] In this sense, by the choice of the information signal, the social planner coordinates consumers on one chosen equilibrium, as it is common in the mechanism design literature.

the density of the prior *must not be too fast decreasing* on $[\frac{1}{C}, \frac{1}{\beta C}]$. Note that the condition only depends on β and C in so far as they define the range $[\frac{1}{C}, \frac{1}{\beta C}]$ in which the density must not decrease too fast. The following Corollary puts a bound on δ for priors that are nondecreasing in this range, which includes the natural benchmark of a uniform prior.

Corollary 4.10. *If the prior is nondecreasing on $[\frac{1}{C}, \frac{1}{\beta C}]$, and $\delta \geq \frac{1}{2}$, then Condition 1 in Proposition 4.9 is satisfied and the PMPE with costless learning coincides with the equilibrium with full informational control as described in Proposition 4.2.*

As one can interpret δ as the discount due to a common market interest rate between two consumption decisions, this condition is easily satisfied.

More generally, a sufficiently high δ always relaxes Condition 1, made precise in the following Corollary.[10]

Corollary 4.11. *For any prior distribution there exists a $\overline{\delta} < 1$ such that for all $\delta \in [\overline{\delta}, 1]$ Condition 1 in Proposition 4.9 is satisfied and the PMPE with costless learning coincides with the equilibrium with full informational control as described in Proposition 4.2.*

5 Conclusion

In this paper I showed how a paternalistic social planner can use information design to persuade consumers with quasi-hyperbolic preferences to better consumption decisions. The optimal information signal takes the remarkably simple form of a cut-off signal, which can be easily implemented by the use of threshold-based recommendations or certified labels.

The results provide intuitive benchmarks for the novel perspective of using information design as a paternalistic policy.

The are several natural ways to extend this baseline model, including information cost for information acquisition, consumer naïveté about their present bias, or consumer heterogeneity in the degree of present bias.

Altogether, this chapter may be regarded as a starting point for interesting future research.

[10]Note that this is not immediate as the value of the externality C depends on δ.

6 Appendix

Proof of Proposition 4.2. Let S be the signal structure chosen by the regulator, let $G(\theta, s)$ be its joint distribution on $[0,1] \times S$. Since the consumer has no access to further information at $t \geq 1$, a strategy for the consumer consists of a sequentially rational consumption decision $x(s) \in \{0,1\}$ for all signal realizations $s \in S$. Sequential utility maximization requires

$$x(s) = \begin{cases} 1, & \beta C \mathbb{E}[\theta | s] < 1, \\ 0, & \beta C \mathbb{E}[\theta | s] \geq 1. \end{cases}$$

The regulator's problem is to find an F_0-consistent signal structure (S, G) which maximizes total discounted utility from the consumer's consumption decision

$$U_0((S,G)) = \frac{\beta\delta}{1-\delta} \int_{(\theta,s)\in[0,1]\times S} x(s)(1 - C\theta) dG(\theta, s). \tag{2}$$

This problem is a classical persuasion problem as defined in Kamenica and Gentzkow (2011) with continuous state space $[0,1]$. In Proposition 3 of their Web Appendix they show that an optimal signal exists for such persuasion problems. Further, according to their Proposition 1, we can restrict to information signals $S = \{s_1, s_2\}$, where realization s_1 induces consumption, while realization s_2 induces abstention. Let from now $S = \{s_1, s_2\}$ with distribution $G(\theta, s)$ be an optimal F_0-consistent signal structure.

First, we show that S can be described by a cut-off $y \in [\frac{1}{C}, \frac{1}{\beta C})$ such that the realization is s_1 if and only if $\theta < y$. Suppose this is not the case. Due to the strictly positive density of the prior distribution, there exists by the intermediate value theorem some unique $y \in [0,1]$ such that $\mathbb{P}(S = s_2) = \mathbb{P}(\theta > y)$. We show that the cut-off signal which displays s_1 if and only if $\theta \leq y$ improves upon S.

Since

$$\begin{aligned} \mathbb{P}(S = s_2, \theta \leq y) &= \mathbb{P}(S = s_2) - \mathbb{P}(S = s_2, \theta > y) \\ &= \mathbb{P}(S = s_2) - \big(\mathbb{P}(\theta > y) - \mathbb{P}(S = s_1, \theta > y)\big) \\ &= \mathbb{P}(S = s_1, \theta > y), \end{aligned}$$

we have

$$
\begin{aligned}
\mathbb{E}[\theta|S=s_2] &= \frac{1}{\mathbb{P}(S=s_2)} \int_{\theta\in[0,1],S=s_2} \theta dG(\theta,s) \\
&= \frac{1}{\mathbb{P}(S=s_2)} \left(\int_{\theta\le y,S=s_2} \theta dG(\theta,s) + \int_{\theta>y,s\in S} \theta dG(\theta,s) - \int_{\theta>y,S=s_1} \theta dG(\theta,s) \right) \\
&< \frac{1}{\mathbb{P}(S=s_2)} \left(\mathbb{P}(\theta\le y, S=s_2)y + \int_{\theta>y,s\in S} \theta dG(\theta,s) - \mathbb{P}(\theta>y,S=s_1)y \right) \\
&= \frac{1}{\mathbb{P}(S=s_2)} \int_{\theta>y,s\in S} \theta dG(\theta,s) \\
&= \frac{1}{\mathbb{P}(\theta>y)} \int_{\theta>y} \theta dF_0(\theta) \\
&= \mathbb{E}[\theta|\theta>y].
\end{aligned}
$$

Consequently, the consumer abstains for $\theta > y$ under the cut-off signal, whenever he abstains for $S = s_2$ under signal S. Analogously, one can show that since $\mathbb{E}[\theta|S = s_1] > \mathbb{E}[\theta|\theta \le y]$, the consumer consumes for $\theta \le y$ under the cut-off signal whenever he consumes for $S = s_1$ under signal S. Plugging this decision rule into the regulator's objective (2), and using $\mathbb{E}[\theta|S = s_1] > \mathbb{E}[\theta|\theta \le y]$, we see that the regulator's utility under the cut-off signal

$$
\begin{aligned}
U_0 &= \frac{\beta\delta}{1-\delta} \int_{\theta\le y} (1-C\theta) dF_0(\theta) \\
&= \frac{\beta\delta}{1-\delta}\, \mathbb{P}(\theta\le y)\, \mathbb{E}[1-C\theta|\theta\le y] \\
&= \frac{\beta\delta}{1-\delta}\, \mathbb{P}(S=s_1)\, \big(1-C\mathbb{E}[\theta|\theta\le y]\big) \\
&> \frac{\beta\delta}{1-\delta}\, \mathbb{P}(S=s_1)\, \big(1-C\mathbb{E}[\theta|S=s_1]\big) \\
&= \frac{\beta\delta}{1-\delta} \int_{(\theta,s)\in[0,1]\times\{s_1\}} x(s)(1-C\theta) dG(\theta,s) \\
&= U_0((S,G))
\end{aligned}
$$

exceeds her utility under S.

We have shown that the optimal signal is a cut-off signal. To determine the optimal threshold y, recall that the regulator prefers abstention for any $\theta \in [\frac{1}{C}, \frac{1}{\beta C})$, and the consumer abstains for all $\theta > y$ if and only if $\mathbb{E}[\theta|\theta > y] \ge \frac{1}{\beta C}$. Therefore, the optimal y satisfies

$$\min\left\{y \in \Big[\frac{1}{C}, \frac{1}{\beta C}\Big)\Big| \mathbb{E}[\theta|\theta > y] \geq \frac{1}{\beta C}\right\}.$$

Consequently, whenever $\mathbb{E}[\theta|\theta \geq \frac{1}{C}] \geq \frac{1}{\beta C}$ the constraint is not binding and we get the boundary solution $y = \frac{1}{C}$. Otherwise the constraint binds, thus $\mathbb{E}[\theta|\theta > y] = \frac{1}{\beta C}$. Finally, since the optimal signal induces abstention on $[y, 1]$ with $y < \frac{1}{\beta C}$ whereas full information induces abstention on $\frac{1}{\beta C}$, self-0's utility under the optimal cut-off signal strictly exceeds his utility under full information. □

Proof of Lemma 4.3. Utility of self-t from consumption with risk type θ is

$$\begin{aligned} U_t(\theta) &= (1 - \beta C\theta) + \beta\delta(1 - C\theta) + \beta\delta^2(1 - C\theta) + ... \\ &= \left(1 + \frac{\beta\delta}{1-\delta}\right) - \frac{\beta}{1-\delta}C\theta. \end{aligned}$$

It follows that $U_t(\theta) \geq 0$ if and only if

$$\left(1 + \frac{\beta\delta}{1-\delta}\right) \geq \frac{\beta}{1-\delta}C\theta,$$

thus if and only if

$$\theta \leq (1-\delta)\frac{1}{\beta C} + \delta\frac{1}{C}.$$

□

Proof of Corollary 4.4. The proof is identical to the proof of Proposition 4.2, where we replace 'regulator' with 'self-1' and the regulator's objective $U_0((S, G))$ by self-1's objective

$$U_1((S, G)) = \int_{(\theta,s)\in[0,1]\times S} x(s)\left(\left(1 + \frac{\beta\delta}{1-\delta}\right) - \frac{\beta}{1-\delta}C\theta\right) dG(\theta, s).$$

In particular, the optimal signal is a cut-off signal, which aims to induce the most possible abstention for types $\theta \leq (1-\delta)\frac{1}{\beta C} + \delta\frac{1}{C}$ as calculated in Lemma 4.3. Consequently, the optimal cut-off y satisfies

$$\min\left\{y \in \Big[(1-\delta)\frac{1}{\beta C} + \delta\frac{1}{C}, \frac{1}{\beta C}\Big)\Big| \mathbb{E}[\theta|\theta > y] \geq \frac{1}{\beta C}\right\},$$

and the result follows. $\square$

Proof of Lemma 4.6. We start by defining the collapse of two signal structures. Let S be an F-consistent signal structure with distribution G. For some signal realization $s_k \in S$ let $\tilde{F} = G(\cdot|s_k)$ be the posterior distribution. Further, let $\tilde{S}$ be an $\tilde{F}$-consistent signal structure with distribution $\tilde{G}$. The collapse of S and $\tilde{S}$ is the signal structure with signal space $\overline{S} = (S \backslash \{s_k\}) \sqcup \tilde{S}$ and a joint distribution on $([0,1] \times \overline{S}, \mathcal{B}([0,1] \times \overline{S}))$ defined via

$$\mathbb{P}(\theta \leq \hat{\theta}, s = \hat{s}) = \begin{cases} \mathbb{P}_G(\theta \leq \hat{\theta}, s = \hat{s}), & \hat{s} \in S \backslash \{s_k\}, \\ \mathbb{P}_G(s = s_k)\mathbb{P}_{\tilde{G}}(\theta \leq \hat{\theta}, s = \hat{s}), & \hat{s} \in \tilde{S}. \end{cases}$$

Note that acquiring first S and then $\tilde{S}$ whenever the signal realization is $\tilde{s}$ is equivalent to acquiring the collapsed signal of S and $\tilde{S}$.

Consider now a Markov perfect equilibrium and denote for the consumer's Markov strategy with S_F the signal choice for belief F. (If the consumer decides not to learn for belief F, take S_F as the trivial signal consisting of only one state.)

Let $t > 1$. Since in equilibrium self-t with belief F_t acquires S_{F_t}, this implies that the consumer weakly prefers the distribution of posteriors from S_{F_t} to belief F_t. However, since self-$t-1$ acquires $S_{F_{t-1}}$ rather than the collapse of $S_{F_{t-1}}$ with S_{F_t} implies that the consumer weakly prefers belief F_t to the distribution of posteriors from S_{F_t}. It follows that the consumer is indifferent between belief F_t and the distribution of posteriors from S_{F_t}.

Consequently, self-$t-1$ with belief F_{t-1} is indifferent between acquiring $S_{F_{t-1}}$ and acquiring the collapse of $S_{F_{t-1}}$ with S_{F_t}, whereas self-t is indifferent between acquiring S_{F_t} and the trivial signal.

Suppose now S_{F_t} is relevant for $t > 1$. As self-t is indifferent between S_{F_t} and the trivial signal, this implies that self-0 is not. Since by Assumption 3.4 self-t always chooses self-0's preferred action whenever he is indifferent between his preferred action, self-0 strictly prefers self-t to acquire signal S_{F_t} rather than the trivial signal. This implies he prefers future incarnations to have the distribution of posteriors from S_{F_t} rather than belief F_t. In particular, self-0 prefers self-$t-1$ to acquire the collapse of $S_{F_{t-1}}$ with S_{F_t} rather than his equilibrium choice $S_{F_{t-1}}$.

Since self-$t-1$ is indifferent between the two, but chooses $S_{F_{t-1}}$, Assumption 3.4 is violated, a contradiction.

□

Proof of Lemma 4.8. By Corollary 4.7, the full information strategy for the subgame starting at $t+1$ is a Markov perfect equilibrium. Take any other Markov perfect equilibrium of the subgame. Lemma 4.6 states that for the game starting at $t=1$ with belief F_1 there is only information acquisition at $t=1$ in equilibrium. By renaming the time index it is immediate that for any subgame starting at $t+1$ with belief F_{t+1} there is only information acquisition at time $t+1$ in equilibrium. Consequently, the consumption decision is identical at all times starting at $t+1$. Let $S_{F_{t+1}}$ with distribution G be the information signal at $t+1$ and $x(s)$ be the consumption decision for signal realization s. Then the expected utility for self-$t+1$ generated by his own consumption decision under signal $S_{F_{t+1}}$ is

$$v_{t+1} = \int_{(\theta,s)\in[0,1]\times S} x(s)(1-\beta C\theta)dG(\theta,s),$$

whereas the undiscounted per-period utility for self-$t+1$ generated by all future selfs consumption decision is

$$v = \int_{(\theta,s)\in[0,1]\times S} x(s)(1-C\theta)dG(\theta,s).$$

Call accordingly

$$v_{t+1}^{FI} = \int_{\theta\in[0,1]} x(s)(1-\beta C\theta)dF_{t+1}(\theta)$$

and

$$v^{FI} = \int_{\theta\in[0,1]} x(s)(1-C\theta)dF_{t+1}(\theta)$$

the respective expected per-period utilities from full information for self-$t+1$ with belief F_{t+1}. Since the equilibrium strategy must give at least the same utility as deviating to full information and consuming the full information consumption level forever, we have

$$v_{t+1} + \beta(\delta v + \delta^2 v + ...) \geq v_{t+1}^{FI} + \beta(\delta v^{FI} + \delta^2 v^{FI} + ...).$$

Since full information enables self-$t+1$ to his best consumption choice we have $v_{t+1}^{FI} \geq v_{t+1}$ and therefore

$$\beta(\delta v + \delta^2 v + ...) \geq \beta(\delta v^{FI} + \delta^2 v^{FI} + ...).$$

Now, on the left-hand side we have the utility for self-t generated by the equilibrium, whereas on the right-hand side we have the utility for self-t generated by the full information equilibrium, which shows that no equilibrium for the subgame starting at $t+1$ can induce a lower utility to self-t than the full information equilibrium.

$\square$

Proof of Proposition 4.9. By Lemma 4.6 we can restrict to Markov equilibria with no information acquisition on equilibrium path at any time $t \geq 1$. Such an equilibrium induces the same consumption decisions for all incarnations of the consumer.

A necessary condition for a Markov strategy to be part of such a Markov perfect equilibrium different to the full information equilibrium is that self-1 does not benefit from deviating and acquiring a full information signal

$$S = \begin{cases} s_1, & \theta < \frac{1}{\beta C}, \\ s_2, & \theta \geq \frac{1}{\beta C}. \end{cases}$$

Such a deviation would yield self-1 an expected utility of

$$\overline{U} = \int_0^{\frac{1}{\beta C}} (1 - \beta C\theta) dF_1(\theta) + \frac{\beta\delta}{1-\delta} \int_0^{\frac{1}{\beta C}} (1 - C\theta) dF_1(\theta).$$

Note that $\overline{U}$ depends on the updated belief F_1, thus on the information realization in $t = 0$.

Hence, a solution to the relaxed problem, where the regulator maximizes her utility under the constraint that self-1's utility without further information acquisition weakly exceeds $\overline{U}$ for all realizations of the regulator's signal, puts an upper bound on the utility which the regulator can achieve in any Markov perfect equilibrium. We determine this upper bound and show how to implement it as a Markov perfect equilibrium.

First, note that for a solution to the relaxed problem we can again restrict attention to signals in $S_{F_0} = \{s_1, s_2\}$ where realization s_1 induces consumption and s_2 induces abstention. Indeed, if self-1 does not benefit from full information for any signal realization of a signal $S = \{s_1, ..., s_n\}$, then he does not benefit from full information in expectation for all states that induce consumption or abstention. Consequently, combining all realizations that induce consumption and all that induce abstention into one each yields a signal with two states for which self-1 does not benefit from full information.

Next, we show that a solution to the relaxed problem exists, if and only if it exists in the class of cut-off signals. Let S be a non-cutoff signal and let again $y \in [0, 1]$ be such that $\mathbb{P}(S = s_2) = \mathbb{P}(\theta > y)$. We showed in the proof of Proposition 4.2 that the cut-off signal which displays s_1 if and only if $\theta \leq y$ improves the regulator's utility compared to S and does not change consumer's consumption decision. Moreover this cut-off signal improves self-1's utility: The cut-off signal changes the consumer's action from abstention to consumption whenever $S = s_2$ and $\theta \leq y$. It changes the consumer's action from consumption to abstention whenever $S = s_1$ and $\theta > y$. Since

$$
\begin{aligned}
\mathbb{P}(S = s_2, \theta \leq y) &= \mathbb{P}(S = s_2) - \mathbb{P}(S = s_2, \theta > y) \\
&= \mathbb{P}(S = s_2) - \big(\mathbb{P}(\theta > y) - \mathbb{P}(S = s_1, \theta > y)\big) \\
&= \mathbb{P}(S = s_1, \theta > y),
\end{aligned}
$$

the same share of consumers consume under S and under the cut-off signal. Hence consumption utility is the same, however as lower risk types consume, the expected externality cost is lower. Consequently, self-1 finds it suboptimal to deviate to full information under the cut-off signal, whenever he finds it suboptimal under the signal S. This concludes the argument that for a solution to the relaxed problem we can focus on cut-off signal.

Before we determine the optimal cut-off as the solution to the relaxed problem, we show how such a cut-off signal in $t = 0$ can be implemented as a Markov Perfect equilibrium, in which no consumer acquires information at $t > 0$. Call $F_{\theta > y}$ and $F_{\theta < y}$ the posterior distributions from the cut-off signal, i.e. the two possible beliefs at $t = 1$. Consider the Markov strategy where the consumer acquires full information whenever he has a belief inconsistent with the regulator's signal, and no information otherwise, i.e.

$$S_F = \begin{cases} S(\theta) = \begin{cases} s_1, & \theta < \frac{1}{\beta C}, \\ s_2, & \theta \geq \frac{1}{\beta C}, \end{cases} & F \notin \{F_{\theta > y}, F_{\theta < y}\}, \\ \{s\}, & F \in \{F_{\theta > y}, F_{\theta < y}\}. \end{cases}$$

Together with the sequentially optimal consumption decisions (consume whenever the belief is $F_{\theta<y}$ or $F_{\theta<\frac{1}{\beta C}}$) this is indeed a Markov perfect equilibrium: Whenever the belief is not $F_{\theta>y}$ or $F_{\theta<y}$, each self-t anticipates that the next incarnation will acquire full information, so he finds it optimal to do so himself, as full information allows the best myopic consumption choice. Whenever the belief is $F_{\theta>y}$ or $F_{\theta<y}$, any information acquisition would end up in a different posterior and would induce full information next period. As the best myopic deviation would be full information, such a deviation would generate at most a utility of $\overline{U}$, and is therefore by assumption not improving upon the trivial signal.

Having established that the PMPE consists of a cut-off signal from the regulator and no consumer information acquisition on equilibrium path we now calculate the optimal cut-off y.

Consider a cut-off signal with cut-off y which induces abstention for $\theta > y$, i.e. $\mathbb{E}[\theta|\theta > y] \geq \frac{1}{\beta C}$. Such a signal induces a Markov perfect equilibrium for the above Markov strategy if and only if no consumer incarnation benefits from deviating to full information, ie. if and only if

$$\int_0^{\frac{1}{\beta C}} (1-\beta\theta C) dF_0(\theta) + \frac{\beta\delta}{1-\delta} \int_0^{\frac{1}{\beta C}} (1-\theta C) dF_0(\theta) \leq \int_0^{y} (1-\beta\theta C) dF_0(\theta) + \frac{\beta\delta}{1-\delta} \int_0^{y} (1-\theta C) dF_0(\theta),$$

or differently if and only if

$$\int_y^{\frac{1}{\beta C}} (1-\beta\theta C) dF_0(\theta) + \frac{\beta\delta}{1-\delta} \int_y^{\frac{1}{\beta C}} (1-\theta C) dF_0(\theta) \leq 0.$$

Rearraning this condition yields

$$\int_y^{\frac{1}{\beta C}} \left(1 + \frac{\beta\delta}{1-\delta}\right) dF_0(\theta) \leq \frac{\beta}{1-\delta} \int_y^{\frac{1}{\beta C}} \theta C dF_0(\theta).$$

Dividing by $F_0(\frac{1}{\beta C}) - F_0(y)$ gives us

$$1 + \frac{\beta\delta}{1-\delta} \leq \frac{\beta}{1-\delta}\mathbb{E}\left[\theta \middle| \theta \in \left[y, \frac{1}{\beta C}\right]\right] C,$$

which is equivalent to

$$\mathbb{E}[\theta|\theta \in [y, \frac{1}{\beta C}]] \geq \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}.$$

Since the regulator prefers abstention for all types $\theta > \frac{1}{C}$ the optimal cut-off therefore satisfies

$$\min\left\{y \in \left[\frac{1}{C}, 1\right] \;\middle|\; \mathbb{E}[\theta|\theta > y] \geq \frac{1}{\beta C}, \quad \mathbb{E}[\theta|\theta \in [y, \frac{1}{\beta C}]] \geq \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}\right\}.$$

Since the optimal cut-off from Proposition 4.2 where the consumer cannot acquire information satisfies

$$\overline{y} = \min\left\{y \in \left[\frac{1}{C}, 1\right] \;\middle|\; \mathbb{E}[\theta|\theta > y] \geq \frac{1}{\beta C}\right\},$$

our optimal cut-off coincides with the cut-off in the case where the consumer cannot acquire information if and only if

$$\mathbb{E}[\theta|\theta \in [\overline{y}, \frac{1}{\beta C}]] \geq \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C},$$

which shows 2.

Otherwise, if

$$\mathbb{E}[\theta|\theta \in [\overline{y}, \frac{1}{\beta C}]] < \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}$$

then the constraint $\mathbb{E}[\theta|\theta > y] \geq \frac{1}{\beta C}$ for the optimal cut-off is not binding. The optimal cut-off is then given by

$$\min\left\{y \in \left[\frac{1}{C}, 1\right] \;\middle|\; \mathbb{E}[\theta|\theta \in [y, \frac{1}{\beta C}]] \geq \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}\right\}.$$

Since for $y = \frac{1}{C}$ we have

$$\mathbb{E}[\theta|\theta \in [y, \frac{1}{\beta C}]] \leq \mathbb{E}[\theta|\theta \in [\overline{y}, \frac{1}{\beta C}]] < \delta \frac{1}{C} + (1-\delta)\frac{1}{\beta C},$$

the minimum is not at the boundary $\frac{1}{C}$, but has an inner solution satisfying

$$\mathbb{E}[\theta|\theta \in [y, \frac{1}{\beta C}]] = \delta \frac{1}{C} + (1-\delta)\frac{1}{\beta C}.$$

In particular, this implies that

$$y > \delta \frac{1}{C} + (1-\delta)\frac{1}{\beta C},$$

where the right hand side is by Corollary 4.4 the self-1 preferred cut-off. Thus, abstention is higher on $[\frac{1}{C}, \frac{1}{\beta C}]$ than under the self-1 preferred signal from Corollary 4.4 and induces higher welfare. $\square$

Proof of Corollary 4.11. By assumption, the prior has a continuous and strictly positive density f_0, thus attains its minimum $f_0^{\min}$ and its maximum $f_0^{\max}$ on $[0,1]$. Let $\overline{\delta} = \frac{f_0^{\max}}{f_0^{\max} + f_0^{\min}}$. Rearranging yields

$$\frac{f_0^{\max}}{f_0^{\min}} = \frac{\overline{\delta}}{1-\overline{\delta}}.$$

To save notation denote in the following $\mathbb{E}[\theta|\theta \in [\frac{1}{C}, \frac{1}{\beta C}]]$ with $\mathbb{E}$. Now, we have

$$\begin{aligned}
& \mathbb{P}\left(\theta \in \left[\frac{1}{C}, \frac{1}{\beta C}\right]\right)\mathbb{E} = \int_{\frac{1}{C}}^{\frac{1}{\beta C}} \theta f_0(\theta) d\theta \\
\Leftrightarrow \quad & \int_{\frac{1}{C}}^{\mathbb{E}} (\mathbb{E} - \theta) f_0(\theta) d\theta = \int_{\mathbb{E}}^{\frac{1}{\beta C}} (\theta - \mathbb{E}) f_0(\theta) d\theta \\
\Rightarrow \quad & \int_{\frac{1}{C}}^{\mathbb{E}} (\mathbb{E} - \theta) f_0^{\max} d\theta \geq \int_{\mathbb{E}}^{\frac{1}{\beta C}} (\theta - \mathbb{E}) f_0^{\min} d\theta \\
\Leftrightarrow \quad & \frac{\overline{\delta}}{1-\overline{\delta}} \int_{\frac{1}{C}}^{\mathbb{E}} (\mathbb{E} - \theta) d\theta \geq \int_{\mathbb{E}}^{\frac{1}{\beta C}} (\theta - \mathbb{E}) d\theta \\
\Leftrightarrow \quad & \frac{\overline{\delta}}{1-\overline{\delta}} \frac{\mathbb{E} - \frac{1}{C}}{2} \geq \frac{\frac{1}{\beta C} - \mathbb{E}}{2} \\
\Leftrightarrow \quad & \mathbb{E} \geq \overline{\delta}\frac{1}{C} + (1-\overline{\delta})\frac{1}{\beta C}.
\end{aligned}$$

Since $y \geq \frac{1}{C}$, this implies that for all $\delta > \overline{\delta}$ we have

$$\mathbb{E}\left[\theta\middle|\theta \in \left[y, \frac{1}{\beta C}\right]\right] \geq \mathbb{E} \geq \overline{\delta}\frac{1}{C} + (1-\overline{\delta})\frac{1}{\beta C} \geq \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}$$

□

Proof of Corollary 4.10. If the prior is nondecreasing on $[\frac{1}{C}, \frac{1}{\beta C}]$, then the conditional distribution of the prior on $[\frac{1}{C}, \frac{1}{\beta C}]$ first order stochastically dominates the uniform distribution on $[\frac{1}{C}, \frac{1}{\beta C}]$. Hence, for all $\delta \geq \frac{1}{2}$ we have

$$\mathbb{E}\left[\theta\middle|\theta \in \left[y, \frac{1}{\beta C}\right]\right] \geq \mathbb{E}\left[\theta\middle|\theta \in \left[\frac{1}{C}, \frac{1}{\beta C}\right]\right] \geq \frac{1}{2}\frac{1}{C} + \frac{1}{2}\frac{1}{\beta C} \geq \delta\frac{1}{C} + (1-\delta)\frac{1}{\beta C}.$$

□

References

Ainslie, G. (1992). *Picoeconomics: The strategic interaction of successive motivational states within the person*. Cambridge University Press.

Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics*, 117(3):871–915.

Brocas, I. and Carrillo, J. D. (2000). The value of information when preferences are dynamically inconsistent. *European Economic Review*, 44(4-6):1104–1115.

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

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic literature*, 47(2):315–72.

DellaVigna, S. and Malmendier, U. (2006). Paying not to go to the gym. *American Economic Review*, 96(3):694–719.

Diamond, P. and Kőszegi, B. (2003). Quasi-hyperbolic discounting and retirement. *Journal of Public Economics*, 87(9-10):1839–1872.

Frederick, S., Loewenstein, G., and O'donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of economic literature*, 40(2):351–401.

Kalinowski, A. and Humphreys, K. (2016). Governmental standard drink definitions and low-risk alcohol consumption guidelines in 37 countries. *Addiction*, 111(7):1293–1298.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *The American Economic Review*, 101(6):2590–2615.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478.

Laibson, D. (1998). Life-cycle consumption and hyperbolic discount functions. *European economic review*, 42(3-5):861–871.

Laibson, D. I. (1996). Hyperbolic discount functions, undersaving, and savings policy. Technical report, Harvard University.

Li, F. and Norman, P. (2017). Sequential persuasion.

Mariotti, T., Schweizer, N., and Szech, N. (2018). Information nudges and self control.

O'Donoghue, T. and Rabin, M. (2002). Addiction and present-biased preferences. *Department of Economics, UCB*.

O'Donoghue, T. and Rabin, M. (2003). Studying optimal paternalism, illustrated by a model of sin taxes. *The American Economic Review*, 93(2):186–191.

O'Donoghue, T. and Rabin, M. (2006). Optimal sin taxes. *Journal of Public Economics*, 90(10-11):1825–1849.

Phelps, E. S. and Pollak, R. A. (1968). On second-best national saving and game-equilibrium growth. *The Review of Economic Studies*, 35(2):185–199.

Rayo, L. and Segal, I. (2010). Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987.

Terstiege, S. and Wasser, C. (2017). Buyer-optimal robust information structures.

Thaler, R. H. and Sunstein, C. R. (2003). Libertarian paternalism. *American economic review*, 93(2):175–179.

Wertenbroch, K. (1998). Consumption self-control by rationing purchase quantities of virtue and vice. *Marketing science*, 17(4):317–337.