

Incentivizing self-control effort

Original

Incentivizing self-control effort / Abrardi, Laura; Cambini, Carlo. - In: JOURNAL OF ECONOMIC BEHAVIOR & ORGANIZATION. - ISSN 0167-2681. - STAMPA. - 164:(2019), pp. 13-30. [10.1016/j.jebo.2019.05.030]

Availability:

This version is available at: 11583/2734632 since: 2019-08-26T11:41:34Z

Publisher:

Elsevier

Published

DOI:10.1016/j.jebo.2019.05.030

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

Elsevier postprint/Author's Accepted Manuscript

© 2019. This manuscript version is made available under the CC-BY-NC-ND 4.0 license
<http://creativecommons.org/licenses/by-nc-nd/4.0/>. The final authenticated version is available online at:
<http://dx.doi.org/10.1016/j.jebo.2019.05.030>

(Article begins on next page)

Incentivizing Self-Control Effort ^{*}

Laura Abrardi[†], Carlo Cambini [‡]

May 2019

Abstract

We analyze the determinants of the incentives that a paternalistic social planner should provide to improve the behavior of time-inconsistent consumers. We focus on the specific case in which consumers can reduce their future consumption by exerting self-control effort in non-observable and non-contractible activities. For example, they can reduce their future healthcare costs by practicing physical exercise. Differently from the results obtained by the sin taxes theory, we find that incentives that address the under-exertion of effort by present-biased consumers are always socially costly even when consumers are homogeneous. Moreover, we show that incentives are first increasing and then decreasing in the level of time-inconsistency, due to the ineffectiveness of incentive policies. Incentives are also affected by the degree of naivety in a non-monotonic way.

Keywords: Optimal paternalism, time-inconsistency, sin taxes, present-bias.

JEL codes: D11, D91

^{*}We would like to thank the editor, Prof. Daniela Puzzello, three anonymous referees, as well as Maria Bigoni, Christos Genakos, Uri Gneezy, Marit Hinnosaar, Giovanni Immordino, Elisabetta Iossa, David E.M. Sappington, Yossi Spiegel, Dmitri Vinogradov, the participants in the International Industrial Organization Conference (Boston) and the iCare conference (Perm, Russia) for the valuable discussions and suggestions made to previous versions of the paper.

[†]Politecnico di Torino, DIGEP. Email: laura.abrardi@polito.it.

[‡]Politecnico di Torino, DIGEP; and EUI. Email: carlo.cambini@polito.it.

1 Introduction

One of the fundamental tenets of the theory on sin taxes is that, when self-control problems lead to over-consumption of unhealthy items, such as cigarettes or fatty foods, people's choices can be corrected by raising the prices of these goods. With homogeneous consumers, a simple Pigouvian tax-and-transfer scheme can induce optimal consumption with no welfare loss, as tax proceeds are returned to consumers (O'Donoghue and Rabin, 2006).

However, present-biased consumption decisions might not be the only issue. In some cases, idleness should also be blamed. For example, our state of health might be the result of too much fatty food, as well as insufficient physical exercise. Having a healthy lifestyle requires self-control efforts such as practising regular exercise or following an healthy diet, which produce benefits in terms of improved health. These behaviors require no financial outlay, but they still generate a cost, which is mainly related to the efforts adopted to avoid wasteful actions or to engage in constructive endeavors. In this case, can a price mechanism correct the *effort* component of a consumer's behavior? And what are the social cost and the effectiveness of such a policy? The goal of this paper is to address these questions.

From a policymaking point of view, the correction of time-inconsistent effort decisions is both problematic and relevant. It is problematic because a social planner may discourage the excessive consumption of unhealthy food by introducing taxes on it, but he cannot directly incentivize people to practice physical activities, at least outside the realm of experiments. The crux of the problem is that habits, unlike consumption, are *not contractible* and *not observable*. Incentives can only be conditioned on the observable consequences of these habits (i.e. efforts)¹. Hence, whenever the consumer's present-bias pertains to a non-observable effort decision, the consumer's behavior cannot be incentivized directly. This leads to a dichotomy between the observable aspects of our present-biased behavior (such as consumption decisions) and the unobservable aspects (such as effort decisions).

Moreover, the fact that the consumer's present-bias pertains to a non-observable effort has a further implication: in many cases, a social planner can only intervene *after* the

¹For example, a social planner cannot directly induce people to practice physical activities, but he can indirectly incentivize them by offering rewards for weight loss milestones.

consequences of those efforts materialize. For example, rewards conditional to losing weight can only be paid after a weight loss has actually been achieved, and the effort necessary to achieve that goal has already been exerted. This leads us to why these policies are relevant. The possible asynchrony between exerting effort and obtaining the corresponding rewards implies that the effects of time-inconsistency can be particularly pervasive in the case of effort decisions. Hence, efforts constitute an even more fertile ground for present-bias than consumption decisions.

Although several types of incentives could be introduced to correct effort, such as subsidizing gym entrance or the purchase of training equipment, people might still hesitate, and thus some residual under-exertion of effort persists. This residual misbehavior could also be corrected by introducing incentives that increase the future benefits of the present effort. We focus on the specific case in which the benefits of present behavior accrue in the future, in order to highlight the main trade-offs that arise when incentive policies address the effort decisions of time-inconsistent agents.

Indeed, these types of incentives are increasingly being introduced by policymakers to motivate people to change their behavior, often as part of schemes aimed at reducing obesity, smoking, alcoholism and other harmful habits. In March 2010, the US Patient Protection and Affordable Care Act introduced the possibility for employers of varying the health insurance premiums employees are charged, by as much as 30 percent, on condition that certain health status factors, such as body mass index, tobacco use, physical fitness or activity levels, are met. In 2015, such workplace wellness programs were offered by 57% of US employers (Kaiser Family Foundation, 2015), and by more than 80% of those with more than 1,000 employees (Mattke et al., 2015). In the UK, more than one out of three Clinical Commissioning Groups are currently denying or delaying routine surgery under the National Health Service to smokers and obese patients (Royal College of Surgeons, 2016), thereby raising the access cost to healthcare services to those segments of the population that are likely affected by present-bias problems.

The popularity of these policies, that condition incentives on the future consequences of our behavior, raises some important questions about their costs and effectiveness. In this paper, we address these questions, which are inherently of great interest for policymakers.

Our aim is twofold. First, we establish whether the correction of a consumer's effort necessarily entails a welfare loss. Second, our analysis points out to what extent the level of incentives should be raised as a function of the levels of consumers' present-bias and of the awareness of their own present-bias.

In our model, the consumer's utility depends on the quantity consumed of a good, and on a non-contractible and non-observable effort. For example, our health depends on the quantity of the healthcare services that we purchase, and on the effort exerted to achieve healthy behaviors, such as exercising regularly. The effort allows the consumer to achieve a certain level of utility, but with a lower consumption, hence with a lower expenditure. As the benefits of effort (i.e., the savings in consumption) only materialize after the effort has been made, the effort and consumption decisions are taken sequentially in our model: in the initial stage, the social planner sets the price of the good, the consumer then decides on the effort (and pays the effort cost), and she then decides on the consumption, pays the price of the quantity consumed and achieves the utility. Given that the benefit of the effort, in terms of savings, is accrued after the effort has already been exerted and its cost has been paid, a present-biased consumer misbehaves by underexerting effort. In order to induce the consumer to improve her behavior, a paternalistic social planner should increase the value of these benefits. This is achieved by raising the price of the good whose consumption is reduced by the effort². Therefore, a trade-off takes place between the deadweight loss of the rise in price and the positive effect it has on correcting effort.

We find that, when misbehavior concerns an effort decision and incentives can only be conditioned on future behavior, some welfare loss is unavoidable, even in the presence of homogeneous consumers, as the correction of effort necessarily requires the distortion of a consumption decision. As a consequence, incentives to correct a present-biased effort are always socially costly, and the first best cannot be achieved. This result should be seen in relation to the literature on taxation of present-biased consumption decision (O'Donoghue and Rabin, 1999a, 1999b, 2003 and 2006; Haavio and Kotakorpi, 2011), from which we

²The savings in consumption, obtained thanks to the effort, represent the reward of good behavior; this way of modelling the reward (i.e., through a price distortion rather than through a transfer) allows a direct comparison to be made with the literature on sin taxes.

depart as we consider an untaxable effort, rather than a taxable commodity. This difference is at the origin of our inefficiency even in case of homogeneous consumers, which conversely does not emerge with present-biased, homogeneous consumption decisions. Another strand of literature which is close to the present work is the one on imperfect externality-correcting taxes (Sandmo, 1978; Wijkander, 1985; Cremer et al., 1998; Fullerton and Wolverton, 2005; Christiansen and Smith, 2012). With this literature we share the result that the imperfection of the tax instrument (for example, for the presence of untaxable commodities) causes a welfare loss. However, the fact that the tax is aimed at correcting -though imperfectly- an externality, implies that the tax is higher, the larger the externality; in other words, an externality-correcting tax might be inefficient, but it is always effective. Conversely, if the problem is to correct a present-biased behavior and incentives can only be conditioned on future behavior, we find that the optimal size of incentives should be non-monotonic in the degree of present-bias. In particular, it should be low when present-bias is both low and when it is high. Intuitively, a trade-off arises between the social cost of incentives and their effectiveness. A present-biased consumer has much to gain from an increase in self-control effort, but it may be more difficult to convince her to exert such an effort, whenever the benefits can only be collected in the future, which she values less because of the present bias. Given that, in this situation, the price distortion is inefficient *and* ineffective, policymakers might want to target very present-biased consumers less aggressively than less present-biased individuals, despite the fact that these consumers display severe misbehavior with high social costs. On the other hand, a social planner provides limited incentives to low present-biased consumers as the inefficiency, in their case, is less severe.

We then extend our analysis to encompass a long-term perspective, in order to study the role of the consumer's awareness of her own present-bias, and to shed light on the interplay between intertemporal behavior and incentives. We assume that efforts must be exerted in two different periods. The first-stage effort can be considered as being related to a "participation" decision, while the second-stage effort can be interpreted as being related to self-control decisions. For instance, the effort of attending a wellness programs is generally preceded by the effort of acquiring information and choosing the program. Both efforts are necessary to obtain future savings and are linked in a complementary way.

We distinguish between *sophisticated consumers*, who are well aware of their time-inconsistency, and *naive ones*, who, conversely, do not anticipate their self-control problems (see DellaVigna and Malmendier, 2004; O’Donoghue and Rabin, 2003). Naive consumers overestimate their future ability to exert self-control, and are hence more likely to sink the cost of the participation effort. This has both positive and negative effects on the optimal level of incentives. On the one hand, incentives have a larger outreach with more naive consumers on self-control efforts. On the other hand, naive consumers display higher drop-out rates, thus dissipating the participation effort. The relative weight of these two opposite forces depends on the degree of present-bias. For high levels of present-bias, the large distortion between the effort they naively expect to exert and the effort that they will actually exert implies that the second effect is dominant, hence incentives should be higher for naive consumers than for sophisticated ones. Conversely, for low levels of present-bias, the inefficiency that arises from the dissipation of the participation effort is low, and incentives should mainly be directed toward the correction of the self-control effort. This inefficiency is particularly severe in the case of sophisticated consumers, as a result of the limited outreach of the policy. Hence incentives should be higher for sophisticated consumers.

The rest of the paper is organised as follows. We briefly present some results from the literature on sin taxes in Section 2. We introduce the baseline model and characterise the optimal policy with a time-inconsistent consumer in Section 3. Section 4 extends the model by introducing a long-term perspective. In particular, we identify how incentives drive participation and self-control decisions, and study the role of consumer’s awareness on incentive design. Section 5 presents the conclusions. All the proofs and technical details are relegated to Appendix 1.

2 Related literature

This work contributes to the debate on the optimal policy that should be introduced in the presence of time-inconsistent consumers, and it is related to three main strands of literature.

The first strand of literature investigates the economic incentives, in the form of sin taxes, applied to consumption goods when consumers are present-biased (see, e.g., O’Donoghue and

Rabin, 1999a, 1999b, 2003 and 2006; Haavio and Kotakorpi, 2011). Notably, and differently from our case, in the literature on sin taxes the present-bias is typically studied within the framework of consumption decisions. When the consumer's utility depends on the consumed quantity alone, raising the price (and then returning the proceeds to the consumer by means of a fixed subsidy) fully corrects the distortion in the consumption decision caused by the consumer's present bias, and the first best level can be implemented. We add to this literature by modelling consumers' behavior in the form of an effort, which could be exerted before the consumption decision, and which has an impact on the demand for the goods. We show that some welfare loss cannot be avoided when the consumer's mistake pertains to an effort decision, as opposed to a consumption decision. We depart from the sin taxes literature also along a second direction. When a long-term horizon is considered, in the case of goods, a social planner could impose differentiated prices for any period, in order to address, and possibly resolve, any intertemporal externality (see, e.g., Beshears et al., 2005). Conversely, when the effort is not contractible, incentives cannot separately target efforts exerted at different times. Therefore, some intertemporal externality may arise, so that short-term incentives interact with the long-term ones. Section 4 focuses on the interplay between the effects of efforts exerted over different periods. We show that such an interaction has important policy implications on the design of the incentives that address efforts over different periods.

The second strand of literature examines imperfect externality-correcting taxes (Sandmo, 1978; Wijkander, 1985; Cremer et al., 1998; Fullerton and Wolverton, 2005; Christiansen and Smith, 2012). In this literature, a tax is required to correct an externality. If the tax instrument is subject to some imperfection which prevents its perfect targeting (for example, because of heterogeneous external effects in the presence of uniform taxes, or the impossibility to tax the externality-producing good or to observe directly individual consumption), the price distortion causes a welfare loss. In this respect, there is a similarity with our case, where the inefficiency is also caused by an imperfection (specifically, the impossibility to tax directly the effort). However, the fact that the imperfection originates from an externality rather than from the consumer's present-bias, has an important implication in terms of incentive design. In fact, in an externality framework, the optimal tax is monotonically

increasing in the level of externality. On the contrary, if there is present-bias and incentives can only be conditioned on future behavior, the optimal price distortion is non-monotonic in the degree of present-bias, owing to the ineffectiveness arising from it. In other words, with an externality, the imperfect targeting of the tax only makes it inefficient, but it remains effective in correcting the externality. The incentive offered to correct a present-biased effort is instead both inefficient (because it must be imposed on a related good) and ineffective in correcting the present-bias.

The third strand of literature pertains to the indirect taxes imposed when one sector, for example leisure, is not taxable (Lerner, 1970; Baumol and Bradford, 1970; Diamond and Mirrlees, 1971; Atkinson and Stiglitz, 1972, 1976). In this literature, the price distortion is usually required to raise some fixed revenues from taxes, and one of the typical conclusions is that some inefficiency necessarily arises because of the shifting of resources from the taxed to the untaxed sector (Lerner, 1970; Dixit, 1970; Baumol and Bradford, 1970). We take an entirely different perspective. In our model, the price distortion is necessary to correct a present-biased effort. Then, the shifting of consumption from the taxed to the untaxable sector (the effort) is welfare-improving, as it corrects a detrimental behavior. Indeed, in our model the taxable good should be taxed even though it is a perfect substitute of the untaxable good (i.e., the effort), thus providing a counter-result of Lerner's (1970).

3 Short-term incentives

In this section, we explore the main economic implications of the fact that a consumer's present bias pertains to a non-observable action (effort), rather than to an observable action (consumption). To this aim, we set up a simple two-period model to find the welfare-maximizing price of the consumption good that provides the right incentives to guide consumer's behavior in terms of effort. We focus on the specific case in which incentives can only be conditioned to the future consequences of the present effort. This allows us to understand the costs and benefits of those policies that address misbehavior by increasing the future benefits of the present effort.

3.1 The model

A market is characterized by a unit mass of identical consumers. The representative consumer seeks to maximize her generic utility function $V(x, y)$, with $V_x, V_y > 0$, $V_{xx}, V_{yy} < 0$ and, for technical reasons, $V_{xxx}, V_{yyy} \geq 0$ ³. The utility is a function of the consumed quantity x and of the effort y exerted in order to increase the efficiency of her consumption. The effort y allows the consumer to achieve a given level of utility with a lower consumption x , i.e. x and y are substitutes. To make our point in the most straightforward way, we assume that the utility function is $V(x, y) = V(x + y)$, where one additional unit of effort allows the consumer to save one unit of consumed quantity⁴. For example, x is the quantity of purchased healthcare services, while y is the effort exerted in healthy behavior (exercising regularly, having a healthy diet, etc.) that allows one to save on healthcare costs. This effort is neither observable nor contractible, and it entails an increasing cost $d(y)$ ⁵. We assume that $d(0) = 0$, $d_y > 0$ (the effort is costly), $d_{yy} > 0$ (the cost of the effort is convex), and, for technical reasons, $d_{yyy} \geq 0$ (which is a sufficient condition for a social planner's problem to be concave, see also Laffont and Tirole, 1993). The non-observability of effort introduces a key novelty with respect to the literature on sin taxes (O'Donoghue and Rabin, 2003), where utility only depends on the (observable) quantities consumed.

The production of the good entails a constant marginal cost c , which also includes the social cost imposed on society.

A social planner charges a uniform price p for the consumption x ⁶. Given the price, the consumer first decides on the effort y , and then on the quantity x to purchase; on the basis of these decisions, she achieves the utility $V(x + y)$.

³In the literature on sin taxes, the utility function is usually assumed separable. For example, O'Donoghue and Rabin (2003) assumed that $V(x, y, Z) = \rho \ln x - \gamma \ln y + f(Z)$, where $\rho, \gamma > 0$ are exogenous parameters and Z is the vector of other goods; Haavio and Kotakorpi (2011) assumed that $V(x, y, z) = v(x) - h(y) + z$, where z is composite goods.

⁴This assumption is made with no loss of generality (see footnote 11). The same function was used by Chu and Sappington (2013) to model the consumers' decisions on energy consumption.

⁵We remark that y is a pure effort net of any price effect: in our example on healthy exercise, this means that y is the self-control effort that we need to exert, *in addition* to the fact that, at the time of the effort decision, we also have to pay for a dietician's consultation or an entry ticket to the gym, which could be partially or totally subsidized by alternative policy tools.

⁶In this framework, a fixed fee would merely have redistributive purposes, by returning the proceeds of the production to consumers; therefore, we focus on the optimal per-unit price (the same approach is used in the literature on sin taxes; see, e.g., O'Donoghue and Rabin, 2003 and 2006).

The timing of the game is as follows:

1. in $t = 0$, the social planner announces the price p ;
2. in $t = 1$, the consumer decides on a level of effort $y \geq 0$ and pays the effort cost $d(y)$;
3. in $t = 2$, the consumer decides on the quantity consumed, x , and receives the net surplus $V(x + y) - px$ from the consumption.

In many real situations, our utility is the result of past efforts. For example, we observe a weight loss only after we have exerted effort in staying on a diet. In this respect, effort decisions provide a natural and fertile ground for present-bias, as incentives can only be conditioned on the observable outcomes of behavior, and these outcomes often materialize in the future⁷. Our timing, in which the effort and consumption decisions are sequential, allows us to model the asynchrony between effort decisions and their consequences.

A social planner seeks to maximize welfare, W , defined as the present value of the consumers' utility, net effort costs, and of the profit from production: $W = \delta (V(x + y) - px) - d(y) + \delta(px - cx)$, where δ is the time-consistent discount factor⁸. Following O'Donoghue and Rabin, (1999, 2006), we assume, with no loss of generality, that $\delta = 1$.

3.2 Benchmark: time-consistent consumers

The social planner's problem is solved backwards. In $t = 2$, given the effort y exerted in $t = 1$ and the price p , the quantity x demanded by a consumer maximizes the function $V(x + y) - px$:

$$V'(x + y) = p. \tag{1}$$

⁷Of course, some of the observable outcomes of effort could be concurrent to the effort decision. In this case, they could be incentivized directly. For example, policymakers could subsidize complements or tax substitutes that are consumed at the same time of the effort decision, as evidenced by the literature on indirect taxation with an untaxable sector. In order to highlight the main trade-offs that could arise when policies address effort, we normalize to zero its present consequences, and focus only on the future ones. Under this respect, let us recall that y is the self-control effort that we need to exert, in addition to the fact that complementary goods such as a dietician's consultation or an entry ticket to the gym, are concurrently purchased and are (possibly) being subsidized by alternative policies.

⁸If the pricing scheme entailed a fixed fee plus a per-unit price, any fixed component of the two-parts tariff could be redistributed to the consumer by means of a lump sum transfer, thereby not affecting the expression of the welfare function.

It should be noted that a higher price reduces the gross surplus $V(x + y)$, as a result of the concavity of the utility function. From (1), the consumer's demand is

$$x^* = (V')^{-1}(p) - y. \quad (2)$$

From (2), the demanded quantity is a decreasing function of both the price and y , i.e. the self-control effort. Equation (2) clearly shows that y represents the quantity saved due to the effort for a given price p .

In $t = 1$, the consumer chooses the effort y by maximizing $V(x^* + y) - px^* - d(y)$. Using (2), the optimal solution of this problem is $d_y(y) = p$ and it allows us to obtain the first best effort:

$$y^* = d_y^{-1}(p). \quad (3)$$

The marginal benefit of exerting effort is represented by the value of the savings obtained through such an effort. This value is obtained by multiplying the quantity (i.e., 1) saved by one unit of effort by its value p . In equilibrium, the marginal benefit of the effort, p , is thus equal to its marginal cost $d_y(y)$.

It should be noted, from (3), that the optimal level of effort increases in p , i.e. the price works as an incentive for the consumer to increase her effort. Since the effort is a substitute for consumption, when the price of consumption increases, it is optimal to increase the quantity-saving effort.

The social planner's optimal policy maximizes welfare, and it entails a price equal to the marginal cost, i.e. $p^* = c^9$.

3.3 Time-inconsistent consumers

We now assume that the consumer has present-biased preferences, which the literature (starting from the seminal work of Laibson, 1997, and, more recently, from the work of O'Donoghue

⁹For a formal proof, see the proof of Proposition 1 in the Appendix 1.

and Rabin 1999a, 1999b and DellaVigna and Malmendier, 2004) usually represents through the simple functional form $U^t(u_t, \dots, u_T) = u_t + \beta \sum_{\tau=t+1}^T u_\tau$, where u_τ is the utility in period τ and the parameter β can be interpreted as the coefficient of short-term discounting. A time-consistent consumer has $\beta = 1$, while $\beta < 1$ represents a time-inconsistent preference for immediate gratification, and thus denotes short-term impatience. Since $\beta < 1$, a present-biased consumer excessively discounts future flows. This results in self-control problems and it leads to a bias between what the consumer does and what she thinks she should do, in a long-term perspective¹⁰. In this Section, we determine the incentives as a function of the degree of the consumer's present-bias. We thus assume that the level of β is observable; for example, it can be inferred from the consumers' state of health assessed by a medical examination. In Section 3.5, we will relax this hypothesis and consider the case in which two different types of consumers are present, with different degrees of present-bias, and the social planner cannot price-discriminate between different types (because the present-bias is not observable, or simply because price discrimination is not allowed).

Let us denote the results pertaining to time-inconsistent consumers with *inc*. Given the consumer's demand $x(p, y) = (V')^{-1}(p) - y$, the consumer chooses the effort y^{inc} in $t = 1$ in order to maximize her net surplus $\beta [V((V')^{-1}(p)) - px(p, y)] - d(y)$. The optimal solution is thus given by

$$d_y(y) = \beta p, \tag{4}$$

leading to the following optimal effort:

$$y^{inc} = y^{inc}(p, \beta) = d_y^{-1}(\beta p). \tag{5}$$

Equation (4) shows that, in the consumer's optimum, the marginal cost of effort $d_y(y)$ is equal to the marginal benefit βp . By comparing (5) with (3), it is easy to see that, given the same price p , the present-biased consumer under-exerts effort in $t = 1$, i.e. $y^{inc} < y^*$ for any

¹⁰A substantial body of evidence suggests that people exhibit dynamically inconsistent preferences and that discount functions are approximately hyperbolic (Ainslie, 1992). Because of this non-standard discount structure, a conflict arises between today's preferences, and the preferences that will be held in the future. As a consequence, in some contexts people make errors and fail to behave in their own best interests.

p . The reason for this is that the benefits of effort are accrued only after the effort is exerted, while its costs are immediate, therefore the benefits are underweighted by a present-biased consumer, and this leads to a lower effort than the first best.

Given a price p , the welfare function is given by $W = V((V')^{-1}(p)) - c((V')^{-1}(p) - y^{inc}) - d(y^{inc})$, hence the socially optimal price becomes:

$$p^{inc} = c \frac{d_{yy} - \beta V''}{d_{yy} - \beta^2 V''} > c, \quad (6)$$

for any $\beta \in (0, 1)$. In order to induce a present-biased consumer to exert a higher level of effort in $t = 1$, it is necessary to raise the price. The price here works as an incentive by increasing the marginal benefit of effort, i.e. the value of the quantity saved for each additional unit of effort¹¹. In fact, since $y^{inc} = d_y^{-1}(\beta p)$ and $y^* = d_y^{-1}(c)$, it would be sufficient to impose a price $p = c/\beta$ to achieve the first best level y^* of effort, thereby fully solving the misbehavior problem.

However, a rise in price does not come without cost. A price above the marginal cost also leads to underconsumption of x , given y , thus causing a welfare loss, given that the utility is $V((V')^{-1}(p))$. As a consequence, raising the price has negative side-effects on the consumers' surplus, *even when all the consumers display present-bias to the same extent*. If the consumer's bias involves misbehavior in the form of a non-observable action (effort) rather than in the form of an observable action (consumption), the correction of the distortion of effort y through the price necessarily entails some welfare loss, as stated in the following Proposition¹².

Proposition 1 *The first best cannot be achieved even with identical time-inconsistent consumers, i.e. $W(p^{inc}) < W(p^*)$.*

¹¹The fact that self-control effort and consumption are substitutes does not have an impact on our conclusions, but only on the direction of incentives (rewards instead of penalties). If, for example, the utility is expressed by the function $V(\min\{x, y\})$, the first best level of self-control effort can be achieved by imposing a price equal to $p = c - d_y(1/\beta - 1)$, which is lower than the marginal cost when $\beta < 1$.

¹²It should be noted that this result holds in general. Given a generic utility function $V(x, y)$, a present-biased consumer maximizes $V(x, \beta y) - p_x x - dy$, where p_x is the price of good x , thereby choosing x^{inc} and y^{inc} so that $V_x(x^{inc}, y^{inc}) = p_x$ and $V_y(x^{inc}, y^{inc}) = d/\beta$. When $V_{xy} \neq 0$, the price p_x corrects the choice of y^{inc} , but it also affects the choice of x^{inc} . Conversely, when $V_{xy} = 0$, the choice of y^{inc} is always biased, for any p_x . The first best cannot be achieved in either case ($V_{xy} \neq 0$ or $V_{xy} = 0$).

Proof. See Appendix 1. ■

A typical result in the literature on sin taxes is that if all consumers are affected by self-restraint problems to the same extent (i.e. the same β), then a price distortion allows the first best to be achieved, with no welfare loss¹³. In fact, if the consumer's bias only concerns a consumption decision, the price is merely a market instrument which can achieve the first best.

Proposition 1 shows that this is no longer true in a framework in which the consumer's mistake pertains to effort. If we account for the possibility of consumers exerting a quantity-saving, non-marketed effort, the price is not only a market instrument (through which the consumer decides on the quantity to consume), but it is also an instrument to incentivize good behavior. Therefore, a trade-off emerges between the welfare gain of improving the consumers' behavior and the welfare loss due to the higher price, even when consumers are homogeneous in terms of the time-inconsistency parameter β ¹⁴. Accounting for consumers' behavior in terms of effort dramatically affects the welfare effects of the policy, which has to be carefully calibrated to trade off an inefficient effort and an inefficient consumption.

The result of Proposition 1 closely recalls that obtained within the literature on imperfect externality-correcting taxes (Sandmo, 1978; Wijkander, 1985; Cremer et al., 1998; Fullerton and Wolverton, 2005; Christiansen and Smith, 2012). In the context of optimal taxation in the presence of externalities originated by non-taxable goods, such as leisure, Sandmo (1976) and Wijkander (1985) show that taxing complements (or subsidizing substitutes) for the externality-creating goods may result in inefficient distortions in a general equilibrium framework if complements and substitutes are related to other, non-polluting goods as well. Proposition 1 makes a similar case, with the inefficiency originating from a present-biased effort rather than from an externality. However, the fact that the price distortion is aimed at correcting -though imperfectly- an externality, implies that the tax is higher, the larger

¹³In fact, the concern expressed in the literature on optimal sin taxes is generally not about when all the consumers are equally present-biased, but when they are heterogeneous in their present-bias, so that helping irrational consumers is detrimental to rational ones. See, e.g., O'Donoghue and Rabin (2003, 2006, 2008) and Haavio and Kotakorpi (2011).

¹⁴It should be noted that inefficiency originates from the biased proportion of x and y , which in turn results from the solution to the consumer's problem. Therefore, it cannot be solved by adopting a different tariff structure for x . In fact, even if a two-part tariff were applied, instead of the per unit price p_x , the fixed component of the tariff would have no impact on the combination of x and y

the externality; in other words, an externality-correcting tax might be inefficient, but it is always effective. In the next Section, instead, we show that this result no longer holds in the case of a present-biased effort, as incentives are partially inefficient and ineffective.

3.4 Effect of consumer's present-bias on optimal incentives

In order to better understand how the optimal level of incentives depends on the degree of present-bias, we report the FOC of the welfare function w.r.t. p :

$$\frac{p - c}{V''((V')^{-1}(p))} + (c - d_y(y^{inc})) \frac{\partial y^{inc}}{\partial p} = 0 \quad (7)$$

The first term in equation (7) represents the marginal social cost caused by the price increase (i.e., the deadweight loss). In fact, a price $p^{inc} > c$ lowers the consumer's surplus, due to a reduction in the level of service by $1/V''$. It should be noted that the deadweight loss does not depend on the degree of present-bias β . It follows that, in order to understand the role β plays in the optimal price, we focus on the second term in equation (7), which represents the marginal effect of a price increase on the present-biased consumer's effort. This effect is given by i) the marginal value of effort (i.e., $c - d_y(y^{inc})$) and ii) the effectiveness of the policy, represented by the extent to which the effort increases due to the higher price (i.e., $\frac{\partial y^{inc}}{\partial p}$). Factors i) and ii) both depend on β . However, the degree of present-bias operates in an opposite direction.

First, present-bias has direct positive effects on the marginal value of effort. In fact, the marginal value of effort $c - d_y(y^{inc})$ decreases with $y^{inc} = d_y^{-1}(\beta p)$, which in turn is low when β is low (i.e., high present-bias). The problem of under-exertion of effort is more severe when the degree of present-bias is high, thus causing social inefficiency to become worse.

Second, $\frac{\partial y^{inc}}{\partial p} = \frac{\beta}{d_{yy}}$ increases in β : the higher the consumer's present-bias is, the less effective a price increase is in inducing behavioral change. Present-biased consumers are more hesitant about responding to incentives, as they can only be rewarded in the future (when the benefits of good behavior materializes), and thus discount them excessively.

The previous analysis suggests that the optimal level of incentives is non-monotonic

with respect to the degree of present-bias. The following Corollary formalises this result and shows that the optimal price distortion shows a maximum for internal levels of β , while it is zero for the two extreme cases of $\beta = 1$ (fully rational) and $\beta = 0$ (fully present-biased) consumers.

Corollary 1 *The optimal price p^{inc} is quasiconcave in β and $p^{inc} = c$ for $\beta = 0$ or $\beta = 1$.*

Proof. See Appendix 1. ■

It should be noted that the present-bias is necessary for a non-monotonicity result, as it is responsible for the trade-off between the benefits and effectiveness of incentives. This is in clear contrast with the case of a distortion caused by, for example, a negative externality. In such a case, the optimal incentive would always increase with the level of the distortion¹⁵.

To obtain a better idea of the result expressed by Corollary 1, let us suppose that $d(y) = \gamma y^2$ and $V(x + y) = (x + y) - (x + y)^2$, with $d_{yy} = 2\gamma > 0$. Then, $p^{inc} = c \frac{\gamma + \beta}{\gamma + \beta^2}$. The optimal price has an inverted U shape function in $\beta \in [0, 1]$ with a maximum for $\bar{\beta} = -\gamma + \sqrt{\gamma^2 + \gamma}$, where the two effects (in terms of value and effectiveness of the policy) balance each other. The price is low for $\beta < \bar{\beta}$, because of the limited effectiveness of the price increase, while the price is low for $\beta > \bar{\beta}$ because of the low value of the price increase. Interestingly, it should be noted that $\bar{\beta}$ is a function of the cost of the effort, and it increases in γ . Intuitively, a higher effort cost, γ , decreases the effectiveness of the policy, as $\frac{\partial y^{inc}}{\partial p} = \frac{\beta}{2\gamma}$, thus implying that the threshold level $\bar{\beta}$ shifts to the right.

3.5 Heterogeneous consumers

The previous section assumes that consumers are homogeneous in their present-bias. However, more realistically, consumers may be heterogeneous in their time-inconsistency. This heterogeneity calls for the adoption of person-specific prices which, however, may be difficult to implement in practice. We therefore assume that the social planner imposes a uniform

¹⁵This result can be obtained immediately by modelling a welfare function that differs from the consumer's objective function as a result of the introduction of an externality term as a function of the effort y . For example, assuming a fully rational consumer ($\beta = 1$), the welfare function could be expressed by adding the externality term $\psi(y)$ to the welfare function.

price, irrespectively of the different types of consumers. The goal of this section is to extend our baseline model in order to incorporate heterogeneity in consumers' present-bias and analyze how, in this scenario, a social planner should set a uniform optimal price.

Let us consider a population with two types of consumers: fully rational consumers, and consumers with self-control problems (with $\beta < 1$). Let us denote the share of present-biased consumers in the population with $\lambda \leq 1$ (hence, $1 - \lambda$ is the share of fully rational consumers). As previously shown, $y^{inc} = \frac{\beta p}{d_{yy}}$ is the effort exerted by the present-biased consumers, while $y^* = \frac{p}{d_{yy}}$ is the optimal effort of the fully rational ones. The welfare function is given by

$$W = V((V')^{-1}(p)) - \lambda [c((V')^{-1}(p) - y^{inc}) + d(y^{inc})] - (1 - \lambda) [c((V')^{-1}(p) - y^*) + d(y^*)],$$

hence the socially optimal price is:

$$p^{het} = c \frac{d_{yy}(y^{inc}) - \lambda \beta V'' - (1 - \lambda) V'' \Omega}{d_{yy}(y^{inc}) - \lambda \beta^2 V'' - (1 - \lambda) V'' \Omega}, \quad (8)$$

where $\Omega = \frac{d_{yy}(y^{inc})}{d_{yy}(y^*)}$. It should be noted that (8) is a generalization of (6), and $p^{het} > c$ for any $\lambda \in (0, 1)$. When consumers are all equally present-biased, $\lambda = 1$ and $p^{het} = c \frac{d_{yy}(y^{inc}) - \beta V''}{d_{yy}(y^{inc}) - \beta^2 V''} > c$; conversely, when all consumers are rational, $\lambda = 0$ and $p^{het} = c$. Whenever a share of consumers is not fully rational, the price should be distorted above the marginal cost, which is the optimal policy in the case of rational consumers. Therefore, the optimal policy in the presence of heterogeneous consumers involves the well-known trade-off between incentivizing present-biased consumers while harming rational ones (O'Donoghue and Rabin, 2003). Rational consumers are hurt in two ways. First, they suffer a deadweight loss because of the higher price. Second, they over-exert effort relative to the social optimum, which leads to a social waste of resources¹⁶. At the same time, incentives are less effective on present-biased consumers, since they are less sensitive to price incentives when benefits are accrued in the future. Therefore, an increase in the price might have a first-order cost on rational consumers, but only a second-order benefit on those with self-control problems. The consumers who respond the best to incentives are those who do not need them. This result

¹⁶The social benefit of effort is $cy - d(y)$, which is maximized for $y = d_y^{-1}(c)$. If $p > c$, the effort exerted by rational consumers is higher than $d_y^{-1}(c)$, hence the social benefit of effort is lower.




Figure 1: Welfare as a function of the share λ of present-biased consumers. Marginal costs equal to 1.

complements the one obtained by O'Donoghue and Rabin (2003), who showed that, in the case of present-biased consumption, it may be worth distorting the price, as the behavior of present-biased consumers can be corrected to a great extent.

To clarify this point, let us adopt the same functional forms used in Section 3.4, i.e. $d(y) = \gamma y^2$ and $V(x + y) = (x + y) - (x + y)^2$. The optimal uniform price in (8) becomes $p^{het} = c \frac{\gamma+1-\lambda+\lambda\beta}{\gamma+1-\lambda+\lambda\beta^2}$. When the present-bias is strong (i.e., β close to zero), the optimal price should not depart to any great extent from the marginal cost, regardless of the share of present-biased consumers in the population. In fact, the incentive is ineffective for strongly present-biased consumers (as $\frac{\partial y^{inc}}{\partial p} = \frac{\beta}{2\gamma}$), while it has an important impact on the effort exerted by rational consumers (as $\frac{\partial y^*}{\partial p} = \frac{1}{2\gamma}$). The larger the share λ of present-biased consumers is, the higher the price distortion is.

Figure 1 shows welfare corresponding to the optimal price p^{het} , for different levels of λ . It should be noted that the welfare in $\lambda = 0$ provides the first best benchmark, as all

consumers are rational. Two main issues emerge. First, a welfare loss always occurs for any $\lambda > 0$. Second, welfare does not monotonically decrease with the share of present-biased consumers, as could instead be expected. In fact, welfare is higher when the share of present-biased consumers is very high, rather than at intermediate levels. This result holds for all levels of present-bias, although it is more accentuated when the present-bias is high (i.e., when β is low). In other words, if 100% of consumers are present biased, welfare is higher than when, for example, only 90% of them are present-biased. Intuitively, if all consumers are equally present-biased, the price is optimal for all of them (i.e., p^{het} and p^{inc} coincide). Conversely, if some consumers are rational, heterogeneity imposes further costs on society, as the uniform price p^{het} lies somewhere in between the optimal price of rational consumers (c) and the optimal price of time-inconsistent consumers (p^{inc}). The fact that welfare is higher when all consumers are present-biased, rather than when only some of them are present-biased, suggests that the costs imposed by incentives on rational consumers may override the benefits for those who are present-biased.

4 Long-term incentives

The benefits of good behavior are often the result of a sequence of efforts. For example, we first exert the effort of joining a gym (acquiring information for choosing the health club, or obtaining the medical certificate), and then we exert the effort of attending the health program. Hence, before exerting the self-control effort, the consumer must participate in the program.

In this section, we set up a three-period model to analyze the long-term effects of incentives on behavior. We allow consumers to be time-inconsistent and (possibly) overconfident. In particular, we study the interaction between the efforts exerted at different stages and how this affects the optimal size of price incentives.

We assume that the consumer can reduce the quantity x by exerting effort in two previous stages. We denote the effort exerted in $t = 1$ by y , and we refer to it as a 'self-control effort', and the effort exerted in an initial stage $t = 0$ by z , and we refer to it as a 'participation effort'. Neither effort is observable, and they are assumed complementary. If the consumer does not

exert any effort, the utility is only a function of the quantity consumed x . Conversely, if the consumer exerted z , she has the possibility of exerting y later on. The initial effort z can be interpreted as a sort of participation effort with long-term effects, while y can be seen as a short-term effort for assiduously attending a venture by exerting self-control. Introducing an initial stage allows one to study the role of the consumer's awareness of her own present-bias (O'Donoghue and Rabin, 2001), and how this awareness affects the level of incentives necessary to correct the present-bias.

In order to streamline the analysis, we assume that the cost of z is constant and equal to k . The value of k is perfectly observed by the consumer, while a social planner only knows its distribution. It should be noted that k is, for example, the cost of acquiring information in order to choose which dietician or trainer to consult, or which sport is more suitable for one's specific needs, and not the cost of showing up or signing up to a wellness program: the latter are observable decisions and, as such, they can be incentivized directly through subsidies (Woerner, 2018). Conversely, our focus is on the non-observable components of behavior, such as the effort of acquiring information. Indeed, people typically sign up in a program but then fail to attend, as the empirical literature shows (DellaVigna and Malmendier, 2006). For simplicity, we assume that k is uniformly distributed over the $[0, K]$ interval, with $K > 0$.

The timing of the game is as follows:

1. in $t = 0$, the social planner announces the price p and the consumer privately observes k . The consumer then decides whether to exert z , or not and pays the relative cost k ;
2. in $t = 1$, the consumer decides on the level of effort $y \geq 0$ and pays its cost $d(y)$, if she previously exerted z ; if she did not exert z , $y = 0$;
3. in $t = 2$, the consumer decides on the quantity consumed, x , and receives the net surplus $V(x + y) - px$ from the consumption.

Let us first examine the first best benchmark. The consumer's demand x^* in $t = 2$ and the effort y^* in $t = 1$ are expressed by (2) and (3), respectively. In $t = 0$, the consumer only invests the effort cost k if the value of the net surplus obtained from participation is greater than that obtained from no participation:

$$V((V')^{-1}(p)) - p((V')^{-1}(p) - y^*) - d(y^*) - k \geq V((V')^{-1}(p)) - p(V')^{-1}(p), \text{ i.e.}$$

$$k \leq \epsilon^*, \tag{9}$$

where $\epsilon^* = py^* - d(y^*)$ (or, using (3), $\epsilon^* = pd_y^{-1}(p) - d(d_y^{-1}(p))$). Inequality (9) states that exerting the participation effort z is only socially efficient if its benefit py^* , in terms of quantities saved, is higher than its direct cost k plus its indirect cost $d(y^*)$ due to the self-control effort.

The interpretation of ϵ^* emerges directly from inequality (9). Given that the cost k of the participation effort is distributed uniformly over the range $[0, K]$, the consumer only exerts z if $k \leq \epsilon^*$; then, ϵ^* can be interpreted as the customer base, i.e. as the set of consumer types who decide to sink the participation effort.

Welfare can be expressed as

$$W = \int_0^{\epsilon^*} [[V((V')^{-1}(p)) - c((V')^{-1}(p) - y^*)] - d(y^*) - k] df(k) + \tag{10}$$

$$+ \int_{\epsilon^*}^K [V((V')^{-1}(p)) - c(V')^{-1}(p)] df(k),$$

where the first term is the welfare of the consumer types who participate in $t = 0$, while the second term is the welfare of the consumer types who do not participate. By maximizing expression (10) w.r.t. p , we find that the first best price corresponds to the marginal cost¹⁷.

Marginal cost pricing implies that the first best effort is $y^* = d_y^{-1}(c)$ and the consumer only participates if her cost k is lower or equal to the value $\epsilon^* = cd_y^{-1}(c) - d(d_y^{-1}(c))$.

These conditions represent our benchmark for the analysis of the following sections.

4.1 Time-inconsistent consumers

We now analyze the consumer's choice with particular attention to her behavior in the initial stage $t = 0$. As discussed by Strotz (1955) and by Phelps and Pollak (1968), the behavior of consumers with time-inconsistent preferences depends on their beliefs about their own

¹⁷see Appendix 1 for the proof.

future behavior. Two polar cases have been used in the literature: *sophisticated* agents, who have rational expectations about their future selves and therefore are fully aware of their self-control problems and correctly predict their future behavior, and *naive* agents (see also O'Donoghue and Rabin, 2001), who do not recognize that they cannot make consistent plans over time and therefore believe their future selves will behave exactly according to their long-term preferences. The consumers' unawareness is motivated by the experimental evidence on overconfidence about positive personal attributes (Larwood and Whittaker, 1977; Svenson, 1981) and is consistent with field evidence on investment (Madrian and Shea, 2001), task completion (Ariely and Wertenbroch, 2002) and health club attendance (DellaVigna and Malmendier, 2006).

Given that sophisticated and naive consumers differ, as far as their expectations about their future behavior are concerned, we denote with $\hat{\beta} \in [\beta, 1]$ the consumer's expectation in $t = 0$ about the discount factor β that her future self will apply in $t = 1$. The higher $\hat{\beta}$ is, the higher the consumers' degree of naivety. In fact, a sophisticated consumer is more aware of her time-inconsistency problem than a more naive consumer and anticipates her future behavior more precisely. Therefore, the sophisticated consumer's discount factor $\hat{\beta}$ is closer to β than the $\hat{\beta}$ of a more naive consumer¹⁸.

The belief $\hat{\beta}$ regarding one's own present-bias plays a crucial role in the consumer's decision in $t = 0$ about whether to participate or not. In fact, she will only exert z if the savings are sufficiently high, i.e. only if she expects to exert sufficient effort y in the future to cover the initial effort cost k .

In $t = 0$, the consumer expects that her effort will be $y_0^{inc} = y_0^{inc}(p, \hat{\beta}) = d_y^{-1}(\hat{\beta}p)$ in the next period. It should be noted that $y_0^{inc}(p, \hat{\beta}) \neq y^{inc}(p, \beta)$: the former is the effort which, in $t = 0$, the consumer believes she will exert in $t = 1$, while the latter is the effort that she will actually exert¹⁹.

¹⁸In Section 4.3, we perform a numerical analysis in which we consider the two extreme cases of a fully sophisticated consumer, with $\hat{\beta} = \beta$, and a fully naive one, with $\hat{\beta} = 1$.

¹⁹In order to simplify the notation, from now on we will not report the full functional forms for $y_0^{inc}(\cdot)$ and $y^{inc}(\cdot)$.

Given a generic price p , the consumer invests in the effort cost k in $t = 0$ only if

$$\beta [(V((V')^{-1}(p)) - p((V')^{-1}(p) - y_0^{inc})) - d(y_0^{inc})] - k \geq \beta [(V((V')^{-1}(p)) - p(V')^{-1}(p))]. \quad (11)$$

Through straightforward simplifications, equation (11) can be rewritten as $k \leq \epsilon^{inc}(p, \beta, \hat{\beta})$, where $\epsilon^{inc}(p, \beta, \hat{\beta}) = \beta (py_0^{inc} - d(y_0^{inc}))$. It should be noted that, as $d_y(y_0^{inc}) = \hat{\beta}p$, we have

$$\frac{\partial \epsilon^{inc}}{\partial p} = \beta \left(y_0^{inc} + p \frac{\partial y_0^{inc}}{\partial p} (1 - \hat{\beta}) \right) > 0. \quad (12)$$

A higher price increases the probability of a consumer participating, by shifting the position $\epsilon^{inc}(\cdot)$ of the type of the marginal consumer, i.e. the one whose effort cost k makes her indifferent about participating or not.

In this framework, the benefits of raising the price are twofold. First, the price has a direct positive effect on the effort y^{inc} , which descends from the fact that the price determines the value py of the quantity saved. Second, the price has a positive indirect effect on the consumer's participation in $t = 0$, which descends from the fact that the higher expected effort y_0^{inc} makes it worth sinking the cost k of the participation effort. Note also that if a fixed subsidy were introduced, in addition to the price incentive, it would be added to both sides of inequality (11), due to the non-observability of efforts, thereby not changing the participation decision.

We now look for the optimal price for time-inconsistent consumers. Given a price p , welfare is given by

$$W = \int_0^{\epsilon^{inc}} [(V((V')^{-1}(p)) - c((V')^{-1}(p) - y^{inc})) - d(y^{inc}) - k] df(k) + \int_{\epsilon^{inc}}^K [(V((V')^{-1}(p)) - c(V')^{-1}(p))] df(k). \quad (13)$$

A quick glance at the welfare function in (13) suggests that the incentives on behavior in the short-term (aimed at increasing y^{inc} , with $y^{inc} = y^{inc}(p, \beta)$) and long-term (aimed at increasing ϵ^{inc} , with $\epsilon^{inc} = \epsilon^{inc}(p, \beta, \hat{\beta})$) interact. Intuitively, a person can only exert effort

in the short-term if they decided to participate earlier on; moreover, the social benefit of participating depends on the effort one exerts later on. Therefore, the optimal price needs to account for the intertemporal effect created by each type of incentive.

Before exploring this interaction, we characterize the optimal price.

Proposition 2 *A solution to the welfare maximization problem exists in the case of a time-inconsistent consumer. Moreover, it entails $p^{inc} \geq c$ for all $\beta \in [0, 1]$.*

Proof. See Appendix 1. ■

On the basis of Proposition 2, it is possible to state that when a consumer is present-biased, a higher price than the marginal (social) cost helps to correct the low effort by raising the benefits of good behavior, i.e. by raising the value of the savings achieved by it. A higher price than the marginal social cost is needed to both induce self-control effort, and to provide the incentive to participate.

In order to better understand the forces at play when determining the level of the optimal price incentives, let us express the FOC of the welfare function in (13) w.r.t. p . The FOC is reported in (14), where, for convenience, we denote the second and third terms with Λ_1 and Λ_2 :

$$\underbrace{\frac{p - c}{V''((V')^{-1}(p))}}_{DWL} + \frac{1}{K} \left(\underbrace{\epsilon^{inc} (c - d_y(y^{inc})) \frac{\partial y^{inc}}{\partial p}}_{\Lambda_1} + \underbrace{(c y^{inc} - d(y^{inc}) - \epsilon^{inc}) \frac{\partial \epsilon^{inc}}{\partial p}}_{\Lambda_2} \right) = 0. \quad (14)$$

Equation (14) presents the marginal social benefits (in terms of higher short- and long-term efforts) and the marginal social cost (in terms of higher deadweight loss) of raising the price above c .

The term $\Lambda_1 = \epsilon^{inc} (c - d_y(y^{inc})) \frac{\partial y^{inc}}{\partial p}$ represents the marginal social benefit of incentives that originates from a higher self-control effort. This benefit should account for the ex-ante probability of participating (namely, ϵ^{inc}), as it can only be appropriated if the consumer participated in $t = 0$. A higher price leads the current base ϵ^{inc} of consumer types who participated in stage 0 to increase their self-control effort in $t = 1$ by $\frac{\partial y^{inc}}{\partial p}$, thus gaining the amount $c - d_y(y^{inc})$ from each additional unit of effort.

Conversely, the term Λ_2 represents the marginal social benefit of increasing participation. This depends on i) the value of participation ($cy^{inc} - d(y^{inc}) - \epsilon^{inc}$) of the marginal consumer, i.e. of the $k = \epsilon^{inc}$ consumer type and ii) on how much participation increases due to the price (namely, $\frac{\partial \epsilon^{inc}}{\partial p}$).

The terms Λ_1 and Λ_2 clearly show that, in a long-term perspective, the benefit of incentives is once again given by the trade-off between their value and their effectiveness. In fact, for high present-bias, the effectiveness of incentives (expressed by the terms $\frac{\partial y^{inc}}{\partial p}$ and $\frac{\partial \epsilon^{inc}}{\partial p}$) is low. Conversely, for low present-bias, the value of incentives (expressed by the terms $\epsilon^{inc}(c - d_y(y^{inc}))$ and $(cy^{inc} - d(y^{inc}) - \epsilon^{inc})$) is low.

Once again, the previous analysis suggests that the optimal level of incentives is non-monotonic with respect to the degree of present-bias. The following Corollary formalizes this result and shows that the optimal price distortion shows a maximum for internal levels of β , while it is zero for the two extreme cases of $\beta = 1$ (fully rational) and $\beta = 0$ (fully present-biased) consumers.

Corollary 2 *The optimal price is $p^{inc} = c$ for $\beta = 1$ or $\beta = 0$, and $p^{inc} > c$ for all $\beta \in (0, 1)$.*

Proof. See Appendix 1. ■

In the following section, we explore the relationship between awareness and incentives.

4.2 Effect of consumer's awareness on incentives

The consumer's degree of awareness of her own present-bias does not affect the actual short-term behavior: in fact, the effort y^{inc} is the same, regardless of $\hat{\beta}$. However, the degree of awareness does affect a consumer's earlier expectation y_0^{inc} about effort, with consequences on the participation decision $\epsilon^{inc}(p, \beta, \hat{\beta})$. In particular,

$$\frac{\partial \epsilon^{inc}}{\partial \hat{\beta}} = \frac{\beta p^2}{d_{yy}(\hat{\beta}p)}(1 - \hat{\beta}) > 0. \quad (15)$$

A higher degree of naivety increases the probability of a consumer participating in stage $t = 0$, because naive consumers overestimate their future self-control y_0^{inc} . The more one

(naively) expects to behave rationally in the future, the less severe the participation problem. This has both positive and negative effects on the optimal level of incentives.

On the one hand, the greater willingness of naive consumers to participate means that the benefits of improved self-control (expressed by the term Λ_1 in (14)) can be attained by a larger customer base. This implies that the effects of incentives on self-control effort are socially more valuable when consumers are naive rather than sophisticated.

On the other hand, the optimistic attitude of naive consumers about their future self-control effort means that they participate even when they should refrain from doing so, given the *actual* effort that they will exert later on. In fact, given that $\beta(py_0^{inc} - d(y_0^{inc})) - \epsilon^{inc} = 0$ (from the definition of ϵ^{inc}), then $\beta(py^{inc} - d(y^{inc})) - \epsilon^{inc} < 0$: in the case of naive consumers, there is a loss of surplus due to the fact that the self-control effort that they will actually exert is too low to cover the participation cost.

A policy implication that arises from these two effects is that an increase in the consumers' awareness should shift the focus of policymakers from improving self-control to improving participation. In fact, although a price incentive addresses, at the same time, the self-control and participation decisions, the extent to which it does so is not the same for naive and sophisticated consumers. When consumers are naive, the social benefit of incentives is mainly due to their effects on self-control, as a result of the large outreach of the policy. Conversely, when consumers are sophisticated, the social benefit of incentives is mainly related to the effect on participation, as the under-participation problem is more severe.

The presence of these countervailing effects suggests that the optimal price might not be monotonic in the degree of consumers' naivety. The following proposition formalises this result and provides the condition under which one of the effects prevails over the others.

Proposition 3 *There exists a value of present-bias $\tilde{\beta}$ such that, if $\beta \leq \tilde{\beta}$, the optimal price is higher for naive consumers than for sophisticated ones, while if $\beta > \tilde{\beta}$, the optimal price is lower for naive consumers than for sophisticated ones.*

Proof. See Appendix 1. ■

Proposition 3 states that when a consumer has a strong present bias, the optimal price incentive is higher for a naive consumer than for a sophisticated one. Conversely, when the consumer's present-bias is low, sophisticated consumers should be given a higher incentive than naive ones.

Intuitively, a loss in welfare has two main sources. First, inefficiency is generated from the low participation rate; this inefficiency is more severe for sophisticated consumers, who refrain from participating from the outset as they know that they will not exert self-control later on. Second, there is the inefficiency that originates from the participation of consumers who will later exert insufficient self-control, thus dissipating the participation cost; this inefficiency, which arises from the distortion between y^{inc} and y_0^{inc} , is present when consumers are naive, and it is particularly severe for high levels of present-bias. Hence, with high present-bias, the optimal price incentive is higher for a naive consumer than for a sophisticated one. Conversely, for low levels of present-bias, the inefficiency originating from the dissipation of the participation effort is low, and incentives should mainly be directed toward the correction of the low participation rates of sophisticated consumers.

The threshold level of β at which the optimal price is the same for naive and sophisticated consumers, and which is denoted by $\tilde{\beta}$, depends in a complex way on the effort and utility functions. We provide insight into these interactions by means of a numerical analysis in the following section.

4.3 An illustrative example

In order to gain better insight into the implications of consumer's awareness on the optimal price, a numerical and graphical analysis has been conducted.

To this aim, as in Section 3.4, let $d(y) = \gamma y^2$; to obtain more insight into the impact of the demand elasticity, let us now assume that $V(x + y) = (x + y) - \theta(x + y)^2$. Therefore, in $t = 1$, the consumer exerts the effort $y^{inc} = \frac{\beta p}{2\gamma}$ from (5), but in $t = 0$ she expects to exert $y_0^{inc} = \frac{\hat{\beta} p}{2\gamma}$.

We exploit these functional forms to derive the marginal benefits and costs of a rise in price, as expressed by the FOC in equation (14). After straightforward simplifications, we

obtain:

$$\frac{p-c}{-2\theta} + \frac{1}{K} \underbrace{\frac{\beta}{2\gamma}}_{\frac{\partial y^{inc}}{\partial p}} \underbrace{\frac{\beta p}{\gamma} \left(\hat{\beta} - \frac{\hat{\beta}^2}{2} \right)}_{\frac{\partial \epsilon^{inc}}{\partial p}} p \left[c - \beta p + \frac{c}{2} - p \left(\hat{\beta} - \frac{\hat{\beta}^2}{2} \right) \right] = 0. \quad (16)$$

From (16), it is possible to confirm that optimal incentives are determined from the trade-off between the value of additional efforts (i.e., the term in the square parenthesis) and their effectiveness (i.e., the factors $\frac{\beta}{2\gamma}$ and $\frac{\beta p}{\gamma} \left(\hat{\beta} - \frac{\hat{\beta}^2}{2} \right)$).

The optimal price, which is the solution to equation (16), is represented in Figure 2 on the vertical axis, as a function of the level of present-bias β . The two curves represent the optimal price for two levels of consumer's awareness: a completely naive consumer ($\hat{\beta} = 1$) and a completely sophisticated one ($\hat{\beta} = \beta$). In panels (a) and (b) – Figure 2, we assume a utility function with $\theta = 10$ and $\theta = 1$, respectively. For robustness, we also test the model with alternative functional forms of the utility function. In panels (c) and (d) – Figure 2, we assume a radical – panel (c) – and an exponential – panel (d) – $V(\cdot)$ curve²⁰. We assume $\gamma = 1$ in all the figures. The numerical analysis confirms that the optimal price is always greater than the marginal cost, it is maximized for $\beta \in (0, 1)$, and it converges to the marginal cost both when the consumer is very present-biased (β goes to 0) and when she is almost rational (β goes to 1), in line with the results of Proposition 2 and Corollaries 1 and 2. Moreover, the price is higher for naive consumers when present-bias is high, and for sophisticated consumers when present-bias is low, as predicted in Proposition 4.

It should also be noted that the curve of incentives of sophisticated consumers increases (and decreases) more rapidly than that of naive consumers. This is due to the fact that sophisticated consumers are more difficult to convince, and they display a more severe participation problem, hence the trade-off between value and effectiveness of incentives is more intense in their case.

The numerical analysis also allows some additional information to be gained about the level of $\tilde{\beta}$, which represents the (internal) degree of present-bias at which naive and sophis-

²⁰In particular, $V(x, y) = 20(x + y)^{1/2}$ in panel (c) and $V(x, y) = 1 - ke^{-(x+y)}$, for any $k > 0$, in panel (d).




Figure 2: Optimal price for a naive and sophisticated consumer, quadratic utility functions with constant second derivatives ((a) and (b)), radical (c) and exponential (d) utility functions. Marginal costs equal to 10.

ticated consumers are provided with the same level of incentives (see Proposition 4). Specifically, we perform some simulations in order to understand how $\tilde{\beta}$ depends on the shape of the utility function (θ) and on the cost of effort (γ). The results of these simulations are provided in the following table.

	$\theta = 1$	$\theta = 5$	$\theta = 10$
$\gamma = 1$	0.32	0.18	0.14
$\gamma = 5$	0.55	0.47	0.41
$\gamma = 10$	0.57	0.55	0.52

The simulations show that the value of threshold $\tilde{\beta}$ increases when the effort cost increases and the concavity of the utility function decreases. Intuitively, a higher θ , which is associated to a more inelastic demand, decreases the deadweight loss for a unit increase of the price. This calls for higher prices for both types of consumers, but especially so for sophisticated ones, due to their more rapid responsiveness of incentives to present-bias. In turn, this implies that threshold $\tilde{\beta}$ decreases. Similarly, when the effort becomes more expensive (i.e., γ increases), the optimal effort decreases; this reduces the inefficiency caused by the consumer's present-bias, and calls for lower incentives for both types of consumers, but especially so for sophisticated ones. Hence, threshold $\tilde{\beta}$ increases.

5 Conclusion

The literature on consumer's present bias is generally concerned with the effects of people's present-biased decisions on consumption. We contribute to this literature by modelling present-biased decisions that involve the distortion of an effort, rather than of the demand of goods. This analysis provides direction to policymakers about the optimal size of financial incentives, whenever they are conditioned on the future outcomes of a self-control effort, such as the weight loss achieved after we exert effort on a diet. Our aim is to highlight the main trade-offs of these kinds of policies.

We show that a large incentive provided to correct the severe misbehavior of a consumer might be suboptimal, if the degree of present-bias and awareness of the individual are not

taken into account. In other words, we find that financial incentives should be larger for intermediate levels of present-bias, while they should be negligible for both very high and very low levels of present-bias.

A non-monotonic profile of the socially optimal level of financial incentives arises from the trade-off between the value of incentives and their effectiveness. A very present-biased consumer obtains a greater benefit from correcting her effort, but is also more difficult to convince, because she inconsistently disregards the future benefits of behavior. Conversely, a consumer with very low present-bias obtains limited advantages from exerting effort, but incentives are more effective because she recognises their importance.

Moreover, in a long-term perspective, this trade-off is complicated by the interplay of the levels of effort exerted at different stages. For example, an inconsistently low attendance rate of a wellness program dissipates the benefits of inducing program uptake, and, symmetrically, the low rate of program uptake of present-biased consumers dissipates the benefits of promoting regular attendance.

Incentives conditioned on the future outcomes of current effort are currently being employed extensively by policymakers to encourage people to initiate healthy habits and continue them on a long-term basis. They conveniently complement other policy instruments such as subsidizing prevention initiatives, or bans and restrictions, which can prevent misbehavior in some situations, but they cannot eradicate it completely. For example, smoking regulations, or the prohibition of providing junk food in schools, cannot be enforced in private homes.

As recent as October 2017, a heated debate arose about the decision of an NHS commission in the UK to restrict surgery for obese people and smokers, unless they lost weight and stopped smoking, thereby promoting long-term behavior changes by raising the access cost to healthcare services. These controversial policies have been questioned, and not only on the grounds of discriminatory, ethical and clinical arguments. Doubts have also been expressed about their effectiveness in actually tackling problem behavior²¹.

²¹In October 2017, Robert West, professor of health psychology at UCL, commenting on the Hertfordshire CCG decision to make the ban indefinite for their patients, stated that "Rationing treatment on the basis of unhealthy behaviors betrays an extraordinary naivety about what drives those behaviors", see <http://edition.cnn.com/2017/10/31/health/smokers-obese-no-surgery-nhs-uk/index.html>.

This paper provides useful insights into the causes of the disappointing results of such policies. Although we do not attempt to suggest prescriptions on the design of optimal incentive schemes, our analysis can still offer useful policy directions by pointing out the shortcomings of financial incentives when the self-control problem concerns an effort decision. Incentives conditioned on the future results of effort may be scarcely effective on the people who need them the most and, in addition, they may induce over-exertion of effort in rational individuals. In fact, if the present bias is very high, or there is a definite predominance of either present-biased or rational consumers, financial incentives can entail large social costs and few benefits. Although a price distortion might be the best policy response to present-biased consumption decisions, our results show that a policy that simply keeps increasing the financial gains from effort might be too naive, socially costly and, more importantly, ineffective in inducing behavioral changes.

Naturally, this raises questions about what the optimal policy should be in the presence of an effort. As is the case with consumption goods, also in the case of effort, policymakers could combine different policy instruments, such as bans and regulations, incentives on goods consumed at the same time of the effort decision, or at a later time. While we do not make the case for any one of these specifically, the optimal combination of different instruments to incentivize effort could be an interesting venue for future research.

References

- [1] Ainslie, G.W., 1992. *Picoeconomics: The strategic interaction of successive motivational states within the person*, Cambridge University Press, Cambridge.
- [2] Ariely, D., and K. Wertenbroch, 2002. Procrastination, Deadlines, and Performance: Self-Control by Precommitment. *Psychological Science*, XIII: 219-224.
- [3] Atkinson, A.B. and Stiglitz, J.E., 1972. The structure of indirect taxation and economic efficiency. *Journal of Public Economics*, 1(1): 97-119.
- [4] Atkinson, A.B. and Stiglitz, J.E., 1976. The Design of Tax Structure: Direct versus Indirect Taxation, *Journal of Public Economics*, 6: 55-75.

- [5] Baumol, W.J. and Bradford, D.F., 1970. Optimal Departures From Marginal Cost Pricing. *The American Economic Review*, 60(3) 265-283
- [6] Beshears, J., Choi J., Laibson D. and Madrian B., 2005. Early Decisions: A Regulatory Framework. *Swedish Economic Policy Review*, 12(2): 41-60.
- [7] Christiansen, V. and Smith, S., 2012. Externality-correcting taxes and regulation. *Scandinavian Journal of Economics*, 114(2): 358-383.
- [8] Chu, L. Y., Sappington, D.E.M., 2013. Motivating Energy Suppliers to Promote Energy Conservation. *Journal of Regulatory Economics*, 43(3), 229-247.
- [9] Cremer, H., Gahvari, F. and Ladoux, N. (1998). Externalities and optimal taxation, *Journal of Public Economics*, 70(3): 343-364
- [10] Diamond, P. and Mirrlees, J., 1971. Optimal Taxation and Public Production II: Tax Rules. *The American Economic Review*, 61(3): 261-78
- [11] Dixit, A., 1970. On the Optimum Structure of Commodity Taxes. *The American Economic Review*, 60(3): 295-301
- [12] DellaVigna, S., and Malmendier, U., 2006. Paying Not to Go to the Gym. *American Economic Review*, 96 (3), 694-719.
- [13] DellaVigna, S., and Malmendier, U., 2004. Contract Design and Self-Control: Theory and Evidence. *Quarterly Journal of Economics*, CXIX (2), 353-402.
- [14] Fullerton, D. and Wolverton, A. (2005). The two-part instrument in a second-best world, *Journal of Public Economics*, 89(9-10): 1961-1975
- [15] Haavio, M. and Kotakorpi, K., 2011. The political economy of sin taxes. *European Economic Review*, 55(4): 575-594.
- [16] Kaiser Family Foundation Health Research and Educational Trust. Employer Health Benefits: 2015 Annual Survey.

- [17] Laffont, J.J., and Tirole J., 1993. A Theory of Incentives in Procurement and Regulation *MIT Press*.
- [18] Laibson, D.I., 1997. Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics*, CXII, 443-477.
- [19] Larwood, L., and W. Whittaker, 1977. Managerial present-bias: Self-Serving Biases in Organizational Planning. *Journal of Applied Psychology*, 62(2): 194-198.
- [20] Lerner, A.P., 1970. On Optimal Taxes with an Untaxable Sector. *The American Economic Review*, 60:284-94.
- [21] Madrian, B.C., and D. Shea, 2001. The Power of Suggestion: Inertia in 401(k) Participation and Savings behaviour. *Quarterly Journal of Economics*, CXVI, 1149-1187.
- [22] Mattke, S., Kapinos, K., Caloyeras, J.P., Taylor, E.A., Batorsky, B., Liu, H., Van Busum, K.R. and Newberry, S., 2015. Workplace Wellness Programs. Services Offered, Participation, and Incentives. *Rand Health Quarterly*. 5(2):7.
- [23] O'Donoghue, T.D., and M. Rabin, 1999a. Doing It Now or Later. *The American Economic Review*, LXXXIX, 103-124.
- [24] O'Donoghue, T.D., and M. Rabin, 1999b. Incentives for Procrastinators. *Quarterly Journal of Economics*, CXIV, 769-816.
- [25] O'Donoghue, T.D., and M. Rabin, 2001. Choice and Procrastination. *Quarterly Journal of Economics*, CXVI, 121-160.
- [26] O'Donoghue, T.D., and M. Rabin, 2003. Studying Optimal Paternalism, illustrated by a Model of Sin Taxes. *The American Economic Review*, 93, 186-191.
- [27] O'Donoghue, T.D., and M. Rabin, 2006. Optimal Sin Taxes. *Journal of Public Economics*, 90, 1825-1849.
- [28] O'Donoghue, T.D., and M. Rabin, 2008. Procrastination on long-term projects. *Journal of Economic Behavior and Organization*, 66(2), 161-175.

- [29] Phelps, E.S. and Pollak, R.A., 1968. On Second-Best National Saving and Game-Equilibrium Growth. *Review of Economic Studies*, 35(2), 185-199.
- [30] Royal College of Surgeons. Smokers and overweight patients: Soft target for NHS savings? London, RCS; 2016.
- [31] Sandmo, A. (1978). Direct versus indirect pigovian taxation. *European Economic Review*, 7(4): 337-349
- [32] Sandmo, A. (1976). Optimal taxation: An introduction to the literature. *Journal of Public Economics*, 6(1-2): 37-54
- [33] Strotz, R.H., 1955. Myopia and Inconsistency in Dynamic Utility Maximization. *Review of Economic Studies*, XXIII, 165-180.
- [34] Svenson, O., 1981. Are We All Less Risky and More Skilful than Our Fellow Drivers? *Acta Psychologica*, XLVII, 143-148
- [35] Wijkander, H. (1985). Correcting externalities through taxes on/subsidies to related goods, *Journal of Public Economics*, 28(1): 111-125
- [36] Woerner, A., 2018. Overcoming Time Inconsistency with a Matched Bet: Theory and Evidence from Exercising. Mimeo.

Appendix 1: Proof

Proof of Proposition 1. In order to determine a social planner's optimal policy, let us express welfare as

$$W = V((V')^{-1}(p)) - c((V')^{-1}(p) - y^*) - d(y^*). \quad (17)$$

The FOC of (17) w.r.t. p is $\frac{V'}{V''} - \frac{c}{V''} + \frac{c}{d_{yy}} - \frac{d_y}{d_{yy}} = 0$. Given that $V' = d_y = p$ from (1) and (3), we immediately obtain the optimal price $p^* = c$.

Let us now consider the case of time-inconsistent consumers. First note that $c < p^{inc} < c/\beta$. Indeed, $p^{inc} = c \frac{d_{yy} - \beta V''}{d_{yy} - \beta^2 V''} = \frac{c}{\beta} \frac{\beta d_{yy} - \beta^2 V''}{d_{yy} - \beta^2 V''} < \frac{c}{\beta}$. Given that the first best price is $p^* = c$, the welfare in the first best is

$$W(p^*) = V((V')^{-1}(c)) - c((V')^{-1}(c) - d_y^{-1}(c)) - d(d_y^{-1}(c)).$$

Conversely, with time-inconsistent consumers, the welfare achieved through the optimal policy p^{inc} is $W^{inc} = V((V')^{-1}(p^{inc})) - c((V')^{-1}(p^{inc}) - d_y^{-1}(\beta p^{inc})) - d(d_y^{-1}(\beta p^{inc}))$.

We want to verify that $W^{inc} < W(p^*)$. This is true iff

$$[V((V')^{-1}(p^{inc})) - c(V')^{-1}(p^{inc})] - [V((V')^{-1}(c)) - c(V')^{-1}(c)] \quad (18)$$

$$+ [cd_y^{-1}(\beta p^{inc}) - d(d_y^{-1}(\beta p^{inc}))] - [cd_y^{-1}(c) - d(d_y^{-1}(c))] < 0 \text{ iff} \quad (19)$$

We now prove that both lines (18) and (19) are negative. Let us study the function $f(m) = V((V')^{-1}(m)) - c(V')^{-1}(m)$. We have that $f'(m) = \frac{V' - c}{V''} = \frac{m - c}{V''} < 0$ for $m > c$; hence, (18) is negative.

Let us now study the function $g(m) = cd_y^{-1}(m) - d(d_y^{-1}(m))$. We have that $g'(m) = \frac{c - d_y}{d_{yy}} = \frac{c - m}{d_{yy}} < 0$ for $m > c$; hence, (19) is negative. This completes the proof. ■

Proof of Corollary 1. By studying the implicit function $G(p^{inc}, \beta) = p^{inc} - c \frac{d_{yy} - \beta V''}{d_{yy} - \beta^2 V''}$ expressed by (6), we obtain

$$\frac{dp^{inc}}{d\beta} = - \frac{\partial G / \partial \beta}{\partial G / \partial p^{inc}} = - \frac{\partial G / \partial \beta}{1} = c \frac{\beta(1 - \beta)y_\beta^{inc}(V''d_{yyy} - V'''d_{yy}) - \beta^2(V'')^2 + d_{yy}V''(2\beta - 1)}{(d_{yy} - \beta^2 V'')^2}. \quad (20)$$

The sign of (20) is given by its numerator. At the numerator of (20), the first two terms are always negative (owing to the fact that $y_\beta^{inc} > 0$ from (5), and from the assumptions $V'' < 0$, $V''' \geq 0$, $d_{yy} > 0$, $d_{yyy} > 0$) for any $\beta \in (0, 1)$. The last term, $d_{yy}V''(2\beta - 1)$ is negative if and only if $\beta > 1/2$. Given that, in $\beta = 0$, $\frac{dp^{inc}}{d\beta} = c \frac{-d_{yy}V''}{(d_{yy})^2} > 0$, then there exists a $\bar{\beta} \in (0, 1/2)$ such that $\frac{dp^{inc}}{d\beta} > 0$ for any $\beta < \bar{\beta}$, while $\frac{dp^{inc}}{d\beta} \leq 0$ for any $\beta \geq \bar{\beta}$. ■

First best in the long-term (Section 4). By substituting (1) and (3) in (10), the latter can be rewritten as

$$W = [V((V')^{-1}(p)) - c(V')^{-1}(p)] + \frac{1}{K} \left(cd_y^{-1}(p)\epsilon^* - d(d_y^{-1}(p))\epsilon^* - \frac{(\epsilon^*)^2}{2} \right).$$

In order to find the optimal price, we find the FOC of W w.r.t. p :

$$\frac{V' - c}{V''} + \frac{1}{K} \left[c \left(\frac{1}{d_{yy}(y^*)}\epsilon^* + d_y^{-1}(p) \frac{\partial \epsilon^*}{\partial p} \right) - \left(\frac{d_y(y^*)}{d_{yy}(y^*)}\epsilon^* + d(d_y^{-1}(\delta p)) \frac{\partial \epsilon^*}{\partial p} \right) - \epsilon^* \frac{\partial \epsilon^*}{\partial p} \right] = 0.$$

Given that $V' = p$ and $d_y(y^*) = p$, the previous expression becomes

$$\frac{p - c}{V''} - \frac{1}{K} \frac{\epsilon^*}{d_{yy}(y^*)} (p - c) + \frac{1}{K} \frac{\partial \epsilon^*}{\partial p} (cd_y^{-1}(p) - d(d_y^{-1}(p)) - \epsilon^*) = 0.$$

Since $\epsilon^* = pd_y^{-1}(p) - d(d_y^{-1}(p))$, the expression becomes

$$\frac{p - c}{V''} - \frac{1}{K} \frac{\epsilon^*}{d_{yy}(y^*)} (p - c) - \frac{1}{K} \frac{\partial \epsilon^*}{\partial p} d_y^{-1}(p) (p - c) = 0,$$

which is verified for $p = c$. ■

Proof of Proposition 2 and Corollary 2. We prove the Proposition and Corollary through two Lemmas, concerning the characterisation (Lemma 1) and the existence (Lemma 2) of a solution to the social planner's problem, respectively.

Lemma 1 *If a solution exists, it entails $p^{inc} = c$ for $\beta = 1$ or $\beta = 0$, and $p^{inc} > c$ for all $\beta \in (0, 1)$.*

Proof. Expression (13) can be rewritten as

$$W = [V((V')^{-1}(p)) - c(V')^{-1}(p)] + \frac{1}{K} \left(cd_y^{-1}(\beta p)\epsilon^{inc} - d(d_y^{-1}(\beta p))\epsilon^{inc} - \frac{(\epsilon^{inc})^2}{2} \right).$$

In order to find the optimal price, we find the FOC of W w.r.t. p :

$$\frac{p - c}{V''} + \frac{\beta \epsilon^{inc}}{d_{yy}(y^{inc})} \frac{c - \beta p}{K} + \frac{\partial \epsilon^{inc}}{\partial p} \frac{cd_y^{-1}(\beta p) - d(d_y^{-1}(\beta p)) - \epsilon^{inc}}{K} = 0, \quad (21)$$

where $V' = p$ and $d_y(y^{inc}) = \beta p$. The proof proceeds in two steps. We first prove that $p^{inc} = c$ for $\beta = 0$ or $\beta = 1$. Secondly, we prove that $p^{inc} \leq c$ cannot be a solution to the social planner's problem for $\beta \in (0, 1)$.

When $\beta = 0$, $\epsilon^{inc} = \frac{\partial \epsilon^{inc}}{\partial p} = 0$. Therefore, (21) becomes $K \frac{p-c}{V''} = 0$, which is verified for $p^{inc} = c$. When $\beta = 1$, we also have $\hat{\beta} = 1$; therefore, $p^{inc} = c$.

Let us now prove that $p^{inc} > c$ for $\beta \in (0, 1)$. Let us suppose, by contradiction, that $p^{inc} \leq c$. Thus, the sum of the first two terms of condition (21) would be strictly positive. Since $\frac{\epsilon^{inc}}{\partial p}$ is also positive, the welfare maximization problem has an internal solution when the last term, which we denote with $L(\beta, \hat{\beta})$, is strictly negative, i.e.

$$L(\beta, \hat{\beta}) = cd_y^{-1}(\beta p) - d(d_y^{-1}(\beta p)) - \beta \left(pd_y^{-1}(\hat{\beta} p) - d(d_y^{-1}(\hat{\beta} p)) \right) < 0.$$

We now study the function $L(\beta, \hat{\beta})$ and show that, even when it is minimized, it is still positive for $p \leq c$, thus contradicting the hypothesis that $p^{inc} \leq c$ is an equilibrium when $\beta \in (0, 1)$. To this aim, let us note that the function $L(\beta, \hat{\beta})$ always decreases in $\hat{\beta}$. This means that it is always minimized when $\hat{\beta} = 1$, for all p and β . Moreover, when $\hat{\beta} = 1$, we obtain

$$\begin{aligned} \frac{\partial L(\beta, 1)}{\partial \beta} &= \frac{cp}{d_{yy}(\beta p)} - \frac{\beta p^2}{d_{yy}(\beta p)} - (pd_y^{-1}(p) - d(d_y^{-1}(p))) \\ \frac{\partial^2 L(\beta, 1)}{\partial \beta^2} &= p \frac{-pd_{yy}(\beta p) - (c - \beta p)pd_{yyy}(\beta p)}{(d_{yy}(\beta p))^2}. \end{aligned}$$

Since from hypothesis $p \leq c$, then $c - \beta p > 0$; moreover, $d_{yy} > 0, d_{yyy} \geq 0$. Hence, $\frac{\partial^2 L}{\partial \beta^2} < 0$, which implies that $L(\beta, 1)$ is a strictly concave function in the support of β . As a consequence, it is minimized for either $\beta = 0$ or $\beta = 1$. In $\beta = 0$, $L(0, 1) = 0$. Conversely, in $\beta = 1$, $L(1, 1) = (c - p)d_y^{-1}(p)$, which is positive or null for $p \leq c$. It follows that when $p \leq c$, $\min L(\beta, \hat{\beta}) \geq 0$ for all $\beta \in (0, 1)$. As a consequence, when $p \leq c$, $L(\beta, \hat{\beta}) \geq 0$ for all $\beta, \hat{\beta} \in (0, 1) \times [\beta, 1)$. Therefore, the l.h.s. of (21) is strictly positive for $p \leq c$, which implies that there does not exist a solution with $p^{inc} \leq c$ for any $\beta \in (0, 1)$. ■

Lemma 1 only shows that the optimal price must be higher than the marginal cost. However, in principle, the problem might not be closed. The following Lemma ensures that

the problem is closed, and thus a solution always exists.

Lemma 2 *A solution to the welfare maximization problem always exists.*

Proof. When $\beta = 0$ or $\beta = 1$, a solution exists and it is equal to $p^{inc} = c$ from Lemma 1. Therefore, we only need to show that a solution exists when $\beta \in (0, 1)$.

First, let us observe that all the functions in (21) are continuous in p . Moreover, in $p = c$, the l.h.s. of (21) is strictly positive when $\beta \in (0, 1)$ by Lemma 1. Should the l.h.s. of (21) be strictly negative for some $p > c$ when $\beta \in (0, 1)$, then there exists some value of p such that the FOC in (21) is equal to zero by continuity, i.e a solution to the welfare maximization exists.

Let us consider the price $p = c/\beta$. The sum of the first two terms in (21) is strictly negative, while $\frac{\partial \epsilon^{inc}}{\partial p}$ is always positive. Let us focus on the third term, i.e. $L(\beta, \hat{\beta})$. We want to find the values of $\beta, \hat{\beta}$ so that $L(\beta, \hat{\beta})$ is maximized, and show that L is still negative for these values when p is sufficiently high. It should be noted that $L(\beta, \hat{\beta})$ decreases in $\hat{\beta}$. This implies that $L(\beta, \hat{\beta})$ is maximized when $\hat{\beta} = \beta$. By substituting $\hat{\beta} = \beta$ and $p = c/\beta$ in the expression of $L(\beta, \hat{\beta})$, the latter becomes

$$L(\beta, \beta) = cd_y^{-1}(c) - d(d_y^{-1}(c)) - \beta \left(\frac{c}{\beta} d_y^{-1}(c) - d(d_y^{-1}(c)) \right)$$

i.e. $L(\beta, \beta) = -(1 - \beta)d(d_y^{-1}(c)) < 0$ for all $\beta \in (0, 1)$. Hence, the l.h.s. of (21) is strictly negative for $p = c/\beta > c$.

Since the l.h.s. of (21) is strictly positive for $p = c$, it is strictly negative for $p = c/\beta > c$ and it is continuous, there exists at least one price $p \in (c, c/\beta)$ so that the l.h.s. is equal to zero by continuity, i.e. a solution to the welfare maximization problem exists. ■ ■

Proof of Proposition 3. From the FOC condition, the term

$$\Lambda = \left(\frac{\beta c}{d_{yy}(y^{inc})} - \frac{\beta^2 p}{d_{yy}(y^{inc})} \right) \epsilon^{inc} + (cd_y^{-1}(\beta p) - d(d_y^{-1}(\beta p)) - \epsilon^{inc}) \frac{\partial \epsilon^{inc}}{\partial p}$$

represents the marginal social benefit of increasing the price. Therefore, the higher the

marginal benefit Λ is, the larger the optimal price distortion. Let us study function Λ :

$$\begin{aligned}\frac{\partial \Lambda}{\partial \beta} &= (c - 2\beta p) \frac{\epsilon^{inc}}{d_{yy}(y^{inc})} + (\beta c - \beta^2 p) \frac{\partial}{\partial \beta} \left(\frac{\epsilon^{inc}}{d_{yy}(y^{inc})} \right) + \\ &+ \left(\frac{pc}{d_{yy}(y^{inc})} - \frac{\beta p^2}{d_{yy}(y^{inc})} - \frac{\partial \epsilon^{inc}}{\partial p} \right) \frac{\partial \epsilon^{inc}}{\partial p} + (cy^{inc} - d(y^{inc}) - \epsilon^{inc}) \frac{\partial^2 \epsilon^{inc}}{\partial p \partial \beta}\end{aligned}$$

$$\begin{aligned}\frac{\partial^2 \Lambda}{\partial \beta^2} &= -2p \frac{\epsilon^{inc}}{d_{yy}(y^{inc})} + 2(c - 2\beta p) \frac{\partial}{\partial \beta} \left(\frac{\epsilon^{inc}}{d_{yy}(y^{inc})} \right) + (\beta c - \beta^2 p) \frac{\partial^2}{\partial \beta^2} \left(\frac{\epsilon^{inc}}{d_{yy}(y^{inc})} \right) + \\ &+ \left(-cp^2 \frac{d_{yyy}(y^{inc})}{(d_{yy}(y^{inc}))^3} - p^2 \frac{d_{yy}(y^{inc}) - \beta p \frac{d_{yyy}(y^{inc})}{d_{yy}(y^{inc})}}{(d_{yy}(y^{inc}))^2} - \frac{\partial^2 \epsilon^{inc}}{\partial \beta^2} \right) \frac{\partial \epsilon^{inc}}{\partial p} + \\ &+ 2 \left(\frac{pc}{d_{yy}(y^{inc})} - \frac{\beta p^2}{d_{yy}(y^{inc})} - \frac{\partial \epsilon^{inc}}{\partial \beta} \right) \frac{\partial^2 \epsilon^{inc}}{\partial p \partial \beta} + (cy^{inc} - d(y^{inc}) - \epsilon^{inc}) \frac{\partial^3 \epsilon^{inc}}{\partial p \partial^2 \beta},\end{aligned}$$

where $\frac{\partial}{\partial \beta} \left(\frac{\epsilon^{inc}}{d_{yy}(y^{inc})} \right) = \frac{\frac{\partial \epsilon^{inc}}{\partial \beta} d_{yy}(y^{inc}) - p \epsilon^{inc} \frac{d_{yyy}(y^{inc})}{d_{yy}(y^{inc})}}{(d_{yy}(y^{inc}))^2}$. When β goes to zero, the values of function Λ and of its first derivative are also zero. However, the second derivative is positive. In fact,

$$\left. \frac{\partial^2 \Lambda}{\partial \beta^2} \right|_{\beta=0} = 2 \frac{c}{d_{yy}(y^{inc})} \frac{\partial \epsilon^{inc}}{\partial \beta} + 2 \left(\frac{cp}{d_{yy}(y^{inc})} - \frac{\partial \epsilon^{inc}}{\partial \beta} \right) \frac{\partial^2 \epsilon^{inc}}{\partial p \partial \beta}. \quad (22)$$

The first component of (22) is positive. The second component is positive, as the term in the parenthesis is positive. To see this, let us denote with $q(\beta) = cy^{inc} - \epsilon^{inc}$. We observe that $\frac{\partial q}{\partial \beta} = \frac{cp}{d_{yy}(y^{inc})} - \frac{\partial \epsilon^{inc}}{\partial \beta}$. The function $q(\beta)$ is always positive (as $q(\beta) > cy^{inc} - d(y^{inc}) - \epsilon^{inc} \geq 0$). Moreover, $\frac{\partial^2 q}{\partial \beta^2} = -cp^2 \frac{d_{yyy}(y^{inc})}{(d_{yy}(y^{inc}))^3} < 0$. Therefore, the function $q(\beta)$ increases in the interval $[0, 1]$, i.e. its first derivative is positive. We conclude that function Λ increases in the neighbourhood of $\beta = 0$. We now prove that it increases faster for naive consumers than for sophisticated ones. The value of the second derivative of Λ in the neighbourhood of $\beta = 0$ is zero for a fully sophisticated consumer ($\hat{\beta} = 0$). The value of the second derivative of Λ in the neighbourhood of $\beta = 0$ is equal to $2 \frac{c}{d_{yy}(y^{inc})} (py_0^{inc} - d(y_0^{inc})) + 2 \left(\frac{cp}{d_{yy}(y^{inc})} - (py_0^{inc} - d(y_0^{inc})) \right) y_0^{inc} > 0$ for a fully naive consumer ($\hat{\beta} = 1$).

Let us now determine the sign of $\frac{\partial \Lambda}{\partial \hat{\beta}}$ for a sufficiently high β :

$$\begin{aligned} \frac{\partial \Lambda}{\partial \hat{\beta}} = & \frac{\beta p(1 - \hat{\beta})}{d_{yy}} \cdot \left[\frac{\beta p}{d_{yy}} \left(c - \beta p - p\hat{\beta}(1 - \hat{\beta}) \right) - \beta p d_y^{-1}(\hat{\beta}p) + \right. \\ & \left. + 2 \left(c d_y^{-1}(\beta p) - d(d_y^{-1}(\beta p)) + \beta \delta d(d_y^{-1}(\hat{\beta}p)) - \beta p d_y^{-1}(\hat{\beta}p) \right) \right]. \end{aligned}$$

When β goes to 1, $\hat{\beta}$ also goes to 1, as $\hat{\beta} \geq \beta$. Therefore, $\frac{\partial \Lambda}{\partial \hat{\beta}} \Big|_{\beta=1} = 0^+ \cdot [-p d_y^{-1}(p)] = 0^-$: the optimal price is lower for naive consumers than for sophisticated ones. ■