

Model-based Evaluation of Cooling-Off Policies*

Christian Michel André Stenzel

October 1, 2020

Abstract

This paper studies the ex-ante prediction and ex-post evaluation of the effects of cooling-off policies when consumers may exhibit a projection bias. We set up a theoretical model in which a firm optimally reacts to consumers' preferences and the regulatory framework and show that neither the adoption of a mandatory cooling-off period nor a return policy is generically superior or consumer welfare improving. We then analyze how market-level data can help to evaluate the policies ex post using baseline statistics. This exploits the firm's optimal response to a policy which depends on its cost characteristics. We discuss robustness of the model predictions and ex-post assessment to a variety of modifications, and discuss the benefits of the approach for policy design.

JEL Classification: D04, D18, D90

Keywords: Cooling-off, Equilibrium-based Inference, Identification of non-standard preferences, Projection bias, Behavioral economics

*We thank Larbi Alaoui, José Apesteguia, Mircea Epure, Jana Friedrichsen, Paul Heidhues, Volker Nocke, David Rojo Arjona, Eleftheria Triviza, Chris Wallace, the advisory editor and two very constructive anonymous referees, and seminar participants for valuable comments. Michel gratefully acknowledges financial support from the Spanish Ministry of Economy and Competitiveness Grant ECO2016-76998. Stenzel gratefully acknowledges funding by the German Research Foundation (DFG) via the CRC TR 224 (Project C03).
Christian Michel, Universitat Pompeu Fabra and Barcelona GSE. E-mail: christian.michel@upf.edu.
André Stenzel, University of Mannheim and MaCCI. E-Mail: andre.stenzel@uni-mannheim.de.

1 Introduction

Economic research has shown that consumer behavior is oftentimes not well explained by standard preferences and rational choices. Consequently, a large number of economic models have emerged in which consumers deviate from standard preferences in different forms.¹ While such models are being used to guide policymakers, their normative analysis poses some problems. First, firms' equilibrium responses to a policy introduction can have unintended consequences for the market outcome and in particular consumers, which need to be taken into consideration (Spiegler, 2015). Second, consumers' choices do not necessarily reflect their preferences, which greatly complicates an ex-post assessment of the adopted policy's effectiveness.²

In this paper, we address both of these problems in the context of cooling off policies. We provide a comprehensive theoretical analysis of a setting in which consumers potentially exhibit a projection bias, and consider two distinct policies which are prominently used in practice: a mandatory cooling off period and a return policy. We characterize the dependence of the optimal firm behavior on the regulatory framework and on the firm's marginal cost and use this to show that neither policy is generically optimal with respect to maximizing consumer welfare. This highlights the need for an ex-post assessment of an adopted policy's impact.

Towards this, we analyze the data requirements under which the observed equilibrium changes due to the policy allow us to unambiguously evaluate whether consumers benefited from the policy adoption. We focus on market level data instead of consumer level data as it is typically easier to obtain. Our approach exploits that the combination of pre and post policy market outcomes carries information as the seller differentially responds to the policy adoption based on her characteristics, and can in principle be carried out also for other models with a behavioral component. It relates to a recent discussion of using equilibrium datasets to identify unobservables of models with potentially non-standard preferences (Eliaz and Spiegler, 2015).

We show that the considered policies differ in the specific data requirements for an ex-post evaluation. While these requirements are naturally model-specific, a reasonable set of aggregate data always serves at least as a screening device, and—depending on the model under consideration—may allow for an unambiguous assessment of whether consumers benefited from the policy adoption. As such, an ex ante consideration of the data requirements for a policy to be evaluated ex post can yield valuable insights for policy design.

Our baseline model in Section 2 nests standard preferences and a specific form of nonstandard time-inconsistent preferences. A monopolist offers a product to consumers, whose utility is state-dependent: they can be motivated (“hot”) or unmotivated (“cold”) with respect to

¹Examples of nonstandard preferences include loss aversion, hyperbolic discounting, presence of a projection bias, and consumer myopia. A growing number of models study the normative effects of policy introductions in presence of such preferences; see, for example, Heidhues and Köszegi (2010), and the overview articles by Armstrong (2015) and Grubb (2015).

²While there is a small but growing literature that seeks to structurally identify and estimate consumer preference parameters while accounting for the possibility that consumers have nonstandard preferences (Della Vigna, 2018), identification of such parameters can usually only be obtained under relatively strong data requirements.

the product, and this motivation can change over time. There is a delay between the purchase decision and consumption. This relates to cases in which consumers purchase products that they cannot consume immediately.³ Consumers potentially experience a *projection bias*. When in a “hot” consumption state, the bias makes consumers prone to overpredicting their future expected utility. Correspondingly, when in a “cold” state, the bias makes them prone to underpredicting their future expected utility.

As motivated consumers overpredict their future consumption utility due to the projection bias, there are equilibrium cases in which the firm profitably sets prices that lead to a negative expected consumer welfare. This leaves scope for consumer policies aimed at increasing consumer welfare via changing the firm’s pricing incentives. We consider two policies that are applied in various forms in practice and as such are of direct policy relevance in many countries. A *mandatory cooling off period* requires a consumer to confirm the product purchase after an initial waiting period.⁴ This delays consumption overall, but leaves the final consumption decision based on predicted utilities instead of actual utilities. By contrast, a *return policy* allows for the consumer to return the product immediately before consumption in exchange for a reimbursement of the product price. A consumer can thus reverse his purchase decision when the good arrives and the uncertainty about the consumption utility has been resolved.

We analyze the model in [Section 3](#). We find that neither policy is generically optimal with respect to maximizing consumer welfare. Absent any policy intervention, the firm targets either motivated consumers who overpredict their expected consumption utility because of the bias, or all consumers, which requires pricing below the actual expected utility. The same holds true once a mandatory cooling off period is introduced—however, due to the required confirmation, only twice motivated consumers obtain the product which makes the purchase under this targeting strategy less likely. With a return policy, the final consumption/return decision is based on actual utilities such that consumers never suffer from ex-post regret. However, there still is a tradeoff for the regulator as conditional on the firm targeting only twice motivated consumers, it can extract the full consumption utility. This increases the relative profitability of this strategy, so that exclusive targeting is more likely to materialize.

In [Section 4](#), we assess the data requirements to determine whether a policy has positively affected the market outcome and whether a different policy might have been advantageous. The crucial step in our analysis is to relate the theoretical model’s equilibrium combinations pre and post policy adoption to the movement of different variables such as prices and quantities that a

³There are many examples of such situations, e.g., internet retail shopping for physical goods, or any other form of shopping where a good has to be pre-ordered. In such instances, a consumer’s realized utility when finally consuming a product is oftentimes state-dependent and the projection bias relevant; see [Loewenstein et al. \(2003\)](#) for a broad discussion of empirical and experimental evidence.

⁴In several US states, mandatory cooling off periods are in effect for general door-to-door sales. See [Rowe \(2007\)](#) for a detailed description of the 2007 “Consumer Solicitation Sales Act” that was introduced in the state of Maine. Adoption of a mandatory cooling off period has furthermore been discussed for consumers who request tattoos both in the US and Germany. A different context in which mandatory cooling-off periods are in place is for the processing of divorce filings in many countries ([Lee, 2013](#)).

researcher might observe. This allows us to determine the data requirements that are sufficient to distinguish the different pre and post equilibrium combinations, which in turn allows for an assessment of the benefits of the policy. We show that a mandatory cooling off period has lower data requirements to be evaluated than a return policy. This is predominantly driven by the firm having three instead of two potential targeting strategies following the adoption of a return policy, which increases the number of equilibrium combinations the ex-post evaluation needs to distinguish. Data on aggregate quantities alone is sufficient to determine whether a cooling off period was beneficial for consumer welfare, while a return policy requires data on quantities, return frequencies, and market size. These requirements are unaffected by allowing for motivation states to be correlated over time, or return to be costly to either consumers or the firm.

We discuss the implications of various extensions for both the theoretical prediction and the ex-post evaluation in [Section 5](#). When allowing for additional consumer heterogeneity—either by introducing a second type of unbiased consumers, or by introducing horizontal consumer differentiation—the data requirements for a full evaluation of the impact on consumer welfare increase. However, an analysis based only on aggregate data on prices, quantities, and return frequencies is still powerful as it either allows the identification of cases where consumer welfare was harmed (setting with unbiased consumers), or at least serves as a screening device to isolate cases which warrant an in-depth investigation (setting with horizontally differentiated consumers).

The analyses point out two key challenges of any model-based evaluation: The data requirements to assess the policy ex post depend on both the firm’s number of potential responses to the policy adoption, and whether the behavior of the economic indicators of interest is unambiguous conditional on an identified combination of pre and post policy equilibrium behavior.

We conclude in [Section 6](#). A systematic assessment of the necessary data requirements to conduct policy evaluations in behavioral models has—to our knowledge—not yet been extensively analyzed by the literature. In line with our analysis, it can provide important guidance both for the initial policy design and the data collection process associated with the implementation of the policy.

Literature A key part of our paper focuses on how equilibrium combinations before and after a policy introduction can help making inference on consumer preferences and the efficacy of a policy. This is closely related to a small literature that seeks to provide tests for nonstandard consumer preferences and to make market assessments without relying on structural econometric models. Close to our approach is [Heidhues and Koszegi \(2018\)](#), Section 2.6, who outline how to test for consumers not responding to an increase in add-on prices, which would point to non-salience of such prices. They propose to base the test on a misprediction-augmented Slutsky equation, or to look at the difference between add-on prices and costs relative to the ratio of add-on demand between marginal and average consumers.

[Eliaz and Spiegel \(2015\)](#) discuss several new research directions within the field of behavioral industrial organization. One that relates to our paper is to use the assumption that data was created under an equilibrium market model. They illustrate this in a model in which two firms compete by choosing the number of products and marketing messages that lead consumers to consider a product. When data on consumers' individual marketing exposure are available, exploiting the equilibrium assumption allows for the identification of a consideration function that depends on a consumer's marketing exposure.

[Loewenstein et al. \(2003\)](#) explicitly define the notion of projection bias that we use and give a formal basis for the utility representation of our model. Using a structural econometric model, [Conlin et al. \(2007\)](#) find evidence that consumers exhibit a projection bias when making clothing catalog orders surrounding rapid weather changes. [Fisher and Rangel \(2014\)](#) find that consumers mispredict their future preference for being a hot state when currently being in a cold state, and vice versa. [Chang et al. \(2018\)](#) find that daily air pollution levels have a significant effect on the decision to purchase or cancel health insurance that is inconsistent with rational choice theory. They argue that these findings give most support for salience effects or for consumers exhibiting a projection bias.

There is substantial related work that theoretically analyzes firms' responses to consumer biases and consumer protection policies, see [Armstrong \(2015\)](#) and [Grubb \(2015\)](#) for broad overviews. [Spiegler \(2015\)](#) analyzes the equilibrium effects of nudging and shows that different nudges can have unintended negative consequences when both demand and supply are accounted for. [Inderst and Ottaviani \(2013\)](#) study the relationship between a financial adviser and a private investor and show that granting a cancellation right to rational consumers who foresee the adviser's self-interest can make his cheap talk credible. [Michel \(2017\)](#) analyzes the effects of consumer protection policies in a competitive retail model in which firms can offer add-on products at the point of sale and finds that a return period for extended warranty contracts weakly increases consumer surplus. [Camerer et al. \(2003\)](#) broadly discuss advantages and disadvantages of different protection policies.

A small but growing stream of the literature seeks to empirically identify utility functions while accounting for the possibility that they are nonstandard; see, for example, [Conlin et al. \(2007\)](#) and [Laibson et al. \(2008\)](#) using firm- and consumer-level data. [DellaVigna \(2018\)](#) provides an extensive discussion of the advantages and possible avenues of using structural models with potential nonstandard components.

[Abaluck and Adams \(2017\)](#) show how to identify consumer inattention with respect to choice sets by testing for a violation of Slutsky symmetry. [Heidhues and Strack \(2019\)](#) provide identification results for partial naivety in an optimal stopping problem when data on both the stopping probabilities and on the continuation value are available. [Mahajan et al. \(2020\)](#) use agents' elicited beliefs about the future together with signals that are correlated with being present-biased to identify time preferences of types with different degrees of time-inconsistency.

There are further related empirical approaches to assess public policies in the presence of

behavioral agents; see [Bernheim and Taubinsky \(2018\)](#) for a broad overview. We particularly relate to work on taxation in the presence of potentially non-rational consumers; see, for example [Gruber and Köszegi \(2001\)](#), [O'Donoghue and Rabin \(2006\)](#), and [Allcott et al. \(2014\)](#). [Hinnosaar \(2016\)](#) finds that restricting the number of hours a day in which time-inconsistent consumers can buy alcohol can be welfare improving, but is inferior to an alcohol tax. [Chetty et al. \(2009\)](#) look at consumers' responses to different consumption taxes. They find that posting tax-inclusive prices decreases demand for products relative to posting non-tax-inclusive prices, which they attribute to taxes not being salient to some consumers in the latter case. [Shui and Ausubel \(2004\)](#) focus on the difference in introductory and regular credit card interest rates chosen by borrowers and find that the ratio between these rates is too low for the choices to be consistent with consumer rationality.

From a normative viewpoint, our paper is further related to work in behavioral welfare economics, see [Bernheim \(2009\)](#) for an overview. To overcome the problem of not observing revealed preferences, [Bernheim and Rangel \(2009\)](#) propose a behavioral welfare criterion based on an unambiguous choice relation: it relies on a specific choice being unambiguously chosen over a different one if both are available. This systematically differs from our approach in that we focus on the potential of firms to strategically change the choice set in response to changes in the institutional setting.

2 The Model

We consider a firm that produces a single good of fixed quality at constant marginal cost $c \geq 0$. The firm chooses the product price p and faces a mass one of consumers. For simplicity, consumers have a single opportunity to purchase the product, and the product can be consumed exactly once. Consumption takes place one period after purchase, e.g. because of shipping, such that the model has an intertemporal component.

Time is discrete and covers at most three periods: $t = 1, 2, 3$. Consumers are risk-neutral and have a common discount factor $\delta_t \in (0, 1]$ between periods t and $t+1$. Ignoring discounting and the product price, a consumer's gross consumption utility when consuming the product in period t , $u(s_t)$, depends on her motivation state $s_t \in \{\underline{s}, \bar{s}\}$. We denote by $\underline{u} \equiv u(\underline{s}) > 0$ the gross consumption utility of an unmotivated consumer, and by $\bar{u} \equiv u(\bar{s}) > \underline{u}$ the gross consumption utility of a motivated consumer. $\Delta \equiv \bar{u} - \underline{u}$ parametrizes the utility difference.

The product price affects consumer utility negatively and linearly. A consumer's final consumption utility when purchasing a product in period t and consuming it in period $t+1$ depends on her gross consumption utility $u(s_{t+1})$ in motivation state $s_{t+1} \in \{\underline{s}, \bar{s}\}$ in period $t+1$, on the product price p , and on her intertemporal discount factor δ_t : It is given by $-p + \delta_t u(s_{t+1})$. Because of the delayed consumption, consumers must predict their future utility $u(s_{t+1})$ when making the purchase decision in period t . This prediction is complicated by the fact that a

consumer’s motivation may change over time, and that we allow for consumers to experience a *projection bias*, as we explain next.

Each consumer correctly assigns probability $\mu \in (0, 1)$ to the event of being motivated in any future period, and a probability $1 - \mu$ to the event of being unmotivated. The probability of being in a specific motivation state in the next period is independent of the current state.⁵ Under the projection bias, consumers partially incorporate the utility associated with their current state into the prediction of the consumption utility for any other state. This implies a correct assessment of the utility in the same state in the future, but a biased prediction for a different state. Following [Loewenstein et al. \(2003\)](#), the predicted gross utility of a consumer for state s given a current motivation s' when experiencing a projection bias is given by

$$\tilde{u}(s|s') = \alpha u(s') + (1 - \alpha)u(s), \quad (1)$$

where $\alpha \in [0, 1]$ captures the degree of the bias. For $\alpha = 0$, consumers are fully rational. For $\alpha > 0$, consumers overpredict their consumption utility in the unmotivated state in the future whenever they are currently motivated, and underpredict their consumption utility in the motivated state in the future when they are currently unmotivated:

$$\tilde{u}(\bar{s}|\underline{s}) = \underline{u} + (1 - \alpha)\Delta \leq \bar{u} \quad (2)$$

$$\tilde{u}(\underline{s}|\bar{s}) = \underline{u} + \alpha\Delta \geq \underline{u}. \quad (3)$$

Both inequalities hold strictly provided $\alpha > 0$. As is common in the literature, we assume that a firm is aware of consumers’ preferences. This means that it knows the degree of the projection bias and can set its price accordingly to maximize its profits.

Timing of the game In the baseline model, the firm sets a price p at the start of the first period. Given the price and motivation state, consumers predict their second period consumption utility based on their current motivation state, and decide whether to purchase the good. In the case of purchase, they pay p and receive the good in the second period, at which point they consume it.

We analyze two model variants that incorporate “cooling off” policies motivated by real world policy interventions. First, we consider a *mandatory cooling off period*. Under such a policy, consumers cannot make a final purchase decision in the first period. Instead, based on their first-period motivation and the price p , they may choose to sign a letter of intent. If they do so, they can either confirm the purchase or step back from it in the second period *after* observing their second-period motivation. A transaction takes place only if the intent to purchase is confirmed, in which case consumers pay p in the second period and consume the good once it arrives in the third period. As such, consumption is delayed by one period.

⁵Our results are robust to allowing for autocorrelated motivation states, see [Section 5](#).

Second, we consider a voluntary *return policy*. Under this policy, the first period is identical to the baseline model. However, consumers have the opportunity to return the good before consumption occurs but after they observe their second-period motivation state. By returning the good, they are reimbursed the price p and forgo consumption.⁶ In the main analysis, we abstract from return costs, but discuss the implications of costly return in [Section 5](#).

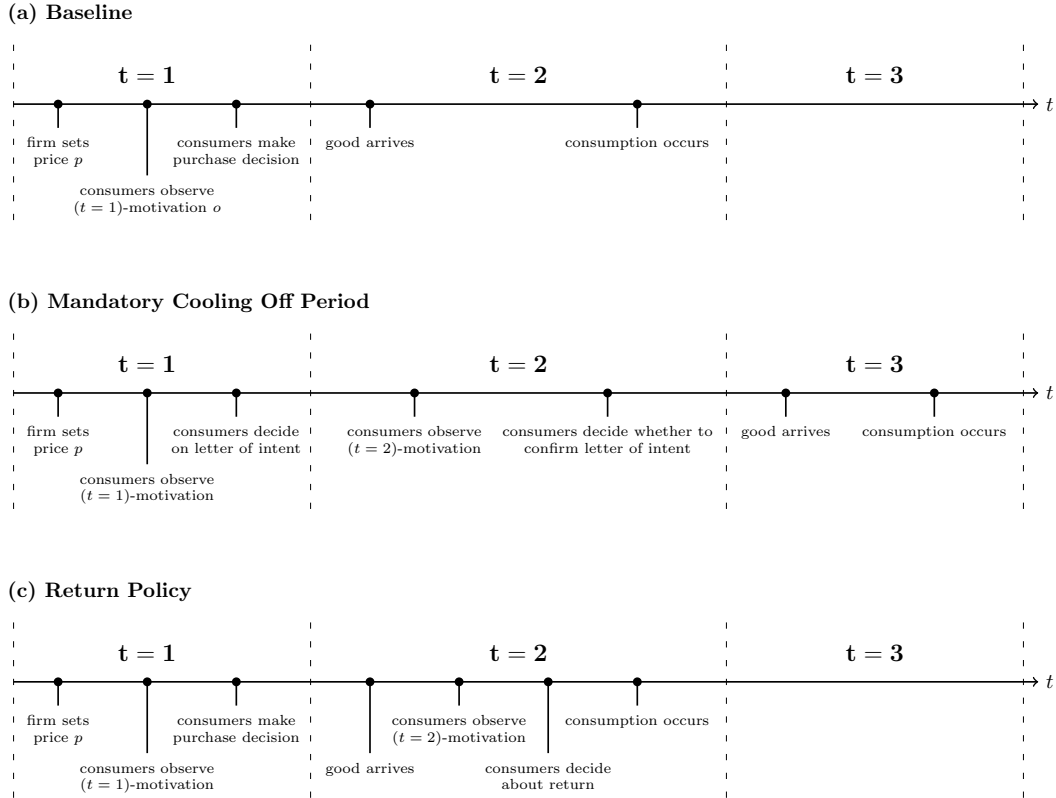


Figure 1: Timing of the Game

A key difference between the two policies lies in what factors into consumers' final decision of whether to consume the good. In both the baseline and the cooling off variant, consumption occurs one period after the final purchase decision and confirmation decision, respectively. The final decision consumers make is thus based on predicted consumption utilities - the projection bias leaves scope for firms extracting more than consumers' expected utility leading to negative expected consumer welfare.⁷ By contrast, only the initial purchase decision is based on predicted utilities when a return policy is in place. When consumers decide whether to return the good or consume it in the second period, they know their second-period motivation and correctly infer their consumption utility. Negative expected consumer welfare will thus never occur under this

⁶For analytical simplicity, we make the assumption that consumers perfectly know their utility in case of consumption when making the return decision. This assumption can easily be relaxed as long as the prediction about the actual consumption utility becomes more accurate because of the return period.

⁷Our analysis abstracts from informational gains for the consumer during the cooling off period, such that changes in preferences are solely driven by a change in their motivational state. Incorporating such gains would magnify the effect of a cooling off period without altering the results qualitatively.

policy. [Figure 1](#) contrasts the timeline of the game in the three settings.

Overall, our model allows for a simple analysis of the outcomes of two cooling-off policies that have real-world applicability and work via different channels, i.e. either via delaying the purchase decision (mandatory cooling-off period), or by allowing to reverse the purchase decision (return policy). The model nests both standard rational preferences and nonstandard preferences and can capture purchases based on rational state-dependent expectations, biased impulse purchases, and overly high consumer skepticism.

3 Model Analysis

In this section, we provide the analysis of our model. Detailed proofs are relegated to [Appendix A](#), which also contains explicit derivations for demand as a function of price. For expositional purposes, we fix the discount factor between periods $t = 1$ and $t = 2$ as $\delta_1 = 1$. Because consumption never occurs before the second period, this allows us to eliminate the carrying of discount factors in the pricing, profit, and welfare expressions without affecting any qualitative insights. This leaves a trade-off between consumption in period $t = 2$, and a delay in consumption until period $t = 3$ triggered by a mandatory cooling off period. To simplify notation, we denote the remaining discount factor by $\delta \equiv \delta_2$. Throughout, we present our results in terms of marginal cost thresholds which determine the firm's optimal pricing. This relates to a large part of the empirical industrial organization literature which assumes that marginal costs are not directly observable to researchers ([Bresnahan, 1982](#); [Berry et al., 1995](#)).

Both in the baseline setting, and with a cooling off period, consumer decisions are based on the predicted *expected* consumption utility. We denote by \tilde{u} and \underline{u} the predicted expected gross utility of consumption in a future period for a motivated and unmotivated consumer, respectively. As previously discussed, consumers correctly assess the probability μ of being motivated, but mispredict the utility in the state they are not currently experiencing due to the projection bias. This gives

$$\tilde{u} \equiv \mu \tilde{u}(\bar{s}|\bar{s}) + (1 - \mu) \tilde{u}(\underline{s}|\bar{s}) = \underline{u} + \alpha \Delta + \mu(1 - \alpha) \Delta \quad (4)$$

$$\underline{u} \equiv \mu \tilde{u}(\bar{s}|\underline{s}) + (1 - \mu) \tilde{u}(\underline{s}|\underline{s}) = \underline{u} + \mu(1 - \alpha) \Delta. \quad (5)$$

For $\alpha > 0$, a motivated consumer's predicted expected utility thus exceeds the actual expected utility, $\tilde{u} > E[u]$, because of the bias in predicting the consumption utility when unmotivated. Similarly, a currently unmotivated consumer's predicted expected utility is lower than the actual expected consumption utility, $\underline{u} < E[u]$.

Equilibrium concept We look for Subgame Perfect Nash Equilibria given that consumers have a projection bias, that is, the firm maximizes its profits while consumer behavior is optimal given the potentially biased predicted expected and state-contingent utilities. This implies that

the predicted second-period decisions which shape the first-period purchase decision may be different than the decisions consumers actually take.⁸

Assumptions We make the tie-breaking assumption that in case of equal profits, the firm prefers the strategy that yields the larger market share. For consumers, we assume that following a policy intervention they initially purchase only if they predict to consume the good with positive probability. Finally, we assume that the marginal cost c is such that the firm always has at least one profitable pricing strategy in each policy regime. Sufficient for this is $c < \delta \tilde{u}$. Removing the assumption would not qualitatively alter the main results, but leads to additional trivial cases in which the firm does not sell.

Notation Throughout the analysis, both a subscript or superscript \sim denotes that pricing is based on predicted utilities, while both a subscript or superscript $-$ indicates that pricing is based on actual consumption utilities. A superscript generally denotes firms targeting motivated consumers, while a subscript refers to firms targeting both consumer types. To illustrate, \tilde{p} denotes the price when motivated consumers are targeted and pricing is based on *predicted* utilities, while \underline{CW} denotes the consumer welfare when all consumers, including unmotivated, are targeted and pricing is based on *actual* utilities.

3.1 Baseline

Absent any policy, consumers make their purchase decision based on their first-period motivation. This leads to two possible pricing strategies for the firm. The firm can target first-period-motivated consumers only, or it can price such that all consumers purchase the product. In both cases, the firm fully extracts consumers’ future predicted expected utility. This is because consumers are homogeneous conditional on their motivation state. We discuss extensions with heterogeneous consumers in [Section 5](#).

Targeting only motivated consumers (“exclusive targeting”) allows for the firm to benefit from their upward projection bias. However, this comes at the cost of only serving a fraction μ of the potential market. In contrast, targeting all consumers (“full targeting”) captures a larger market share, but requires pricing below the actual expected utility because unmotivated consumers have a downward bias. The firm prefers to target motivated consumers only and thus forego the additional demand at a lower price if and only if the marginal cost of production is sufficiently high, i.e., if it exceeds a threshold \tilde{c} . We obtain the following proposition.

Proposition 1 (Baseline Pricing & Welfare) *There exists a unique threshold $\tilde{c} = \underline{u} + \frac{\mu}{1-\mu}\Delta(1 - 2\alpha - (1 - \alpha)\mu) < E[u]$ which determines the firm’s pricing strategy.*

⁸For example, in case of a return policy, initially unmotivated consumers underpredict the consumption utility in the motivated state and hence *predict* to return the good regardless of motivation if the price exceeds $\tilde{u}(\bar{s}|\underline{s})$ —however, they would only return the good when motivated if the price exceeds \bar{u} as they know their consumption utility at this stage.

- (i) If $c \leq \tilde{c}$, the firm targets all consumers and charges $\underline{p} = \underline{u}$, leading to firm profits of $\underline{\pi} = \underline{u} + \mu(1 - \alpha)\Delta - c$, consumer welfare $\underline{CW} = \mu\alpha\Delta$, and total welfare $\underline{TW} = \underline{u} + \mu\Delta - c$.
- (ii) Otherwise, the firm targets the mass μ of first-period-motivated consumers only and charges $\tilde{p} = \tilde{u}$. Profits are $\tilde{\pi} = \mu(\underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta - c)$, consumer welfare is $\widetilde{CW} = -\mu(1 - \mu)\alpha\Delta \leq 0$, and total welfare is given by $\widetilde{TW} = \mu(\underline{u} + \mu\Delta - c)$.

Proof. See [Appendix A.1](#). ■

The firm's targeting decision has direct implications for welfare. As motivated consumers overpredict their expected utility relative to the actual expected utility because of the projection bias, and unmotivated consumers underpredict it, consumer welfare is weakly (strictly for $\alpha > 0$) negative in the case where only motivated consumers are targeted, and weakly (strictly for $\alpha > 0$) positive if the full market is covered. Moreover, total welfare may be negative if motivated consumers are exclusively targeted—their overprediction of the future consumption utility allows for the firm to profitably produce even if the marginal cost exceeds the actual expected utility of consumption. As exclusive targeting always leads to negative consumer welfare for $\alpha > 0$, introducing a policy can be necessary to prevent consumers from having negative expected consumption utility.

3.2 Mandatory cooling off period

With a mandatory cooling off period the initial decision to register for purchase in the first period and the confirmation decision in the second period are based on predicted expected utilities. The projection bias thus leaves scope for the firm extracting more than the expected consumption utility from motivated consumers even with the policy in place. The registering and confirmation decisions are straightforward. A consumer registers if and only if the predicted consumption utility given her first-period motivation weakly exceeds the price, and confirms a purchase if and only if the same holds given her second-period motivation. This is because in the first period, unmotivated consumers do not realize their projection bias and thus do not think it is worth registering because they do not think that their prediction about the usefulness of the product in the next period can increase.

It follows that the firm has two potential targeting strategies; it may either induce only first-period-motivated consumers to register, anticipating that they only confirm if they are also second-period-motivated. Alternatively, the firm can initially target all consumers, anticipating that they all subsequently confirm their purchase. The candidate prices are thus the same as in the baseline case, but weighted with the discount factor δ as consumption is delayed by one period. Similar to the baseline case, there is a unique threshold \tilde{c}_c such that the firm prefers the exclusive targeting if and only if the marginal cost of production exceeds \tilde{c}_c .

Proposition 2 (Pricing & consumer welfare with cooling off period)

There exists a unique threshold \tilde{c}_c with $\tilde{c}_c = \underline{u} + \frac{\mu}{(1-\mu)}\Delta \frac{(1-(1+\mu)\alpha-(1-\alpha)\mu^2)}{1+\mu} \in [\tilde{c}, E[u]]$ which determines the firm's pricing strategy in the presence of a cooling off period.

(i) If $c \leq \tilde{c}_c$, the firm targets all consumers and charges $\underline{p}_c = \delta\underline{u}$, leading to purchase confirmation of all consumers and firm profits of $\underline{\pi}_c = \delta(\underline{u} + \mu(1-\alpha)\Delta - c)$, consumer welfare $\underline{CW}_c = \delta\mu\alpha\Delta$, and total welfare $\underline{TW}_c = \delta(\underline{u} + \mu\Delta - c)$.

(ii) Otherwise, the firm targets first-period-motivated consumers only who confirm their purchase if and only if they are also motivated in the second period. The firm charges $\tilde{p}_c = \delta\tilde{u}$, which leads to profits $\tilde{\pi}_c = \delta\mu^2(\underline{u} + \alpha\Delta + \mu(1-\alpha)\Delta - c)$, consumer welfare $\widetilde{CW}_c = -\delta\mu^2(1-\mu)\alpha\Delta \leq 0$, and total welfare $\widetilde{TW}_c = \delta\mu^2(\underline{u} + \mu\Delta - c)$.

Proof. See [Appendix A.2](#). ■

The threshold \tilde{c}_c lies above the one in the baseline case, $\tilde{c}_c \geq \tilde{c}$, with $\tilde{c}_c > \tilde{c}$ for $\alpha > 0$. Only twice-motivated consumers confirm their purchase decision with exclusive targeting, which decreases the attractiveness of this strategy relative to the baseline case as the quantity decreases. Consumer welfare and total welfare behave similar to the baseline case: for $\alpha > 0$, consumer welfare is negative under exclusive targeting, and positive with full market coverage. Total welfare may be negative in case motivated consumers are exclusively targeted.

When assessing the introduction of a cooling off period, [Figure 2](#) illustrates the possible effects. If the market is fully covered both pre and post policy introduction, effects are limited in that prices, consumer welfare, and total welfare are only negatively affected by the discount factor δ . Similarly, the price change in case exclusive targeting materializes both pre and post policy introduction is limited to the discount factor but still induces a substantial positive change in consumer welfare—motivated consumers still have a negative expected consumer welfare, but there is a significant drop in demand as only twice-motivated consumers end up purchasing the good. The policy may succeed in switching the firm's targeting strategy from exclusive targeting to full market coverage, which leads to a decrease in price, and an increase in quantity and consumer welfare.

Similarly, total welfare is increased by the policy in this case as full coverage post introduction requires $c < E[u]$ such that the increased coverage is beneficial for total welfare. By contrast, the effect is not clearly determined if exclusive coverage materializes even post introduction—depending on the relation between the production cost c and actual expected utility $E[u]$, the decrease in quantity can be either harmful ($c < E[u]$) or beneficial ($c > E[u]$).

Cooling Off Period

		Full Targeting	Exclusive Targeting
B a s e l i n e	Exclusive Targeting	$\Delta q > 0$ $\Delta p < 0$ $\Delta CW > 0$ $\Delta TW > 0$ no stepping back	$\Delta q < 0$ $\Delta p \approx 0$ $\Delta CW > 0$ ΔTW ambiguous stepping back
	Full Targeting	$\Delta q = 0$ $\Delta p \approx 0$ $\Delta CW \approx 0$ $\Delta TW \approx 0$ no stepping back	

The figure depicts the directional changes in prices, consumed quantities, consumer welfare, and total welfare for the different combinations of targeting strategies when moving from the baseline setting to a setting including a mandatory cooling off period.

Figure 2: Effects of the introduction of a cooling off period

3.3 Return policy

The return decision is based on consumers' actual consumption utility as it takes place after the second-period motivation has been realized. As such, consumers keep the good if and only if their actual consumption utility is at least as high as the price. By contrast, the initial purchase decision is driven by the predicted utilities in each motivation state, which incorporate the projection bias.

A consumer initially purchases if and only if she believes to consume the good in at least one motivation state, i.e., if the price the firm charges is equal to or below the predicted consumption utility in the high motivation state. This highlights an important difference to the baseline and the cooling off period: the candidate prices are given by the (predicted) utilities in a given motivation state instead of the (predicted) expected utilities. The firm's pricing depends on the marginal cost of production, and involves two cutoffs \underline{c}_r and \tilde{c}_r , where $\underline{c}_r \leq \tilde{c}_r$. If the marginal cost of production is below \underline{c}_r , the firm prefers to serve the full market. This implies that even unmotivated consumers do not return the good in the consumption period and thus imposes a maximal price of \underline{u} .

If the marginal cost is above \tilde{c}_r , the firm instead targets only initially motivated consumers, who return the good unless they are again motivated in the second period. The firm fully extracts the consumers' rent by charging \bar{u} —initially motivated consumers purchase as they

predict the utility in the motivated state to be $\tilde{u}(\bar{s}|\bar{s}) = \bar{u}$, and confirm upon again being motivated in the second period as the price equals their actual consumption utility. Finally, the firm may use a third strategy for intermediate marginal cost: it initially sets the price such that even unmotivated consumers purchase the good but induces return unless consumers are second-period-motivated. Pricing is therefore based on the utility unmotivated consumers predict for the motivated state, $\tilde{u}(\bar{s}|\underline{s})$.

Proposition 3 (Firm’s pricing given return policy) *There exist thresholds \tilde{c}_r and \underline{c}_r where $\underline{c}_r \leq \underline{u}$, $\underline{c}_r \leq \tilde{c}_r$, and $\tilde{c}_r \geq \tilde{c}$, which determine the firm’s pricing decision in presence of a return policy. $\underline{c}_r < \tilde{c}_r$ if and only if $\alpha < \frac{1}{1+\mu}$.*

(i) *If $c \leq \underline{c}_r$, the firm initially targets all consumers who all keep the good. It charges*

$$\underline{p}_r = \underline{u} \text{ to reap profits } \underline{\pi}_r = \underline{u} - c. \text{ Consumer welfare is } \underline{CW}_r = \mu\Delta \text{ and total welfare } \underline{TW}_r = E[u] - c = \underline{u} + \mu\Delta - c.$$

(ii) *If $c \in (\underline{c}_r, \tilde{c}_r]$, the firm initially targets all consumers, who return the good unless they*

$$\text{are motivated in the second period. It charges } \underline{p}_r = \tilde{u}(\bar{s}|\underline{s}) \text{ to reap profits } \underline{\pi}_r = \mu \cdot [\underline{u} + (1 - \alpha)\Delta - c]. \text{ Consumer welfare is } \underline{CW}_r = \mu\alpha\Delta, \text{ and total welfare } \underline{TW}_r = \mu(\underline{u} + \Delta - c).$$

(iii) *If $c \in (\tilde{c}_r, \bar{u}]$, the firm initially targets only motivated consumers, who return the good*

$$\text{unless they remain motivated in the second period. It charges } \bar{p}_r = \bar{u} \text{ to reap profits } \bar{\pi}_r = \mu^2 \cdot (\underline{u} + \Delta - c). \text{ Consumer welfare is } \bar{CW}_r = 0, \text{ and total welfare } \bar{TW}_r = \mu^2(\underline{u} + \Delta - c).$$


Proof. See [Appendix A.3](#). ■

The threshold for exclusive targeting, \tilde{c}_r , lies above that in the threshold \tilde{c} in the baseline case. While the price under exclusive targeting is larger than in the baseline, this is more than offset by the decrease in quantity. The range for targeting all consumers initially while inducing on-path return (“intermediate targeting”), $(\underline{c}_r, \tilde{c}_r]$, may either embed \tilde{c} , lie strictly above \tilde{c} , or be empty.

A return policy ensures that the firm cannot fully extract the motivated consumers’ utility predictions under projection bias—the return decision is based on actual consumption utilities and the bias only matters for the initial purchase decision. As such, consumer and total welfare are necessarily weakly positive. The projection bias is actually beneficial for consumer welfare under the intermediate targeting strategy. A higher bias lowers the price and increases the rent left to second-period-motivated consumers. However, as this decreases the firm’s profits, it simultaneously makes it less likely that this targeting behavior materializes in equilibrium.

The impact of the adoption of a return policy on prices, quantities, consumer and total welfare are summarized in [Figure 3](#). The sign of the changes in price and quantity can be signed conditional on the targeting strategies except for the case where exclusive targeting pre policy is followed by intermediate targeting post policy. This is because the predicted utility for the motivated state by unmotivated consumers $\tilde{u}(\bar{s}|\underline{s})$ may lie both above and below the predicted expected utility of motivated consumers \tilde{u} .

If consumers have a negative expected consumption utility in the baseline case, the introduction of a return policy always increases consumer welfare. It either stops negative expected consumer utility while keeping the same targeting strategy or changes the firm’s targeting strategy in which case consumer welfare post-introduction is positive. Notably, this can materialize even if prices increase, as exclusive and intermediate targeting following the policy introduction ensure that only second-period-motivated consumers actually consume the good. If the firm targets all consumers before the policy introduction, a return policy increases consumer welfare if second-period-unmotivated consumers keep the good, as the associated price decrease to \underline{p}_r leaves a significant rent to second-period-motivated consumers. By contrast, it is consumer welfare neutral if the firm changes its pricing to induce on-path return—the higher rent to second-period-motivated consumers is offset by the lower market coverage.

		<u>Return Policy</u>		
		Full Targeting	Intermediate Targeting	Exclusive Targeting
B a s e l i n e	Exclusive Targeting	$\Delta q > 0$ $\Delta p < 0$ $\Delta CW > 0$ $\Delta TW > 0$ no return	$\Delta q = 0$ Δp ambiguous $\Delta CW > 0$ $\Delta TW > 0$ observe return	$\Delta q < 0$ $\Delta p > 0$ $\Delta CW > 0$ ΔTW ambiguous observe return
	Full Targeting	$\Delta q = 0$ $\Delta p < 0$ $\Delta CW > 0$ $\Delta TW = 0$ no return	$\Delta q < 0$ $\Delta p > 0$ $\Delta CW = 0$ ΔTW ambiguous observe return	

The figure depicts the directional changes in prices, consumed quantities, consumer welfare, and total welfare for the different combinations of targeting strategies when moving from the baseline setting to a setting including a voluntary return policy.

Figure 3: Effects of the introduction of a return policy

The effect of a return policy on total welfare is more ambiguous. Total welfare is unchanged if the full market was covered both pre and post policy, and increases if the return policy leads to a change from exclusive targeting to initial full coverage with induced return. In all other

cases, the total welfare impact cannot be generically signed. This is because it depends on the relation between the marginal cost of production c and the expected consumption utility $E[u]$. If $c < E[u]$, the larger market share in the unregulated case is desirable such that a return policy may be harmful for total welfare. We summarize these considerations in [Figure 3](#).

Voluntary adoption of a return policy One important question is whether a policymaker even needs to consider imposing a mandatory return policy. Costless return is offered by a variety of firms in practice, and it could in principle be the case that there is no scope for such a policy because any time consumer and/or total welfare would be increased, the firm would already find it beneficial to adopt it voluntarily. We therefore demonstrate in [Appendix A.5](#) that even if firms have the option to voluntarily adopt a return policy, they do not do so in all cases where consumer or total welfare would benefit from it.

3.4 Comparison between policies

In practice, oftentimes only either a cooling off period or a return policy is applicable. However, there may be scenarios in which both policies are in principle implementable. It is thus of interest to compare the efficacy of these policies in terms of consumer or total welfare.

From a consumer welfare perspective, a return period might at first glance appear to be the preferred policy. Absent return costs, it never harms welfare and moreover fully avoids negative expected consumer utility by basing the return decision on actual consumption utilities. By contrast, a cooling off period may be harmful to welfare by delaying consumption. However, even in the case where return is privately and socially costless, we can identify cases where a cooling off period outperforms a return policy in terms of consumer welfare.

Proposition 4 (Cooling off, return policy & consumer welfare)

The threshold cost for exclusive targeting of twice motivated consumers is higher under a mandatory cooling off period than under a return policy, $\tilde{c}_c > \tilde{c}_r$. For $c \in (\tilde{c}_r, \tilde{c}_c)$, a mandatory cooling off period always leads to higher consumer welfare than a return policy, and may lead to higher total welfare.

Proof. We formally establish $\tilde{c}_c > \tilde{c}_r$ in [Appendix A.4](#), which also shows that total welfare can be higher under a cooling off period. For $c \in (\tilde{c}_r, \tilde{c}_c)$, a return policy leads to full rent extraction, while a cooling off period leaves positive CW to consumers. ■

This may occur whenever the firm targets only motivated consumer in the baseline, i.e. whenever $c > \tilde{c}$. We establish that it is possible that full market coverage may occur provided that a cooling off policy is in place—which implies that consumer welfare is positive—while under a return policy the firm would still target only twice-motivated consumers and charge a price equal to their actual consumption utility.

The reasoning for why exclusive targeting is more attractive under a return policy than a cooling off period lies in the firm’s pricing. In both cases, exclusive targeting leads to a final quantity of μ^2 as only twice motivated consumers end up consuming the good. However, with a return policy in place, the firm can extract the actual consumption utility \bar{u} , as the consumers *know* their consumption state when making the return decision (and correctly predict it in the initial period). By contrast, the firm charges the predicted expected utility \tilde{u} of motivated types with a cooling off period. Consumer welfare is hence lower under exclusive targeting with a cooling off period— $(1 - \mu)$ of purchasing consumers end up being unmotivated in the consumption period—but it is at the same time less profitable for the firm to use this strategy.

This shows that cooling off periods may be preferable despite the implied delayed consumption even in the *best-case-scenario* of a return policy, that is, absent return costs. But the implementation of a cooling off period also comes with a drawback. Even if the policymaker is able to identify that provision of the good is restricted to motivated consumers before any policy introduction, it is possible that the same holds true after the introduction of the policy, which would preserve negative consumer welfare. This could be avoided entirely with a return policy.⁹

We believe that our model can easily be adapted to capture situations of interest for policymakers such as “pressure selling” or “fooling the elderly” either via telephone or via door-to-door. An easy way to capture such considerations is if we assume that a salesman can increase the probability of being motivated in the first period but not the second one. This in turn increases a firms’ incentive to sell the product at a high price in absence of a policy which—if it materializes—implies scope for a policy intervention to be welfare enhancing.

4 Data requirements for policy evaluation

This section uses our model analysis to determine the data requirements which allow us to unambiguously assess the different policies’ effects on several economic indicators ex-post. In doing so, we exploit the changes in the firms’ optimality condition before and after the policy adoption. As indicators we consider the sign changes in consumer and total welfare, respectively. Furthermore, we discuss under which circumstances the fraction of motivated consumers μ can be recovered, and whether it can be established that consumers experience a positive bias, i.e., whether $\alpha > 0$. Naturally, the requirements always depend on the specific model. We assess the robustness of this approach using different model specifications in [Section 5](#).

Our focus is on the use of aggregate instead of individual-level data. While aggregate data are typically less informative, they are also much easier to obtain. We therefore explore

⁹A return policy never decreases consumer welfare in our model provided that return is privately and socially costless. There is a non-empty parameter range where the policy is simply consumer welfare neutral under these assumptions. In such a case, an adoption of a policy would still lead to additional social costs because of extra administrative processing and political implementation, which is outside of our model. These additional costs could make not adopting a policy a preferable choice.

the extent to which an evaluation can be conducted when only these data are available. We consider the following variables to be available when conducting an ex-post evaluation: prices, final quantities (i.e., quantities sold net of return or stepping back), return frequencies in case of a return policy, and frequency of consumers not confirming purchase in case of a mandatory cooling off period.

To assess the data requirements, we proceed as follows. For a given combination of pre and post policy targeting strategies, the model always allows us to unambiguously sign the change in consumer welfare, ΔCW , and sometimes to sign the change in ΔTW (see [Figure 2](#) and [Figure 3](#)). This is a specific feature of our model, and we discuss the usefulness of the ex-post assessment based on aggregate data when this does not hold in the context of additional consumer heterogeneity in [Section 5](#). We ask whether the available data allows us to identify the combination of targeting strategies, or to at least obtain a restricted set of targeting strategies consistent with the observed data. For this, it is important to take into account the number of potential responses of the firm following the policy adoption—under a return policy, there are three candidate targeting strategies instead of two under a cooling off period. As can be expected, it is more complicated to identify the realized combination of equilibrium behavior from aggregate data the more potential combinations there are. Provided that the combination of targeting strategies can be identified, this can then also be used to recover additional information such as the fraction of motivated consumers or the presence of a projection bias.

One aspect which complicates the assessment based on models with a projection bias is the following. For a given level of the bias, there are always both motivated and unmotivated consumers present in any given period. Because of the projection bias, these consumers—who exhibit the same underlying preferences—behave as if they are of two distinct types. This would e.g. be different in a model with $(\beta - \delta)$ preferences and a single time-inconsistent type.

4.1 Mandatory cooling off period

In [Figure 4](#), we map the effects of the adoption of a cooling off period into a space that relates the implied changes in quantity to the implied changes in price. The different possible combinations that are consistent with our model are labeled as Cells A-C.

For each directional movement in quantity, i.e., for increasing, decreasing, and constant final quantity following a policy introduction, there is a single equilibrium combination of targeting strategies that fits such a movement. This is not the case for the directional movement in price, see Cells A and B. If the price stays relatively constant, this can be because the firm always targets all consumers (Cell B), or because it always targets only motivated consumers (Cell A). We use the term relatively constant because the discount factor δ implied by the waiting period leads to a decrease in the price. The decrease would be the same for both cases and is negligible for a discount factor close to 1.

Only taking into account the stepping back behavior after signing a letter of intent does

	final quantity decreases $\Delta q < 0$	final quantity constant $\Delta q = 0$	final quantity increases $\Delta q > 0$
price increases $\Delta p > 0$			
price relatively constant $\Delta p \approx 0$	(A) observe stepping back Excl \rightarrow Excl: $\Delta CW > 0$ ΔTW ambig.	(B) no stepping back Full \rightarrow Full: $\Delta CW \approx 0$ $\Delta TW \approx 0$	
price decreases $\Delta p < 0$			(C) no stepping back Excl \rightarrow Full: $\Delta CW > 0$ $\Delta TW > 0$

The figure shows the possible cases implied by the model, given different combinations of observable changes in final quantities and prices following the introduction of a mandatory cooling off period.

Figure 4: Ex-post evaluation of cooling off period

not enable us to unambiguously identify the combination of targeting strategies, as multiple combinations feature no stepping back (Cell B and Cell C). This could, however, be facilitated if price data are available provided that the discount factor is such that a (small) price change due to the discount factor (Cell B) can be distinguished from a large price change due to a change in targeting strategies (Cell C).

If the combination of targeting strategies can be identified, the direction of the effect of the policy on consumer welfare is always identified. Moreover, the change in total welfare can be signed unless it is identified that the firm uses the most exclusive targeting strategy both pre and post policy adoption. We can also gather additional information from the ex-post evaluation. If stepping back is observed (occurs only in Cell A), the fraction of motivated consumers μ is given by the ratio of the mass of consumers who consume the good relative to the mass of consumers initially signing a letter of intent. When an increase in quantity is observed (Cell C), μ can also be backed out: in this case, μ exactly equals the ratio of pre and post policy quantities. Finally, partial market coverage pre-policy can only materialize if consumers are biased; as such, identifying Cell A or C allows us to conclude that $\alpha > 0$.

Proposition 5 *Data on quantities or data on both prices and confirmation rates are sufficient to assess the directional change in consumer welfare and the relevant cost range following the adoption of a mandatory cooling off period. In some cases, the existence of a projection bias, i.e. $\alpha > 0$, and the fraction of motivated consumers, μ , can also be recovered.*

Proof. Follows from the previous discussion. ■

Optimality of policy intervention Given that the combination of equilibrium targeting strategies can be identified using the data outlined in [Proposition 5](#), it is important to consider the implications for whether the chosen policy was optimal, or whether another policy may have performed better in terms of consumer welfare. This is a complicated issue as the adoption of a given policy typically precludes learning about the total costs and benefits of other potential policies ([Coase, 1960](#)). In light of [Proposition 4](#), an ex-post assessment is nonetheless partially informative. Unless the cooling off policy has shifted the firm’s targeting strategy from exclusive to full targeting, it was always a suboptimal policy as a return policy would have lead to higher consumer welfare.

Corollary 1 *A mandatory cooling off policy was a potentially optimal intervention if it succeeded in shifting the firm’s targeting strategy from exclusive to full targeting. Otherwise, a return policy would have been better for consumer welfare.*

Proof. Follows from [Proposition 4](#) and [Figure 6](#) in the Appendix. ■

4.2 Return policy

[Figure 5](#) illustrates the different outcomes for the introduction of a return policy. For ease of exposition both [Figure 5](#) and the discussion implicitly assume $\alpha < 1$.

Cell A of the figure shows that if the quantity decreases post-introduction, two cases emerge that are both associated with a price increase. This occurs when motivated consumers are targeted exclusively pre and post policy adoption, and when there is a switch from targeting all consumers to pricing at unmotivated consumers’ predicted utility for the high motivation state $\tilde{u}(\bar{s}|s)$ while inducing return by second-period unmotivated consumers.

When the quantity stays constant after the policy introduction, three pricing patterns can emerge, see Cells B-D. When the firm changes from exclusive to intermediate targeting, this can lead to an increase in the price (Cell B), leave the price constant (Cell C), and lead to a price decrease (Cell D). However, Cell D also contains a second case in which the firm target all consumers both before and after the policy introduction. Cell E shows that if the quantity increases post-introduction, this is associated with the firm switching from targeting only motivated consumers to targeting all consumers, which also leads to a price decrease.

Because Cells A and D each contain multiple cases that lead to the same price-quantity predictions, combinations of price and quantity data alone cannot distinguish among all the emerging equilibrium targeting combinations. Data on return frequencies can distinguish between the two cases in Cell D, i.e., resolve the issue when there is no change in the final quantity following the policy adoption. This is because when all consumers are targeted both with and

	final quantity decreases $\Delta q < 0$	final quantity constant $\Delta q = 0$	final quantity increases $\Delta q > 0$
price increases $\Delta p > 0$	(A) <u>observe return</u> Excl → Excl: $\Delta CW > 0$ ΔTW ambig. Full → Int: $\Delta CW = 0$ ΔTW ambig.	(B) <u>observe return</u> Excl → Int: $\Delta CW > 0$ $\Delta TW > 0$	
price constant $\Delta p = 0$		(C) <u>observe return</u> Excl → Int: $\Delta CW > 0$ $\Delta TW > 0$	
price decreases $\Delta p < 0$		(D) <u>observe return</u> Excl → Int: $\Delta CW > 0$ $\Delta TW > 0$ <u>observe no return</u> Full → Full: $\Delta CW > 0$ $\Delta TW = 0$	(E) <u>observe no return</u> Excl → Full: $\Delta CW > 0$ $\Delta TW > 0$

The figure shows the possible cases implied by the model, given different combinations of observable changes in final quantities and prices following the introduction of a voluntary return policy for $\alpha \in [0, 1)$.

Figure 5: Ex-post evaluation of return policy

without policy, there should not be any return by consumers, while any other form of targeting should involve product return with positive probability. By contrast, the two cases in Cell A are both associated with positive product return. An additional assumption that would enable us to distinguish between the cases in Cell A is that the researcher knows the total market size. This is because before the policy the market is only fully covered in one of the two cases.¹⁰ Provided that the combination of targeting strategies is identified, the direction of the change in consumer welfare is always identified, while that of the change in total welfare is identified unless a switch from full to intermediate targeting is observed.

Finally, whenever the data reveals that the firm used exclusive targeting before the policy introduction it follows that consumers experience a positive projection bias, $\alpha > 0$. Unless the market is fully covered both pre and post policy introduction, the fraction of motivated consumers μ can be recovered. For example, both combinations in Cell A have the feature that the ratio of pre and post policy quantities consumed exactly equals μ , while the same holds for the ration of post and pre policy quantities consumed in Cell E.

Proposition 6 *Data on quantities, return frequencies and market size are sufficient to assess*

¹⁰Modern empirical models in industrial organization also make the assumption that the total market size, i.e. the total number of potential consumers, is known to a researcher. This assumption seems to be even stronger in our context because it implies that in some cases all potential consumers already consume the product.

the directional change in consumer welfare and the relevant cost range following the adoption of a return policy. In some cases, the directional change of total welfare, the presence of a positive projection bias, $\alpha > 0$, and the fraction of motivated consumers, μ , can also be recovered. Data on quantities, prices and return frequencies are not sufficient to always assess the directional change in consumer welfare.

Proof. Follows from the preceding discussion. ■

An interesting observation is that conditional on the availability of data on quantities, return behavior, and market size, data on prices are not necessary for the identification argument. This is because the cases with constant quantities, B-D, obtain from only two equilibrium combinations of targeting strategies which can be distinguished with data on return behavior, and is important once we discuss model extensions featuring costly return in [Section 5](#). Nonetheless, price data by itself can be informative. For example, when prices are found to have decreased, consumer welfare has unambiguously increased due to the policy adoption. In this case, the return policy was also unambiguously (weakly) optimal, as we discuss below.

Optimality of policy intervention As with the cooling-off policy, we can assess the optimality of the adoption of a return policy. A return policy is unambiguously optimal whenever the full market is catered to post policy adoption as this maximizes consumer welfare. If the return policy leads to intermediate targeting ex-post, it was the uniquely optimal policy if the firm used exclusive targeting pre-intervention, and is jointly optimal with no intervention if the market was already fully covered. Finally, a cooling off period may have been better in case the firm targets motivated and twice-motivated consumers exclusively pre and post intervention, respectively.

Corollary 2 *If the firm uses the most exclusive targeting strategy pre and post adoption of a return policy, a cooling off period may have been the better policy in terms of consumer welfare. In all other cases, adopting the return policy was optimal, while no intervention would have led to the same outcome in terms of consumer welfare whenever the firm shifts from full market coverage to intermediate targeting.*

Proof. Follows from [Proposition 4](#) and [Figure 6](#) in the Appendix. ■

5 Robustness & Further Considerations

The main model specification is deliberately stylized as it allows for a clean exposition of both the theoretical model analysis and the evaluation approach. It is therefore important to examine

the robustness of both the model’s predictions and the data requirements to assess the impact of the policy. To address this, we consider several modifications. In what follows, we summarize the results, and refer to the web appendix for the detailed analyses.

Autocorrelated motivations The main model specification assumes that motivation states are independent across periods. In practice, it is entirely possible that motivation states are somewhat persistent, that is, that a consumer who is motivated (unmotivated) today is more likely to be motivated (unmotivated) in the next period. We address this issue by considering a model variant in which consumer motivation is persistent with probability $\rho \in [0, 1)$. We show that the theoretical predictions in the baseline setting and with a mandatory cooling off period are unchanged, and that the same combinations of targeting behavior pre and post policy adoption, as well as the associated price, quantity and consumer welfare movements obtain. This implies that the ex-post evaluation of whether consumer welfare increased or decreased due to the policy can be conducted under the same data requirements as in the main model specification: data on quantities alone, or data on prices and stepping back behavior suffices.

With a return policy, there are two changes relative to the main model specification. First, it is possible that the firm exclusively caters to twice motivated consumers following the adoption of a return policy even if the market was fully covered prior to the adoption. This is because the persistence in motivations makes the exclusive targeting strategy relatively more profitable than the intermediate targeting strategy and hence expands the range such that exclusive targeting is optimal. This combination of strategies is associated with a negative effect on consumer welfare. Second, whenever the firm uses the intermediate targeting strategy post policy adoption while the market was fully covered pre policy, the policy is no longer neutral in terms of consumer welfare but also negatively affects it. This is because the price under full market coverage adjusts downward as unmotivated consumers correctly predict that they are more likely to again be unmotivated in the second period; as this price is the basis of the targeting strategy, consumer welfare increases relative to the main model specification. Despite these changes, the data requirements to identify the sign of the effect of the policy on consumer welfare are unchanged. Quantity data, data on return behavior, and knowledge of market size together allow an assessment of the efficacy of the policy. A given combination of targeting behaviors is always associated with a unique sign of the change in consumer welfare, and the newly arising case can be identified by noting that it is the only combination of targeting strategies where the initially purchased quantity falls below the one sold prior to the policy adoption.

Costly return The main model specification presents the best possible case for a return policy by abstracting from any form of private or social return cost. However, return in practice is likely to be costly and comprise various dimensions such as shipping or hassle costs. We therefore analyze a model variant in which return is costly. There are three main takeaways.

First, the theoretical predictions largely carry over, with the exception of a novel case in which full targeting pre adoption of a return policy is followed by exclusive targeting post

adoption. As with autocorrelated motivations, this is driven by the fact that return costs render the exclusive targeting strategy relatively more profitable for the firm than the intermediate targeting strategy in which more consumers return the good. If this combination of targeting strategies materializes as optimal, a return policy has a negative impact on consumer welfare.

Crucially, these predictions do not depend on whether consumers or the firm directly bear the return costs. As long as consumers correctly anticipate the cost, the fact that the firm fully extracts the (predicted) rents of the targeted consumer group implies that it ultimately is born by the firm so that the targeting strategies, quantities, and also consumer welfare pre and post policy are identical in the two settings. However, there is more ambiguity with respect to the potential direction of price movements if consumers bear return costs, as the price of the product needs to be lowered. This implies that a price decrease following the policy adoption could be driven either by a switch in targeting strategies, or by the need to be compensated for costly return, which potentially complicates the ex-post assessment.

The data requirements to assess the efficacy of a return policy in terms of the impact on consumer welfare are unaffected by the presence of return costs or who bears them. This fact relies on two observations. On the one hand, the potential newly arising case of full market coverage being followed by exclusive targeting can always be identified using data on quantities and return behavior as it is the only combination of targeting strategies associated with a lower initially purchased quantity relative to the pre-policy period. On the other hand, price data are not necessary to assess the efficacy of the return policy for consumers; the identification requires data on quantities, return behavior, and market size. As such, the added ambiguity with respect to price movements due to the policy adoption is not an issue and the data requirements are unchanged relative to the main model specification.

Consumer heterogeneity In the main model specification, heterogeneity in consumers' valuation for the good is driven only by consumers' motivation states. We provide analyses of model variants in which consumers differ in the degree of their bias, and in which their preferences are horizontally differentiated, respectively. In general, incorporating consumer heterogeneity into the theoretical analysis is straightforward; it simply increases the number of candidate strategies for the firm which need to be compared. The resulting increasing number of equilibrium strategies, however, complicates the ex-post assessment, in part because of the aforementioned issue that consumers who exhibit the same bias nonetheless behave as if two different types were present in any given period.

First, we consider a model variant that incorporates a fraction of unbiased consumers in the population. Their presence provides the firm with a new targeting strategy: both in the baseline and if a mandatory cooling off period is adopted, pricing may be based on the (correctly predicted) expected utility of unbiased consumers. In case of a return policy, unbiased consumers behave like motivated biased consumers because the predicted utility in the high state is identical and drives the initial purchase decision, while return decisions are based on

actual consumption utilities so that potential biases do not matter.

The theoretical predictions of the model naturally adapt to the availability of additional targeting strategies for the firm, but overall largely carry over qualitatively. The main difference is that a return policy may now lead to more exclusive targeting post policy adoption, so that it may negatively affect consumer welfare—while the policy avoids negative consumer welfare by aligning the final consumption decision with knowledge about the consumption utility, it may be the case that consumers are losing the rent they enjoyed because pre-policy pricing was based on unmotivated biased consumers who underpredict their expected consumption utility.

With respect to using aggregate data to identify the impact of the policy change, the data requirements to identify the sign of the change in consumer welfare increase but exhibit a similar flavor to the main model. In particular, the requirements to identify the sign change in consumer welfare due to the policy introduction are higher for a return policy than for a mandatory cooling off period. In both cases, aggregate data on prices, quantities and return/stepping back behavior need to be complemented with additional information about the fraction of motivated consumers or even the market size to fully identify the sign of the change in consumer welfare. However, the aforementioned data is sufficient to identify a subset of the potential cases, and always allow an assessment of whether consumer welfare was negatively affected.

Second, we explicitly introduce consumer heterogeneity in terms of idiosyncratic tastes. We model this by considering the firm to be located at the end of a Hotelling line of sufficient length. Consumers are evenly distributed along the line. All consumers exhibit a projection bias, and their consumption utilities are as in the main model. However, they suffer a disutility equal to their distance from the firm.

In this setup, the firm's price is determined not by the identical willingness-to-pay of the consumer group (in terms of bias and motivation) which is targeted with a given pricing strategy, but instead via a first-order condition which trades off the within-targeting-group sensitivity of demand to changes in the price. This complicates both the theoretical analysis as well as the potential identification via aggregate data following a policy adoption. The reason is that not only may the policy induce a change in the targeting behavior (based on which consumer group is the price predominantly set), but may simultaneously affect the within-group responsiveness of demand to price changes. As such, aggregate changes in demand and price together with stepping back/return behavior are no longer sufficient to back out the firm's targeting strategy, which in the previous analyses was sufficient to assess e.g. the directional change in consumer welfare. Moreover, even given a particular combination of pre- and post-adoption targeting the directional change in consumer welfare is not necessarily unambiguously determined.

This highlights two key determinants of how difficult a given policy is to evaluate. The first determinant is how many potential responses to the policy adoption the firm has. As the number of possible combinations increases—which is naturally affected by the model under consideration—the more complicated it is to identify the realized combination of equilibrium behavior from aggregate data. The second determinant is whether the impact on economic indi-

cators such as consumer welfare can be generically assessed *conditional* on a given combination of equilibrium behavior by the firm.

Despite the increased ambiguity, our approach of using aggregate data to identify the combination of targeting strategies can still be used as a screening device. In the case of a mandatory cooling off period, the combination of targeting strategies can always be identified using a combination of price and quantity data. As consumer welfare can only be negatively affected by the policy whenever the firm uses the exclusive targeting strategy both pre and post policy intervention, an in depth investigation using additional individual level data is only necessary if the associated market level outcome of a decreased quantity at a relatively constant price level materializes. Similarly, consumer welfare can only be negatively affected by a return policy if the firm switches from exclusive to intermediate targeting and if the degree of the projection bias is sufficiently low. In this case, the model always predicts an increase in the quantity consumed, as well as a price increase. Thus, only in this case a policy maker would have to look for more detailed data for an in-depth investigation.

Competition between firms We analyze the workings of different consumer policies when only a single firm is present. In many cases, competition among firms can naturally lead to a reduction in prices. This is why we see the presence of cooling off policies to be of particular importance in cases where specific market structures prevent firms from competitive pricing.

There are several examples for such situations. First, initially nonsalient add-on products together with price floors for a base product can lead to quasi-monopoly situations for the add-on market; see, for example, [Heidhues et al. \(2016\)](#). Second, direct selling channels such as door-to-door sales or infomercials at least temporarily result in monopoly conditions that leave consumers only with the decision to purchase a particular product or no product in that category at all.

Relating our model to a competitive setting, it follows that whenever the degree of product differentiation is low, a return policy should be clearly preferred to a mandatory cooling-off policy. This is because in such a case equilibrium prices are close to marginal cost and the return policy in this case aligns the final consumption decision with knowledge about the relevant motivation states, which is beneficial for both consumer and total welfare. If product differentiation is sufficiently high such that each firm has substantial market power, the market is already closer to the one in our main model such that the basic trade-offs between the two policies are present. Data on industry margins, for example from accounting data, or consumer surveys on preferences and consideration behavior could give first indications about which of the cases is more likely, where higher margins should indicate a higher degree of market power.¹¹

¹¹One empirical way to model oligopolistic competition is via the use of a structural model. Most approaches that estimate demand in this literature do not require a formalization of the supply side. In the presence of non-standard preferences it however becomes more difficult to estimate demand because revealed choices do not reveal preferences. In our case of a policy introduction one could combine a demand-estimation with the estimation of the supply side under the assumption that the form of oligopoly competition is known (e.g., multi-product Bertrand Nash pricing) and that the change in pricing of firms both before and after the policy

Additional biases Naturally, consumers could exhibit further biases that affect their behavior towards the cooling off policies, such as inattention towards renewal and return fine print, or present bias towards renewal. This, in combination with whether consumers are aware of their second bias, would move the relative attractiveness of the different policies. For example, depending on whether the confirmation decision is implemented as an opt-in or opt-out (in case of a mandatory cooling off period), the confirmation rates would naturally decrease or increase, respectively. The tradeoffs between the different policies would still be present, however, even though disentangling multiple biases will not be possible using our approach.

Firms' anticipation of policy evaluation When tailoring a policy and data gathering towards ex-post evaluation, there can be concerns that firms may respond strategically to a policy and change their behavior to the detriment of consumers *after* the ex-post evaluation by a policymaker. Firms could thus first behave in a way that a review would suggest there is no further need to intervene in a market after which they would revert back to behavior that is detrimental to consumer welfare. We believe, however, that such a strategy is relatively hard to adopt for firms, and that policymakers have options to make such a strategy even more difficult to sustain. First, such a strategy already implies a sufficient amount of patience from a firm to forego higher profits in the short term to obtain potentially higher profits after an ex-post evaluation has cleared the industry. Second, even if data are gathered by policymakers, an ex-post analysis does not necessarily have to be instituted, and even if it is intended, no particular date needs to be announced ex-ante. This uncertainty lowers the incentives for firms to strategically hide their true incentives in the short run for higher long-run profitability.

Negative model test In the event that the observable data before and after a policy introduction are not consistent with any of the equilibrium outcomes of a specified model, it is implied that the underlying model is not a good industry predictor. An example of this in our case would be an observed price increase following the introduction of a mandatory cooling off period. As this is not consistent with the predictions, it would “falsify” our model. This can be seen as a negative test for the underlying assumptions, and can be helpful in selecting among different model variants when they lead to different predictions. If the data are consistent with one or more combinations of firm strategies and consumer preferences, this naturally does not imply that the model is correct but simply that it is consistent with the observations. With multiple policy introductions or exogenous changes in industry structure, a repeated analysis may allow for distinguishing among different theory models. This can be particularly helpful when trying to distinguish between different types of nonstandard preferences.¹²

introduction is optimal. This constraint leads to at least one extra moment condition. Whether this condition can help to (point- or set-) identify non-standard preferences is an important topic for future research.

¹²This consideration also extends to other policies or regulations than consumer policies. For example, one could think of exploiting changes in wage bonus regulations to test for loss-averse agents (Herweg et al., 2010).

6 Conclusion

Recent years have seen a growing popularity of advocating for and introducing consumer policies based on a “behavioral” justification such as the presence of consumers deviating from fully rational behavior. Because the empirical identification of nonstandard consumer preferences is in many contexts very challenging and oftentimes impossible, a large part of the ex-ante prediction of the effects of policy introductions on market outcomes has to be based on theoretical analyses of models with a behavioral foundation. We argue that an important feature of these models is that they allow an exploration of the combinations of aggregate data pre and post intervention which enable an assessment of the effects of a policy.

We explore the ex ante prediction and ex post evaluation in the context of cooling-off policies when consumers may experience a state dependent projection bias. To do so, we provide a detailed theoretical analysis of two policies prevalent in practice. We characterize the dependence of the firm’s optimal pricing and resulting market outcomes on the regulatory framework and the firm’s cost characteristics. With the characterization of equilibrium strategies pre and post policy adoption, we show which aggregate data are able to distinguish between different potential equilibrium combinations.

The approach has important challenges and limitations. First, such evaluations are inherently model-specific. While this issue is shared with other approaches, it also implies that different model variants can lead to potentially different predictions. Second, in some cases, i.e., with heterogeneity in consumer valuations, aggregate data is not always able to make unambiguous predictions, but can nonetheless flag cases in which a policy is potentially detrimental to consumer welfare and thus warrants further in-depth investigation. Third, while it is easier to obtain aggregate data in most cases, the conclusions that can be obtained from it are naturally much less detailed than those from individual consumer data. We consider a model based evaluation with individual data to be promising for future work. With individual-level data, exploiting equilibrium optimality conditions before and after a policy introduction could be useful to combine with structural econometric models to identify non-standard preferences.

Overall, we believe that taking into consideration the testability of behavioral models after having assessed their theoretical predictions is an important topic. Because of the simplicity of looking at the necessary data requirements from a theoretical viewpoint, an analysis like ours can in principle be easily applied to many if not most models with a behavioral foundation in which policies are assessed, and can provide guidance on gathering the relevant data.

References

- ABALUCK, J. AND A. ADAMS (2017): “What do consumers consider before they choose? Identification from asymmetric demand responses,” Tech. rep., National Bureau of Economic Research.

- ALLCOTT, H., S. MULLAINATHAN, AND D. TAUBINSKY (2014): “Energy policy with externalities and internalities,” *Journal of Public Economics*, 112, 72–88.
- ARMSTRONG, M. (2015): “Search and ripoff externalities,” *Review of Industrial Organization*, 47, 273–302.
- BERNHEIM, B. D. (2009): “Behavioral welfare economics,” *Journal of the European Economic Association*, 7, 267–319.
- BERNHEIM, B. D. AND A. RANGEL (2009): “Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics,” *Quarterly Journal of Economics*, 124, 51–104.
- BERNHEIM, B. D. AND D. TAUBINSKY (2018): “Behavioral public economics,” in *Handbook of Behavioral Economics: Applications and Foundations 1*, Elsevier, vol. 1, 381–516.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile prices in market equilibrium,” *Econometrica: Journal of the Econometric Society*, 841–890.
- BRESNAHAN, T. F. (1982): “The oligopoly solution concept is identified,” *Economics Letters*, 10, 87–92.
- CAMERER, C., S. ISSACHAROFF, G. LOEWENSTEIN, T. O’DONOGHUE, AND M. RABIN (2003): “Regulation for Conservatives: Behavioral Economics and the Case for “Asymmetric Paternalism”,” *University of Pennsylvania Law Review*, 1211–1254.
- CHANG, T. Y., W. HUANG, AND Y. WANG (2018): “Something in the air: Pollution and the demand for health insurance,” *The Review of Economic Studies*, 85, 1609–1634.
- CHETTY, R., A. LOONEY, AND K. KROFT (2009): “Salience and taxation: Theory and evidence,” *American economic review*, 99, 1145–77.
- COASE, R. H. (1960): “The problem of social cost,” in *Classic papers in natural resource economics*, Springer, 87–137.
- CONLIN, M., T. O’DONOGHUE, AND T. VOGELSANG (2007): “Projection bias in catalog orders,” *American Economic Review*, 97, 1217–1249.
- DELLAVIGNA, S. (2018): “Structural behavioral economics,” in *Handbook of Behavioral Economics: Applications and Foundations 1*, Elsevier, vol. 1, 613–723.
- ELIAZ, K. AND R. SPIEGLER (2015): “Beyond “Ellison’s matrix”: New directions in behavioral industrial organization,” *Review of Industrial Organization*, 47, 259–272.
- FISHER, G. AND A. RANGEL (2014): “Symmetry in cold-to-hot and hot-to-cold valuation gaps,” *Psychological science*, 25, 120–127.

- GRUBB, M. D. (2015): “Overconfident consumers in the marketplace,” *Journal of Economic Perspectives*, 29, 9–36.
- GRUBER, J. AND B. KÖSZEGI (2001): “Is addiction “rational”? Theory and evidence,” *The Quarterly Journal of Economics*, 116, 1261–1303.
- HEIDHUES, P. AND B. KÖSZEGI (2010): “Exploiting naivete about self-control in the credit market,” *American Economic Review*, 100, 2279–2303.
- HEIDHUES, P. AND B. KOSZEGI (2018): “Behavioral Industrial Organization,” *Handbook of Behavioral Economics*, Elsevier.
- HEIDHUES, P., B. KÖSZEGI, AND T. MUROOKA (2016): “Inferior products and profitable deception,” *The Review of Economic Studies*, 84, 323–356.
- HEIDHUES, P. AND P. STRACK (2019): “Identifying Present-Bias from the Timing of Choices,” *Available at SSRN*.
- HERWEG, F., D. MÜLLER, AND P. WEINSCHENK (2010): “Binary payment schemes: Moral hazard and loss aversion,” *American Economic Review*, 100, 2451–77.
- HINNOSAAR, M. (2016): “Time inconsistency and alcohol sales restrictions,” *European Economic Review*, 87, 108–131.
- INDERST, R. AND M. OTTAVIANI (2013): “Sales Talk, Cancellation Terms and the Role of Consumer Protection,” *The Review of Economic Studies*, 80, 1002–1026.
- LAIBSON, D., A. REPETTO, AND J. TOBACMAN (2008): “Estimating Discount Functions with Consumption Choices over the Lifecycle,” *mimeo*.
- LEE, J. (2013): “The impact of a mandatory cooling-off period on divorce,” *The Journal of Law and Economics*, 56, 227–243.
- LOEWENSTEIN, G., T. O’DONOGHUE, AND M. RABIN (2003): “Projection Bias In Predicting Future Utility,” *Quarterly Journal of Economics*, 118, 1209–1248.
- MAHAJAN, A., C. MICHEL, AND A. TAROZZI (2020): “Identification of time-inconsistent models: The case of insecticide treated nets,” *NBER Working Paper (No. w27198)*.
- MICHEL, C. (2017): “Market regulation of voluntary add-on contracts,” *International Journal of Industrial Organization*, 54, 239–268.
- O’DONOGHUE, T. AND M. RABIN (2006): “Optimal sin taxes,” *Journal of Public Economics*, 90, 1825–1849.

ROWE, G. S. (2007): “Legal Guide to Door-to-Door Criminals,” *Office of Attorney General Consumer Protection Division Report*, <http://http://maine.gov/ag/consumer/doortodoorguide.pdf/>.

SHUI, H. AND L. AUSUBEL (2004): “Time inconsistency in the credit card market,” *mimeo*.

SPIEGLER, R. (2015): “On the equilibrium effects of nudging,” *The Journal of Legal Studies*, 44, 389–416.

A Model Analysis

A.1 Baseline Model

Absent any policy, consumers make their purchase decision based on their first-period motivation. A fraction μ of consumers is motivated and predicts their expected utility next period to be \tilde{u} , while a fraction $1 - \mu$ is unmotivated with prediction \underline{u} . The resulting demand as a function of price is easily derived:

$$D(p) = \begin{cases} 1 & \text{if } p \leq \underline{u} = \underline{u} + \mu(1 - \alpha)\Delta \\ \mu & \text{if } p \in (\underline{u}, \tilde{u}] = (\underline{u} + \mu(1 - \alpha)\Delta, \underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta] \\ 0 & \text{if } p > \tilde{u} = \underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta. \end{cases} \quad (\text{A.1})$$

As such, the firm has two possible pricing strategies. It can target first-period-motivated consumers only, or it can price such that all consumers purchase. In both cases, the firm extracts consumers’ predicted expected utility in full.

Targeting only motivated consumers allows the firm to fully extract the motivated consumers’ predicted consumption utility under projection bias. However, this comes at the cost of only serving a fraction μ of the potential market. By contrast, targeting all consumers captures a larger market share, but requires pricing below the actual expected utility. We show that the firm prefers to target motivated consumers if and only if the marginal cost of production is sufficiently high, which is captured by [Proposition 1](#).

Proof of [Proposition 1](#) Profits are denoted by $\underline{\pi}$ when targeting all consumers, and $\tilde{\pi}$ when targeting motivated consumers only, respectively. We obtain

$$\begin{aligned} \underline{\pi} &= \underline{u} - c = \underline{u} + \mu(1 - \alpha)\Delta - c \\ \tilde{\pi} &= \mu(\tilde{u} - c) = \mu(\underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta - c), \end{aligned} \quad (\text{A.2})$$

where $c < \delta\tilde{u} \leq \tilde{u}$ ensures that $\tilde{\pi}$ is weakly positive. Given the tie-breaking assumption, the

firm prefers to target all consumers iff

$$\begin{aligned} \mu \cdot (\underline{u} + (\mu + (1 - \mu)\alpha)\Delta - c) &\geq \underline{u} + (\mu - \alpha\mu)\Delta - c \\ \Leftrightarrow c &\geq \underline{u} + \frac{\mu}{1 - \mu}\Delta(1 - 2\alpha - (1 - \alpha)\mu) \equiv \tilde{c}. \end{aligned} \quad (\text{A.3})$$

Note that

$$\begin{aligned} E[u] \geq \tilde{c} &\Leftrightarrow \underline{u} + \mu\Delta \geq \underline{u} + \frac{\mu}{1 - \mu}\Delta(1 - 2\alpha - (1 - \alpha)\mu) \\ &\Leftrightarrow -(2 - \mu)\alpha \leq 0 \end{aligned} \quad (\text{A.4})$$

which is generically true due to $\mu \in (0, 1)$ and $\alpha \in [0, 1]$. For comparative statics,

$$(i) \quad \partial\tilde{c}/\partial\alpha = -(2 - \mu)\frac{\mu}{1 - \mu}\Delta < 0$$

$$(ii) \quad \partial\tilde{c}/\partial\underline{u} = 1 > 0$$

$$(iii) \quad \partial\tilde{c}/\partial\Delta = \underbrace{\frac{\mu}{1 - \mu}}_{>0} \cdot \underbrace{(1 - 2\alpha - (1 - \alpha)\mu)}_{\geq 0}$$

$$(iv) \quad \partial\tilde{c}/\partial\mu = \underbrace{\frac{1}{(1 - \mu)^2}}_{>0} \Delta \underbrace{[1 - 2\alpha - (1 - \alpha)\mu(2 - \mu)]}_{\geq 0}$$

Finally, consumer welfare is simply the difference between price and actual expected utility, weighted by the mass of purchasing consumers. Thus,

$$\begin{aligned} \widetilde{CW} &= \mu \cdot (E[u] - \tilde{u}) = -\mu(1 - \mu)\alpha\Delta \leq 0 \\ \underline{CW} &= E[u] - \underline{u} = \mu\alpha\Delta \geq 0, \end{aligned} \quad (\text{A.5})$$

with total welfare obtained by summing consumer welfare and the firm's profits. ■

As the markup $p - c$ that the firm charges is larger for the exclusive targeting strategy, an increase in marginal cost has a lower relative impact on profits. Thus, for sufficiently high marginal cost c , the firm prefers to target motivated consumers only. The firm's targeting decision has direct implications for welfare. As motivated consumers overpredict their expected utility relative to the actual expected utility due to the projection bias, while unmotivated consumers underpredict it, consumer welfare is weakly (strictly for $\alpha > 0$) negative in case only motivated consumers are targeted, and weakly (strictly) positive if the full market is covered. If c lies above the actual expected utility $E[u]$, but below the predicted expected utility of motivated consumers \tilde{u} , the firm is able to profitably produce and sell the product even though the marginal cost exceeds the expected utility of consumption, rendering total welfare negative. As exclusive targeting always leads to negative consumer welfare, a policy introduction may be necessary to prevent consumers from having negative expected consumption utility when

$\alpha > 0$. Moreover, a policy introduction is necessarily desirable irrespective of the chosen welfare standard whenever total welfare is negative.

The threshold \tilde{c} is increasing in the utility of unmotivated consumers \underline{u} as a higher \underline{u} implies a larger loss of foregone extracted rent due to exclusive targeting. Similarly, a higher projection bias α decreases the threshold \tilde{c} as it decreases the price which can be charged when targeting all consumers, while it increases the price chargeable when targeting motivated consumers only.¹³ Finally, note that it is possible that the firm will prefer exclusive targeting irrespective of its marginal cost of production c . This materializes if the projection bias α is sufficiently large, the fraction of motivated consumers sufficiently high, and the utility difference $\Delta = \bar{u} - \underline{u}$ large relative to \underline{u} . $\tilde{c} \xrightarrow{\alpha \rightarrow 1} \underline{u} - \frac{\mu}{1-\mu}\Delta$, which is negative in this case.

A.2 Cooling off

Both the initial decision to register for purchase in the first period, and the confirmation decision in the second period are based on predicted expected utilities. As in the baseline model, the projection bias thus leaves scope for a firm to extract more than consumers' expected utility. The registering and confirmation decision are straightforward. A consumer registers if and only if the predicted consumption utility given her first-period motivation weakly exceeds the price, and confirms a purchase if and only if the same holds given her second-period motivation. Thus, we obtain for the mass of consumers who actually purchase the good, $D_c(p)$, and for the mass of consumers which step back after their initial letter of intent, $S(p)$:

$$D_c(p) = \begin{cases} 1 & \text{if } p \leq \delta \underline{u} = \delta(\underline{u} + \mu(1 - \alpha)\Delta) \\ \mu^2 & \text{if } p \in (\delta \underline{u}, \delta \tilde{u}] = (\delta(\underline{u} + \mu(1 - \alpha)\Delta), \delta(\underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta)] \\ 0 & \text{if } p > \delta \tilde{u} = \delta(\underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta) \end{cases} \quad (\text{A.6})$$

$$S(p) = \begin{cases} 0 & \text{if } p \leq \delta \underline{u} = \delta(\underline{u} + \mu(1 - \alpha)\Delta) \\ \mu(1 - \mu) & \text{if } p \in (\delta \underline{u}, \delta \tilde{u}] = (\delta(\underline{u} + \mu(1 - \alpha)\Delta), \delta(\underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta)] \\ 0 & \text{if } p > \delta \tilde{u} = \delta(\underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta), \end{cases} \quad (\text{A.7})$$

¹³The effect of changes of the utility difference Δ , as well as the fraction of motivated consumers μ , by contrast, is ambiguous. This is because increases in both Δ and μ increase the prices which can be charged in both strategies. Note that this increase is always weakly larger for the exclusive targeting strategy which can be seen from the sensitivities of the prices with exclusive targeting (\tilde{p}) and full coverage (p) to changes in μ and Δ :

$$\begin{aligned} \frac{\partial \tilde{p}}{\partial \mu} &= (1 - \alpha)\Delta &= \frac{\partial p}{\partial \mu} \\ \frac{\partial \tilde{p}}{\partial \Delta} &= \alpha + (1 - \alpha)\mu &\geq (1 - \alpha)\mu = \frac{\partial p}{\partial \Delta}. \end{aligned}$$

Nonetheless, this larger increase has to be traded off with the lower market share. Which effect dominates depends on the bias α , as well as the current values of μ and Δ . If consumers are unbiased, i.e. for $\alpha = 0$, the firm will never target only motivated consumers. This is because both consumer types predict to have the same expected utility $E[u] = \underline{u} + \mu\Delta$ in period 2. The firm thus chooses to either serve the full market in case it is profitable, i.e. for $c \leq E[u]$, or else abstain from production. The assumption $c < \delta \tilde{u}$ in this case reduces to $c < E[u]$.

where $D_c(p) + S(p)$ gives the mass of consumers who initially sign a letter of intent.

It follows that the firm has two potential targeting strategies. It may either induce only first-period-motivated consumers to register, which implies that they only confirm if they are also second-period-motivated. Alternatively, the firm can initially target all consumers, which implies that all consumers subsequently confirm their purchase. The candidate prices are thus the same as in the baseline case, but weighted with the discount factor δ as consumption is delayed by one period. Similar to the baseline case, there is a unique threshold \tilde{c}_c such that the firm prefers the exclusive targeting if and only if the marginal cost of production exceeds \tilde{c}_c . This leads to [Proposition 2](#).

Proof of Proposition 2 Given the candidate prices, the firm's profits are straightforward and given by

$$\begin{aligned}\tilde{\pi}_c &= \mu^2 \cdot \delta(\tilde{u} - c) = \delta\mu^2(\underline{u} + \alpha\Delta + (1 - \alpha)\mu\Delta - c) \\ \underline{\pi}_c &= \delta(\underline{u} - c) = \delta(\underline{u} + (1 - \alpha)\mu\Delta - c).\end{aligned}\tag{A.8}$$

As such, the firm prefers exclusive targeting if and only if

$$\begin{aligned}\tilde{\pi}_c &\geq \underline{\pi}_c \\ \Leftrightarrow \delta\mu^2 \cdot (\underline{u} + (\mu + (1 - \mu)\alpha)\Delta - c) &\geq \delta(\underline{u} + (\mu - \alpha\mu)\Delta - c) \\ \Leftrightarrow c &\geq \underline{u} + \frac{\mu}{(1 - \mu)}\Delta \frac{(1 - (1 + \mu)\alpha - (1 - \alpha)\mu^2)}{1 + \mu} \equiv \tilde{c}_c\end{aligned}\tag{A.9}$$

where

$$\begin{aligned}\tilde{c}_c &> \tilde{c} \\ \Leftrightarrow \frac{(1 - (1 + \mu)\alpha - (1 - \alpha)\mu^2)}{1 + \mu} &> (1 - 2\alpha - (1 - \alpha)\mu) \\ \Leftrightarrow \alpha + (1 + 2\alpha)\mu &> 0,\end{aligned}\tag{A.10}$$

which holds generically. Similarly, we have that

$$\begin{aligned}\tilde{c}_c \leq E[u] &\Leftrightarrow \underline{u} + \frac{\mu}{(1 - \mu)}\Delta \frac{(1 - (1 + \mu)\alpha - (1 - \alpha)\mu^2)}{1 + \mu} \leq \underline{u} + \mu\Delta \\ &\Leftrightarrow -(1 + \mu(1 - \mu))\alpha \leq 0\end{aligned}\tag{A.11}$$

Regarding the comparative statics of \tilde{c} , we obtain

$$(i) \quad \partial\tilde{c}_c/\partial\alpha = \frac{\mu}{1-\mu^2}\Delta(-\mu - (1 - \mu^2)) < 0$$

$$(ii) \quad \partial\tilde{c}_c/\partial\underline{u} = 1 > 0$$

$$(iii) \quad \partial \tilde{c}_c / \partial \Delta = \frac{\mu}{1-\mu^2} (1 - (1 + \mu)\alpha - (1 - \alpha)\mu^2)$$

$$(iv) \quad \partial \tilde{c}_c / \partial \mu = -\frac{\Delta}{(1-\mu^2)^2} \cdot [\mu(1 - \mu^2)(\alpha + (2 - \alpha)\mu) - (1 + \mu^2)((1 - \alpha)(1 - \mu^2) - \alpha\mu)].$$

Consumer welfare immediately follows by comparing the price with the actual expected consumption utility $\delta E[u]$, similar to welfare in the baseline case, and total welfare is obtained by summing consumer welfare and the firm's profits. ■

\tilde{c}_c lies above that in the baseline case, $\tilde{c}_c \geq \tilde{c}$ ($\tilde{c}_c > \tilde{c}$ for $\alpha > 0$), as only twice-motivated consumers confirm their purchase decision with exclusive targeting, which decreases the attractiveness of this strategy relative to the baseline case—the market share is even lower. The comparative statics of \tilde{c}_c are similar to those of the baseline threshold \tilde{c} : \tilde{c}_c increases in \underline{u} , decreases in α , and is ambiguously affected by changes in Δ and μ . Similarly, the firm may prefer to generically target only motivated consumers, and total welfare becomes negative for $c \in (E[u], \delta \tilde{u}]$.

Impact of a Cooling Off Policy Table 1 summarizes the firm's targeting behavior if a mandatory cooling off period is instituted, and contrasts it with the baseline case when $\alpha > 0$.

Cost Range	$c \leq \tilde{c}$	$c \in (\tilde{c}, \tilde{c}_c]$	$c > \tilde{c}_c$
Δ Price	≈ 0 (by δ)	$\delta \underline{u} - \tilde{u} < 0$	≈ 0 (by δ)
Δ Market Coverage	-	$1 - \mu > 0$	$\mu^2 - \mu < 0$
Δ Profit	≈ 0 (by δ)	< 0 otherwise low target in baseline	< 0 as foregone market share
Δ Consumer Welfare	≈ 0 (by δ)	> 0 as no longer $CW < 0$	> 0 as increase in CW (but < 0)
Δ Total Welfare	≈ 0 (by δ)	$(\delta - \mu)(\underline{u} + \mu\Delta - c), > 0$ for $\delta \rightarrow 1$	$-\mu(1 - \delta\mu)(E[u] - c), < 0 \iff E[u] > c$

Table 1: Comparison of unregulated outcome and mandatory cooling off period ($\alpha > 0$)

If only motivated consumers are targeted without the policy, the firm fully extracts the motivated consumers' predicted utility under projection bias which leads to negative consumer welfare. Introducing a cooling off period in this case may either lead to a decrease in such practices by lowering the number of consumers with a negative expected consumption utility and delaying consumption ($c \geq \tilde{c}_c > \tilde{c}$), or alter the firm's targeting decision to serve the full market ($c \in (\tilde{c}, \tilde{c}_c]$). In both cases, consumer welfare increases. By contrast, the firm always covers the full market after a cooling off period has been introduced if it already does so without the policy, $c \leq \tilde{c} \leq \tilde{c}_c$. In that case, the cooling off period only unnecessarily delays consumption, which reduces consumer welfare. With regards to total welfare, the policy is beneficial if a switch in the targeting behavior is induced ($c \in (\tilde{c}, \tilde{c}_c]$), and detrimental if the market was already fully covered pre policy introduction ($c \leq \tilde{c}$). If exclusive targeting persists post policy introduction, the total welfare effect depends on the relation of the marginal cost c and the actual expected utility $E[u]$. If $c \leq E[u]$, the lowered market coverage due to the policy's implementation is actually detrimental to total welfare (albeit beneficial for consumer welfare), whereas the reverse holds if $c > E[u]$.

A.3 Return policy

When a return policy is in place, consumers have the option to return the product instead of consuming it, in which case they are refunded the cost. Crucially, this means that the return decision is based on their actual utility from consumption as it takes place after the second-period motivation has been realized. However, the initial purchase decision is still driven by the conditional predicted utilities which experience the projection bias.

As the return decision is based on the actual utility of consumption, return behavior is straightforward. In the absence of return costs, the consumer keeps the good if and only if her actual consumption utility is at least as high as the price, i.e. if and only if $p \leq \bar{u}$ for motivated consumers, or $p \leq \underline{u}$ for unmotivated consumers.

By contrast, the initial purchase decision depends on predicted utilities. A consumer initially purchases if and only if the price the firm charges is below the predicted utility in at least one motivational state. As the predicted conditional utility in the motivated state is weakly higher than in the unmotivated state, $\tilde{u}(\bar{s}|s) \geq \tilde{u}(\underline{s}|s)$, a consumer thus purchases if and only if the price is below the predicted utility for the motivated state. Note that these predicted utilities incorporate the projection bias; the predicted return behavior can differ from the actual one.

Formally, we obtain for the mass of consumers who actually purchase and keep the good, $D_r(p)$, and for the mass of consumers which return the good $R(p)$:

$$D_r(p) = \begin{cases} 1 & \text{if } p \leq \underline{u} \\ \mu & \text{if } p \in (\underline{u}, \tilde{u}(\bar{s}|s)] = (\underline{u}, \alpha\underline{u} + (1-\alpha)\bar{u}] \\ \mu^2 & \text{if } p \in (\tilde{u}(\bar{s}|s), \bar{u}] = (\alpha\underline{u} + (1-\alpha)\bar{u}, \bar{u}] \\ 0 & \text{if } p > \bar{u} \end{cases} \quad (\text{A.12})$$

$$R(p) = \begin{cases} 0 & \text{if } p \leq \underline{u} \\ (1-\mu) & \text{if } p \in (\underline{u}, \tilde{u}(\bar{s}|s)] = (\underline{u}, \alpha\underline{u} + (1-\alpha)\bar{u}] \\ \mu \cdot (1-\mu) & \text{if } p \in (\tilde{u}(\bar{s}|s), \bar{u}] = (\alpha\underline{u} + (1-\alpha)\bar{u}, \bar{u}] \\ 0 & \text{if } p > \bar{u}, \end{cases} \quad (\text{A.13})$$

where $D_r(p) + R(p)$ gives the mass of consumers who initially purchase in the first period.

Having characterized the initial purchase and return decision, we can turn to the firm's pricing. It again depends on the marginal cost of production, but involves two cutoffs \underline{c}_r and \tilde{c}_r . If the marginal cost of production is sufficiently low, the firm prefers to serve the full market. This requires even unmotivated consumers to not return the good, and thus imposes a maximal price of $\underline{p}_r = \underline{u}$. For sufficiently high marginal cost, the firm instead targets only initially motivated consumers, who return the good unless they are again motivated in the second period. The firm fully extracts the consumers' rent by charging $\bar{p}_r = \bar{u}$. Finally, the firm may use a third strategy for intermediate marginal cost. It initially sets the price such that even unmotivated consumers purchase the good, but induces return unless consumers are

second-period-motivated. This is facilitated by charging $\underline{p}_r = \tilde{u}(\bar{s}|\underline{s}) \equiv \underline{p}_r$. For this strategy to be viable in equilibrium, we require $\alpha < \frac{1}{1+\mu}$. As the bias is not too large, the gain in profits from charging a higher price $\underline{p}_r > \underline{p}_r$ is enough to offset the lost sales from returned goods, while the fact that the fraction of motivated consumers is not too high keeps this strategy preferred to fully exclusive targeting. This is captured by [Proposition 3](#).

Proof of Proposition 3 Profits for the three candidates are characterized by

$$\begin{aligned}\bar{\pi}_r &= \mu^2 \cdot (\bar{u} - c) = \mu^2 \cdot (\underline{u} + \Delta - c) \\ \underline{\pi}_r &= \mu \cdot [\underline{u} + (1 - \alpha)\Delta - c] \\ \underline{\pi}_r &= \underline{u} - c,\end{aligned}\tag{A.14}$$

where e.g. $\bar{\pi}_r$ is derived by initially targeting only the fraction μ of motivated consumers, who return the good unless they are also second-period-motivated and the firm charges $\bar{u} = \underline{u} + \Delta$. Pairwise comparison of the profits associated with the targeting strategies yields

$$\begin{aligned}\bar{\pi}_r \geq \underline{\pi}_r &\iff \mu\underline{u} + \mu\Delta - \mu c \geq \underline{u} + (1 - \alpha)\Delta - c \\ &\iff c \geq \underline{u} + \Delta \frac{1 - \alpha - \mu}{1 - \mu} \equiv \tilde{c}_r^1\end{aligned}\tag{A.15}$$

$$\begin{aligned}\bar{\pi}_r \geq \underline{\pi}_r &\iff \mu^2\underline{u} + \mu^2\Delta - \mu^2 c \geq \underline{u} - c \\ &\iff c \geq \underline{u} - \frac{\mu^2}{(1 - \mu)(1 + \mu)}\Delta \equiv \tilde{c}_r^2\end{aligned}\tag{A.16}$$

$$\begin{aligned}\underline{\pi}_r \geq \underline{\pi}_r &\iff \mu\underline{u} + \mu(1 - \alpha)\Delta - \mu c \geq \underline{u} - c \\ &\iff c \geq \underline{u} - \frac{\mu(1 - \alpha)}{(1 - \mu)}\Delta \equiv \tilde{c}_r^3.\end{aligned}\tag{A.17}$$

Note that $\tilde{c}_r^3 \leq \underline{u}$ as $\frac{\mu}{1-\mu} > 0$, $1 - \alpha \geq 0$ and $\Delta > 0$. Regarding the ordering of thresholds, note that

$$\begin{aligned}\tilde{c}_r^1 > \tilde{c}_r^2 &\iff 1 - \mu - \alpha > -\frac{\mu^2}{1 + \mu} \\ &\iff \alpha < \frac{1}{1 + \mu} \\ &\iff \tilde{c}_r^2 > \tilde{c}_r^3.\end{aligned}\tag{A.18}$$

We hence have two cases: If $\alpha < \frac{1}{1+\mu}$, $\tilde{c}_r^1 > \tilde{c}_r^2 > \tilde{c}_r^3$, while for $\alpha \geq \frac{1}{1+\mu}$ we have $\tilde{c}_r^1 \leq \tilde{c}_r^2 \leq \tilde{c}_r^3$. Defining $\tilde{c}_r \equiv \max\{\tilde{c}_r^1, \tilde{c}_r^2\}$ and $\underline{c}_r = \min\{\tilde{c}_r^2, \tilde{c}_r^3\}$, this yields the targeting behavior contained in the proposition.

To see that three different cases may arise, consider the following. If $\tilde{c}_r^1 \leq \tilde{c}_r^2 \leq \tilde{c}_r^3 \iff 1 - (1 + \mu)\alpha < 0$, there is only a unique relevant threshold $\underline{c}_r = \tilde{c}_r = \tilde{c}_r^2$: The firm either targets all consumers without inducing return, $c \leq \tilde{c}_r^2$, or extracts the full rent from twice-motivated consumers, $c > \tilde{c}_r^2$. By contrast, for $\tilde{c}_r^1 > \tilde{c}_r^2 > \tilde{c}_r^3$, we have a second threshold $\underline{c}_r = \tilde{c}_r^3$. But for

the comparison between \tilde{c}_r^3 and \tilde{c} , note that

$$\tilde{c} \geq \tilde{c}_r^3 \Leftrightarrow 1 \leq (3 - \mu)(1 - \alpha) \Leftrightarrow \alpha \leq \frac{2 - \mu}{3 - \mu}, \quad (\text{A.19})$$

which can in principle be both satisfied for low values of α , or not satisfied for high values of α . Moreover, this is consistent with $\alpha < \frac{1}{1+\mu}$ which is a requirement for this case to be relevant—in fact, given $\mu \in (0, 1)$, we have $\frac{2-\mu}{3-\mu} < \frac{1}{1+\mu}$ and hence that both orders are always possible.

Consumer and total welfare follow immediately from the above derivations. Finally, recall that $\tilde{c} = \underline{u} + \frac{\mu}{1-\mu}\Delta(1 - 2\alpha - (1 - \alpha)\mu)$, and that $\tilde{c}_r = \max\{\tilde{c}_r^1, \tilde{c}_r^2\}$. It hence suffices to show that

$$\begin{aligned} \tilde{c}_r^1 \geq \tilde{c} &\iff 1 - \mu - \alpha \geq \mu(1 - 2\alpha - (1 - \alpha)\mu) \\ &\iff (1 - 2\mu + \mu^2)(1 - \alpha) \geq 0, \end{aligned}$$

which holds generically as $1 - 2\mu + \mu^2$ is strictly positive for $\mu \in (0, 1)$. ■

The threshold for exclusive targeting, \tilde{c}_r , lies above that in the baseline, \tilde{c} . The range for targeting all consumers initially, but inducing on-path return of the good by consumers unmotivated in the second period, $(\underline{c}_r, \tilde{c}_r]$, may either embed \tilde{c} , lie strictly above \tilde{c} , or be empty (if $\alpha \geq \frac{1}{1+\mu}$).

Impact of a return policy A return policy ensures that a firm cannot make full use of an upward projection bias—the return decision is based on actual consumption utilities and the bias only matters for the initial purchase decision. For the same reason, total welfare is necessarily (weakly) positive. Moreover, the projection bias is actually beneficial for consumer welfare in case the firm targets all consumers initially, but induces on-path return. Pricing in this case is based on the predicted utility in the motivated state for an unmotivated consumer, a higher bias lowers what the firm can charge and thus increases the rent left to second-period-motivated consumers. However, as this decreases the firm’s profits, it simultaneously makes it less likely that this targeting behavior materializes in equilibrium; if the bias is large enough ($\alpha > \frac{1}{1+\mu}$), this strategy never features in equilibrium. If consumers have negative expected consumption utility in the baseline case because the firm targets only motivated consumers, the introduction of a return policy always increases consumer welfare as it either avoids negative expected consumer utility or changes the firm’s targeting decision in which case consumer welfare post policy introduction is positive. However, if the firm targets all consumers before the introduction, a return policy can be consumer welfare neutral if the firm changes its pricing to induce on-path return and only increases consumer welfare if even second-period-unmotivated consumers keep the good, as the associated price decrease to \underline{p}_r leaves a significant rent to second-period-motivated consumers. The effect of a return policy on total welfare is more ambiguous. Total welfare is unchanged if the full market was covered both pre and post policy, and increases if the return policy leads to a change from exclusive targeting, to initial full coverage with induced return. This is because a fraction μ of consumers consumes

the good in both scenarios, but with a return policy, all of these are motivated. In all other cases, the total welfare impact can not be generically signed. Consider for example the case where $c > \tilde{c}_r > \tilde{c}$. A return policy leads to only twice-motivated consumers (mass μ^2) actually consuming the good, while in the baseline all initially motivated consumers (mass μ) consumed it in the second period irrespective of their period-two-motivation. Whether this is beneficial for total welfare or not depends on the relation of the marginal cost c with the consumption utility of unmotivated consumers \underline{u} , i.e. whether unmotivated consumers should or should not consume the good. The relation of \tilde{c}_r , \tilde{c} , and \underline{u} , however, depends on the fundamentals and cannot be generically determined. We summarize these considerations in Table 2.

Cost Range	(1) $c \leq \tilde{c}, c \leq \underline{c}_r$	(2) $c \leq \tilde{c}, c \in (\underline{c}_r, \tilde{c}_r]$	(3) $c > \tilde{c}, c \leq \underline{c}_r$
Δ Price	$\underline{u} - \underline{u} < 0$	$\tilde{u}(\bar{s} \underline{s}) - \underline{u} > 0$	$\underline{u} - \tilde{u} < 0$
Δ Market Coverage	-	$\mu - 1 < 0; (1 - \mu)$ return	$1 - \mu > 0$
Δ Profit	$-\mu(1 - \alpha)\Delta < 0$	$(1 - \mu)(c - \underline{u}) - \mu\alpha\Delta$	$(1 - \mu)(u - c) - (\alpha + (1 - \alpha)\mu)\Delta$
Δ Consumer Welfare	$(1 - \alpha)\mu\Delta > 0$	-	> 0 (from neg to pos)
Δ Total Welfare	-	$(1 - \mu)(c - \underline{u})$	$(1 - \mu)(E[u] - c) > 0$ as $\underline{u} > c$ for profitability
Cost Range	(4) $c > \tilde{c}, c \in (\underline{c}_r, \tilde{c}_r]$	(5) $c > \tilde{c}_r > \tilde{c}$	
Δ Price	$\tilde{u}(\bar{s} \underline{s}) - \tilde{u}$, ambiguous	$\tilde{u} - \tilde{u} > 0$	
Δ Market Coverage	-; $(1 - \mu)$ return	$\mu(\mu - 1); \mu(1 - \mu)$ return	
Δ Profit	$> 0 \iff 1 - 2\alpha - (1 - \alpha)\mu > 0$	$-\mu(1 - \mu)(u + \alpha\Delta - c)$	
Δ Consumer Welfare	> 0 (from neg to pos)	> 0 (from neg to 0)	
Δ Total Welfare	$\mu(1 - \mu)\Delta > 0$	$\mu(1 - \mu)(c - \underline{u})$	

Table 2: Comparison of unregulated outcome and return policy

A.4 Proof: Comparison of Policies

Proof. Regarding consumer welfare, it is sufficient to establish $\tilde{c}_c > \tilde{c}_r$ as then for $c \in (\tilde{c}_r, \tilde{c}_c)$ consumer welfare is larger under a cooling off period (where it is strictly positive) than under a return policy (where it is zero).

Recall that $\tilde{c}_r = \max\{\tilde{c}_r^1, \tilde{c}_r^2\}$, with the characterizations of \tilde{c}_c , \tilde{c}_r^1 and \tilde{c}_r^2 given as follows.

$$(i) \quad \tilde{c}_c = \underline{u} + \frac{\mu}{(1-\mu)}\Delta \frac{(1-(1+\mu)\alpha-(1-\alpha)\mu^2)}{1+\mu} = \underline{u} + \Delta \frac{\mu(1-\alpha-\mu[\alpha+(1-\alpha)\mu])}{1-\mu^2}$$

$$(ii) \quad \tilde{c}_r^2 = \underline{u} - \Delta \frac{\mu^2}{1-\mu^2}$$

$$(iii) \quad \tilde{c}_r^1 = \underline{u} + \Delta \frac{1-\mu-\alpha}{1-\mu}.$$

$\tilde{c}_c > \tilde{c}_r^1$ immediately follows as $\alpha + (1 - \alpha)\mu < 1$ due to $\mu \in (0, 1)$. Similarly, we obtain

$$\begin{aligned} \tilde{c}_c > \tilde{c}_r^2 &\Leftrightarrow \frac{\mu(1-\alpha-\mu[\alpha+(1-\alpha)\mu])}{1-\mu^2} + \frac{\mu^2}{1-\mu^2} > 0 \\ &\Leftrightarrow \mu(1-\alpha-\mu[\alpha+(1-\alpha)\mu]) + \mu^2 > 0 \\ &\Leftrightarrow \underbrace{(1-\alpha)}_{\geq 0} - \underbrace{[\alpha+(1-\alpha)\mu]}_{< 1} \mu \geq -\mu, \end{aligned} \tag{A.20}$$

which holds generically. So we always have $\tilde{c}_c > \max\{\tilde{c}_r^2, \tilde{c}_r^1\}$. There hence exists a range of costs such that exclusive targeting would prevail under a return policy, leading to zero consumer

welfare, while the full market is served under a cooling off period, leading to positive consumer welfare.

The implications for total welfare are not as straightforward due to the difference in the timing of consumption. However, if we consider the case $c < \tilde{c}_c$ and $c > \tilde{c}_r$, we have $TW_r = \mu^2(\bar{u} - c) > 0$ and $TW_c = \delta(E[u] - c) > 0$. Either policy may do better, as the increased utility from having only motivated consumers consume the good (return policy) is traded off with the increased number of consumers (cooling off). Crucially, a cooling off period is sure to be better for total welfare if $c < \underline{u}$ (provided that δ is not too small). ■

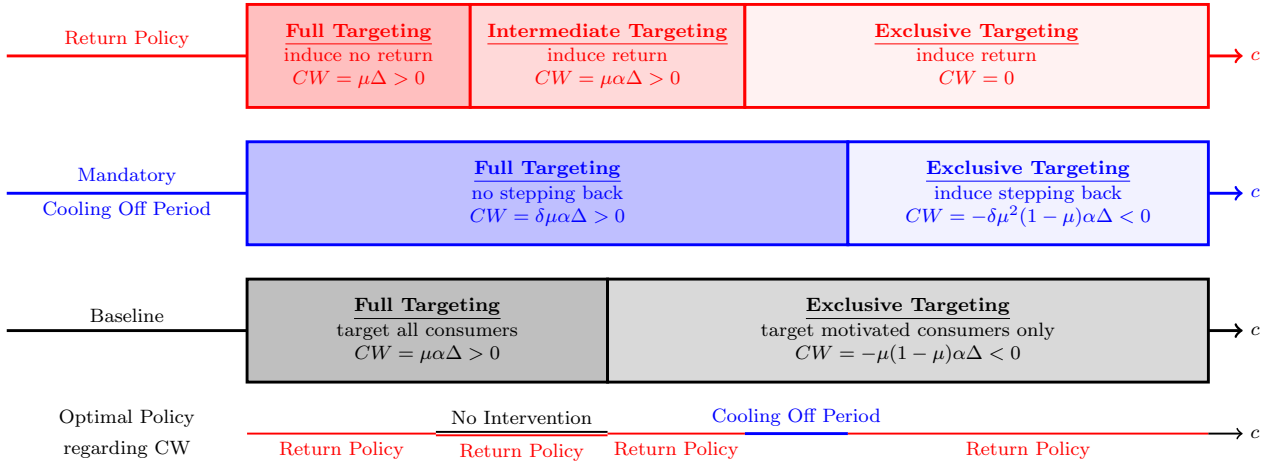


Figure 6: Relevant cost ranges across all policy regimes & CW-maximizing policy

Figure 6 summarizes the results regarding the firm’s pricing as a function of the firm’s marginal cost c by contrasting the three cases—baseline, mandatory cooling off period, and return policy. It also shows the optimal policy from a consumer welfare perspective. The figure implicitly assumes $\tilde{c} \in (\underline{c}_r, \tilde{c}_r) \neq \emptyset$, as well as $\tilde{c} > \underline{c}_r > 0$. Notably, there is a nonempty range of marginal cost such that consumer welfare is negative absent a policy, zero if a return policy is instituted, but positive if a mandatory cooling off period is in place.¹⁴ In all other ranges, a return policy maximizes consumer welfare, albeit under the assumption that return is socially and privately costless. Note that the unregulated market also achieves the maximal consumer welfare for low to intermediate c . In those cases, the absence of a policy intervention would arguably be preferred by a social planner; in fact, even vanishingly small return costs would ensure this.

A.5 Voluntary return policy

A key question is whether a policymaker even needs to consider imposing a mandatory return policy. Costless return is offered by a variety of firms in practice, most notably by amazon.com,

¹⁴Note that this holds for all cases and not just the one depicted in the figure as long as $\tilde{c}_c > 0$.

and it could in principle be the case that there is no scope for a policy because any time consumer and/or total welfare would be increased due to the introduction of a return policy, the firm would already find it beneficial to adopt it voluntarily. However, it turns out that this is not the case.

To establish this, consider the case where $\alpha < \frac{1}{1+\mu}$ such that $(\underline{c}_r, \tilde{c}_r]$ is nonempty and consider c such that $c \leq \tilde{c}$, $c \in (\underline{c}_r, \tilde{c}_r]$. In this case, the adoption of a return policy leads to a change in consumer welfare and total welfare respectively of

$$dCW = 0, dTW = (1 - \mu)(c - \underline{u}). \quad (\text{A.21})$$

However, for the firm to voluntarily adopt a return policy, we would require $d\pi = (1 - \mu)(c - \underline{u}) - \mu\alpha\Delta > 0 \iff c > \underline{u} + \frac{\mu\alpha}{1-\mu}\Delta$. This implies two things. First, if the RP were welfare-destroying (for $c < \underline{u}$), the firm would not offer it voluntarily. However, there is a potentially nonempty range of costs $\underline{u} < c < \underline{u} + \frac{\mu\alpha}{1-\mu}\Delta$ such that we are in case (2) and a RP would be beneficial for TW, but would not be voluntarily offered by the firm as it is not beneficial for the firm's profits. It is only potentially nonempty when there are parametrizations of μ and α such that the intersection of case (2) and $\underline{u} < c < \underline{u} + \frac{\mu\alpha}{1-\mu}\Delta$ is empty. It is not generically empty, however, as $\tilde{c} > \underline{u}$, and $\tilde{c}_r > \underline{u}$ are easily obtained for small α .

Similarly, considering $c > \tilde{c}_r > \tilde{c}$, we have $dCW > 0$ (as consumer welfare is negative due in the baseline), and $dTW = \mu(1 - \mu)(c - \underline{u})$. For voluntary adoption of a return policy, we would require $d\pi = -\mu(1 - \mu)(\underline{u} + \alpha\Delta - c) > 0$, i.e. $c > \underline{u} + \alpha\Delta$. Again, this implies two things. First, if the RP were welfare-destroying ($c < \underline{u}$), it would never be voluntarily offered by the firm. However, there is a potentially nonempty range of costs $c \in (\underline{u}, \underline{u} + \alpha\Delta]$ such that (i) a RP would be beneficial for both consumer and total welfare, but (ii) is not offered voluntarily by the firm.