

Policy Evaluation in Behavioral Models*

Christian Michel André Stenzel

November 29, 2018

Abstract

This paper studies the data requirements to conduct an unambiguous policy evaluation when consumers may exhibit nonstandard preferences. We theoretically analyze how the equilibrium combinations before and after a policy introduction allow to clearly assess different economic indicators. Specifically, we study the data requirements for an unambiguous evaluation of two distinct consumer “cooling off” policies in a model in which consumers may exhibit a projection bias when making consumption decisions. We find that market-level data can help to make simple policy evaluations using baseline statistics and sometimes even discriminate between different classes of consumer preferences. We further discuss the generalizability to other models and the use for ex-ante policy design.

JEL Classification: D04, D18, D90

Keywords: Ex-post policy evaluation, Identification of nonstandard preferences, Projection bias, Cooling off, Behavioral economics

*We thank Larbi Alaoui, José Apesteguia, Jana Friedrichsen, Paul Heidhues, David Rojo Arjona, Elif-terea Trivizia, Chris Wallace, 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.

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 the behavior of consumers is often 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 One drawback of these models is that revealed consumer choices do not necessarily reflect consumer preferences, which greatly complicates a normative analysis.

This is particularly problematic in the context of policy introductions. For example, when evaluating whether a specific policy has increased consumer welfare, the assessment can differ substantially depending on whether consumers exhibit standard preferences or nonstandard preferences. Many consumer policies are explicitly adopted on the premise that consumers deviate from rational behavior. This leads to the question under what conditions an unambiguous evaluation of economic indicators such as consumer welfare is possible when it is not known ex-ante whether consumers exhibit standard or nonstandard preferences.

A small but growing literature seeks to structurally identify and estimate consumer preference parameters while accounting for the possibility that consumers have nonstandard preferences, see [DellaVigna \(2018\)](#) for an overview.² When such approaches are applicable, they allow for a narrow identification of consumer preferences and even for a counterfactual policy analysis. However, the issue of identification has so far been solved only for a relatively small subset of behavioral models and institutional settings. Identification in these models typically requires relatively detailed consumer data across multiple observation periods, which sometimes needs to be augmented by experimental data. Unfortunately, this type of data is not readily available in many real world situations. Furthermore, such structural models are relatively complicated to set up - and require significant time to implement.

In this paper, we focus on a much simpler approach to evaluate policies in a behavioral context. We start with a theoretical model that includes a behavioral foundation. This is motivated by the large number of such models that study the effects of an introduction of economic policies on indicators such as consumer or overall welfare.³ Based on the model, we conduct a theoretical analysis of supply and demand both before and after the introduction of a policy. The crucial step in our analysis is then to relate the equilibrium combinations pre and post policy adoption to the movement of different variables such as prices and quantities that a researcher might observe. This allows to determine the

¹Examples of nonstandard preferences include loss aversion, hyperbolic discounting, presence of a projection bias, and consumer myopia.

²Examples include [Laibson *et al.* \(2008\)](#) and [Mahajan *et al.* \(2018\)](#) for the case of hyperbolic discounting, [DellaVigna *et al.* \(2017\)](#) for reference dependence, and [Conlin *et al.* \(2007\)](#) for projection bias.

³See, for example, [Heidhues and Köszegi \(2010\)](#), [Inderst and Ottaviani \(2013\)](#), and the overview articles by [Armstrong \(2015\)](#) and [Grubb \(2015\)](#).

data requirements in terms of available variables that are sufficient to distinguish the different pre and post equilibrium combinations. In particular, it allows making inference about underlying model parameters, and about the changes of economic indicators such as consumer welfare implied by the theory model.

This type of analysis enables basic policy assessments in a comparatively short time period, and can in principle be conducted for any theory model with a behavioral foundation. We find a systematic study of necessary data requirements to conduct policy evaluations in behavioral models an important avenue that – to our knowledge – has not yet been extensively studied in the literature. When taken into account before a policy’s adoption, these considerations can provide insights regarding the gathering of the required data and thereby contribute to effective policy design.

In general, a researcher may often lack detailed knowledge about whether consumers exhibit nonstandard preferences in a specific market. In contrast, firms are typically much better informed because of their direct market experience. Thus, they can maximize profits given the specific consumer preferences and institutional setting. If there are changes to the institutional setting, the question is whether it is possible to find conditions under which the firms’ response to these changes allows for discrimination between different assumptions made in the underlying economic model, and even to unambiguously evaluate relevant economic indicators such as changes in consumer welfare. We explore one such change in detail, namely the introduction of a policy itself. A policy introduction often leads to a change in a firm’s choice set or to a change in the firm’s profitability, and thus to a different maximization problem. The combination of equilibrium outcomes before and after the policy introduction can reduce the number of cases where available data are consistent with the model. If the approach allows for isolating groups of equilibrium cases that involve the same behavior of an indicator of economic relevance, this indicator can be unambiguously assessed by the model if sufficient data are available. This may include a clear assessment of whether consumers experience standard or a specific form of nonstandard preferences.

[Section 2](#) introduces a simple model that nests standard preferences and a specific form of nonstandard time-inconsistent preferences on the demand side. In the model, a monopolistic firm offers a product to consumers. Consumer utility is state-dependent. Consumers can be motivated (“hot”) or unmotivated (“cold”) with respect to using the product, and this motivation can change over time. We allow for them to potentially experience a *projection bias*. This state-dependent bias makes consumers prone to overpredicting their future utility when they are in a “hot” consumption state and correspondingly prone to underpredicting it in a “cold” state. Crucially, there is a delay between the initial purchase decision and consumption. This relates to cases in which consumers

purchase products that they cannot consume immediately.⁴

As motivated consumers overpredict their future consumption utility, there are equilibrium cases in which the firm profitably sets prices that lead to a negative expected overall utility for consumers. This leaves scope for consumer policies aimed at increasing consumer welfare via changing the firm’s pricing incentives. We consider two distinct policies that are applied in various forms in practice and as such are of direct policy relevance in many countries; see, e.g., [Consumer Affairs Victoria \(2009\)](#). A *mandatory cooling off period* requires a consumer to confirm the product purchase after an initial waiting period.⁵ This delays consumption overall, but leaves the final consumption decision based on predicted utilities, which may suffer from the bias. In 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.⁶

In [Section 3](#), we analyze the data requirements to assess the two policies ex-post with regard to their effect on different indicators of interest, such as the directional changes in consumer and overall welfare, whether consumers experience a bias, or the fraction of motivated consumers. We find important differences in the data requirements across the two policies. The mandatory cooling off period has relatively low data requirements to assess the qualitative change in consumer welfare. The directional effects in consumer and total welfare can be evaluated when data are available on the directional change in quantity, or alternatively on the directional change in prices in combination with the frequency of consumers not confirming a letter of intent. In both cases, the policy further enables a partial assessment, i.e., an unambiguous assessment in a subset of observed

⁴There are many examples of such situations, including internet retail shopping for physical goods, or any other form of shopping where a good has to be pre-ordered before consumption. In such instances, a consumer’s realized utility when finally consuming a product often is state-dependent. Recent research has shown that when predicting their future consumption utility, consumers often put too much weight on their current state; see [Loewenstein et al. \(2003a\)](#) for a broad discussion of empirical and experimental evidence of such patterns. [Conlin et al. \(2007\)](#) estimate a structural model and find evidence of consumers experiencing a projection bias with respect to predicting future utility for clothes that they order from catalogs. [Busse et al. \(2015\)](#) provide evidence for consumers decision to buy a convertible similarly being related to current weather changes. Research from cognitive psychology furthermore shows that people who are in a specific emotional state tend to remember memories experienced in such a state better than those experienced in a different state; see for example [Bower \(1992\)](#).

⁵In 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.

⁶We choose to analyze these policies for several reasons. First, both policies are used in practice. Second, while these policies are often applied in different contexts, in some cases a policymaker may face the choice between different policies. This makes a normative comparison of these policies interesting. In [Section 3](#), we outline how our results can be generalized to other notions of nonstandard preferences, and other policies.

cases, about whether consumers experience a projection bias, and the fraction of motivated consumers.

A return policy has considerably higher data requirements to allow for these evaluations. Even a combination of data on quantities, prices, and return frequencies is not sufficient to clearly distinguish between all potential cases in terms of the directional change in consumer welfare. This is because it is not possible to observationally tell a particular case in which consumers have standard preferences apart from one in which they experience a projection bias. To be able to do so, a researcher must know the total market size, which is arguably a strong assumption. When data on quantities and return frequencies are available, they allow for a partial assessment of the directional change in consumer welfare, as well as of the existence of a projection bias and the fraction of motivated consumers. Furthermore, pricing data never provide added value when quantity data are already available, and can never fully substitute quantity data.

We believe that our paper highlights a simple and fast, yet important avenue to approach the consideration of a future policy evaluation already before an adoption using very basic statistics. Such analyses can be conducted for any behavioral model that theoretically explores the workings of a policy. This point and other generalizations and robustness checks are discussed in [Section 3.3](#). In many cases, such an analysis can also provide important insights regarding an efficient data collection process. We believe that this can contribute to a more systematic approach to consumer policy evaluation which in turn could become an integral part of the initial policy design. We conclude in [Section 4](#).

Literature In this paper, we analyze the data requirements to allow for unambiguous policy assessments using relatively simple statistics. This is based on relating the policy changes to predictions of underlying behavioral theory models. Our paper is thus mainly related to the literature on behavioral models with a policy component, especially those in behavioral industrial organization. From a theoretical viewpoint, we explore in detail the effects of projection bias on firm pricing, and analyze the supply-side responses of strategical firms for different cooling off policies. In our complementary note [Michel and Stenzel \(2018\)](#), we formalize a more general framework for how to assess the necessary data requirements to be applicable for different settings.

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, and [Mahajan *et al.* \(2018\)](#) using a field experiment. [Hinnosaar \(2016\)](#) empirically studies whether time-inconsistent preferences can provide a justification for restricting the number of hours in a day in which consumers can buy alcohol. She finds that while such a sales restriction might be welfare-improving, it is inferior to using an alcohol tax. [DellaVigna *et al.* \(2017\)](#) provides an extensive discussion of the advantages and possible avenues of

using structural models with potential nonstandard components. In general, they make use of relatively detailed consumer-level data and of the specific setting to discriminate between time-consistent and time-inconsistent behavior, as well as between naive and sophisticated behavior. The combination of feasible setting and data requirements are relatively difficult to meet to directly estimate the preferences using a structural approach. As a consequence, we do not focus on the consistency of choices over different time periods for different options, but rather on incorporating a combination of firm behavior and consumer choices to make simple predictions when the class of consumer preferences is unclear ex-ante.

Further work seeks to provide tests for non-standard preferences that do not rely on fully structural models, particularly in the context of salience. [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 \(2004\)](#) use the difference in introductory credit card interest rates and post-introductory interest rates chosen by borrowers as a test of their rationality. They find that the ratio between these two interest rates is too low for the choices being explained by consumer rationality. [Heidhues and Koszegi \(2018\)](#) discuss how to test for consumers not responding to an increase in add-on prices. They propose to use either 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 to test for this property. Such an approach can be used in absence of any policy introduction. However, it requires more detailed data on consumer demand than our approach, and in some cases a full demand estimation.

From a theory viewpoint, [Loewenstein *et al.* \(2003a\)](#) explicitly define the notion of projection bias and give a formal basis for the utility representation of our model. While the authors argue that a mandatory cooling off period might prevent consumers from making mistakes, their analysis abstracts from endogenous pricing of a firm. Using a structural econometric model, [Conlin *et al.* \(2007\)](#) find evidence that consumers experience 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.

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. [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. In their model, a sufficiently high minimum refund level in the presence of credulous consumers

is both a consumer surplus and efficiency-enhancing tool, as it lowers the expected rents that the firm receives from false suitability claims. 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, and underline the non-distorting character of return policies regarding consumer choices.

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, in our case changing the pricing of products given a specific policy.

2 The Model

We first set up a model to theoretically analyze the effects of different cooling off policies. The model is later used to determine the data requirements for being able to unambiguously evaluate these policies.

We consider a firm that produces a single good of fixed quality at constant marginal cost $c \geq 0$. The firm can choose 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. However, the consumption of the product takes place one period after purchase because of shipping, such that the model has an intertemporal component.

Time is discrete and covers at most three periods: $t = 1, 2, 3$. We assume that 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 current motivation state $s_t \in \{\underline{s}, \bar{s}\}$.⁷ 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. $\Delta \equiv \bar{u} - \underline{u}$ characterizes the utility difference.

The product price affects consumer utility negatively and linearly. A consumer's overall utility when purchasing a product in period t depends on her gross consumption

⁷Throughout the paper, we use the state-dependent utility notation $u(s_t)$ for ease of readability instead of a more complicated notation $u(\cdot; s_t)$.

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 when making the purchase decision. 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*.

Each consumer correctly assigns probability $\mu \in (0, 1)$ to the event of being motivated in a future period, and a probability $1 - \mu$ to the event of being unmotivated. The probability of being in a certain motivation state in the next period is independent of the current state.⁸ Under the projection bias, consumers partially incorporate the utility associated with their current state into the prediction of the consumption utility also for any other state. This implies a correct assessment of the utility in the same state in the future, but implies a bias when making a prediction for a different state. Following the definition in [Loewenstein *et al.* \(2003b\)](#), 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 still correctly predict the utility of consuming the good in the same state in the future, $\tilde{u}(s|s) = u(s)$. However, they overpredict their consumption utility for being in the unmotivated state in the future whenever they are currently motivated, and underpredict their consumption utility for being in the motivated state in the future when they are currently unmotivated:

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

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

where 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 Recall that time covers at most three periods. The third period is relevant only for the model variant incorporating a mandatory cooling off period. In the baseline model, at the start of the first period, the firm sets a price p and consumers

⁸Introducing correlation between motivation states would alter the quantitative, but not the qualitative predictions of the model. According to [Consumer Affairs Victoria \(2009\)](#), "behavioural traits suited to cooling off are those in which consumers are likely to change their minds relatively quickly. For example, purchases made in hot or emotional states are likely to be more transient than some other types of hyperbolic discounting."

observe their motivation state.⁹ Given the price and motivation state, consumers predict their second period consumption utility, 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. This decision takes place *after* observing their second-period motivation. Only if the intent to purchase is confirmed, a transaction actually takes place. 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. The cooling off period thus does not eliminate the relevance of the projection bias for consumers’ final purchase decision, but makes it harder for firms to profitably extract more than expected consumer utility. As shown in the model analysis, making use of consumers’ overprediction of expected consumption utilities is only feasible for the firm if they remain in the motivated state for two periods.

Second, we consider a voluntary return policy. Whenever this policy is applicable, the first period is identical to the baseline model. However, when the good arrives at the start of the second period, consumers observe their second-period motivation and have the opportunity to return the good before consumption occurs. By returning the good, they are reimbursed the price p and forego consumption. As such, consumers are able to assess their utility *in the current motivational state* right before potentially consuming the product. A return policy thus gives consumers perfect information about their utility when they make the final consumption decision.¹⁰ In the main analysis, we abstract from costs associated with return both on the consumers’ side, and the firms’ side. The model’s qualitative predictions carry over to a setting where the firm incurs a fraction $\eta \in (0, 1]$ of the production cost for each good that is initially shipped out but returned, which we

⁹We analyze the workings of different consumer policies when only a single firm is available. In many cases, competition among firms can naturally lead to a reduction in prices. However, 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. Third, while there might not be an entire industry specializing – and thus competing – in fraudulent products, single firms can try to profitably use fraudulent practices, which is in itself a potential justification for the adoption of protection policies.

¹⁰For 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, and the return decision is made based on that information.

analyze in detail in the appendix.

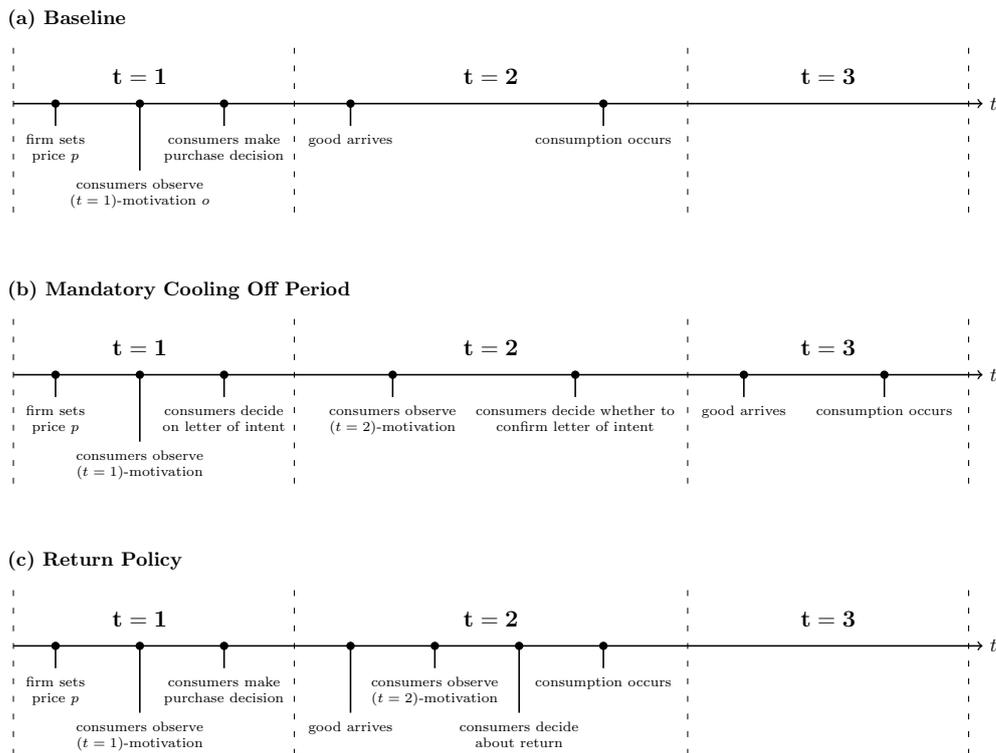


Figure 1: Timing of the Game

A key difference between the 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.¹¹ 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 in this case. Figure 1 contrasts the timeline of the game in the three settings.

¹¹Our analysis abstracts from informational gains for the consumer during the cooling off period, such that changes in consumers' preferences are solely driven by a change in their motivational state. In our model, incorporating this would magnify the effect of a cooling off period without altering the results qualitatively.

2.1 Model Analysis

We next summarize the main results for our model analysis. In [Appendix A](#) we provide the related proofs and additional explanations. 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 still 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 $\delta \equiv \delta_2$.

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.

Targeting only motivated consumers 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 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} . Denoting a motivated consumer's prediction for the next period's expected utility by

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

and, similarly, an unmotivated consumer's prediction by $\underline{u} = \underline{u} + \mu(1 - \alpha)\Delta$, we obtain the following proposition.

Proposition 1 (Baseline Pricing & Welfare) *There exists a unique threshold $\tilde{c} < E[u]$ which determines the firm's pricing strategy.*

- (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*

$$\widetilde{CW} = -\mu(1 - \mu)\alpha\Delta \leq 0, \text{ and total welfare is given by } \widetilde{TW} = \mu(\underline{u} + \mu\Delta - c).$$

The cost threshold \tilde{c} is given by $\tilde{c} = \underline{u} + \frac{\mu}{1-\mu}\Delta(1 - 2\alpha - (1 - \alpha)\mu)$, and is increasing in \underline{u} and decreasing in α .

Proof. Contained in [Appendix A](#). ■

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.

Cooling Off 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 consumers 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 therefore do not think it is worth registering because they do not think that their prediction about the usefulness of a 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 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 .

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

There exists a unique threshold \tilde{c}_c with $\tilde{c}_c \in [\tilde{c}, E[u]]$ 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 $\widetilde{TW}_c = \delta\mu^2(\underline{u} + \mu\Delta - c)$.

The cost threshold \tilde{c}_c is given by $\tilde{c}_c = \underline{u} + \frac{\mu}{(1-\mu)}\Delta \frac{(1-(1+\mu)\alpha-(1-\alpha)\mu^2)}{1+\mu}$, and is increasing in \underline{u} and decreasing in α . For $\alpha > 0$, $\tilde{c}_c > \tilde{c}$.

Proof. Contained in [Appendix A](#). ■

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 $\alpha > 0$, 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. 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 ($c > E[u]$) for total welfare.

		Cooling Off Period	
		full coverage $c \leq \tilde{c}_c$	exclusive targeting $c > \tilde{c}_c$
B a s e l i n e	exclusive targeting $c > \tilde{c}$	$\Delta p < 0$ (motiv. to unmotivated) $\Delta q > 0$ (from μ to 1) no stepping back $\Delta CW > 0$ $\Delta TW > 0$	$\Delta p \approx 0$ (depends on δ) $\Delta q < 0$ (from μ to μ^2) fraction μ steps back $\Delta CW > 0$ ΔTW indetermined
	full coverage $c \leq \tilde{c}$	$\Delta p \approx 0$ (depends on δ) $\Delta q = 0$ (full market coverage) no stepping back $\Delta CW \approx 0$ $\Delta TW \approx 0$	

The figure depicts the directional changes in prices, consumed quantities, consumer welfare, and total welfare for the different combination of marginal cost thresholds 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

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 product price. 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. As such, the consumer keeps the good if and only if her actual consumption utility is at least as high as the price. By contrast, the initial purchase decision is still driven by the predicted utilities in each motivation state, which incorporate the projection bias. A consumer hence initially purchases if and only if the price the firm charges is equal to or below the predicted consumption utility in at least one motivational state.

Having characterized the initial purchase and return decision, the firm's pricing again depends on the marginal cost of production, but 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} . 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. If the bias α is not too

large, the gain in profits from charging a higher price is sufficient to offset the lost sales from returned goods, while the fact that the fraction of motivated consumers is not too high can keep this strategy preferred to fully exclusive targeting.

Proposition 3 (Firm's pricing given return policy) *There exist thresholds \tilde{c}_r and \underline{c}_r where $\underline{c}_r \leq \underline{u}$ and $\tilde{c}_r \geq \tilde{c}$, which determine the firm's pricing decision in presence of a return policy.*

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

The firm charges $\underline{p}_r = \underline{u}$ to reap profits $\underline{\pi}_r = \underline{u} - c$, while consumer surplus is $\underline{CW}_r = \mu\Delta$ 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 are motivated in the second period. The firm charges $\underline{p}_r = \tilde{u}(\bar{s}|\underline{s})$ to reap profits $\underline{\pi}_r = \mu \cdot [\underline{u} + (1 - \alpha)\Delta - c]$, while consumer surplus is $\underline{CW}_r = \mu\alpha\Delta$, 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 unless they remain motivated in the second period. The firm charges $\bar{p}_r = \bar{u}$ to reap profits $\bar{\pi}_r = \mu^2 \cdot (\underline{u} + \Delta - c)$, while consumer surplus is $\bar{CW}_r = 0$, and total welfare $\bar{TW}_r = \mu^2(\underline{u} + \Delta - c)$.*

$(\underline{c}_r, \tilde{c}_r]$ is nonempty if and only if $\alpha < \frac{1}{1+\mu}$, otherwise it holds that $\underline{c}_r = \tilde{c}_r \leq \underline{u}$.

Proof. Contained in [Appendix A](#). ■

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 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. This has implications for the impact of the adoption of a return policy.

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. 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

		Return Policy		
		full coverage $c \leq \underline{c}_r$	intermediate targeting $c \in (\underline{c}_r, \bar{c}_r]$	exclusive targeting $c > \bar{c}_r$
Baseline	exclusive targeting $c > \bar{c}$	$\Delta p < 0$ $\Delta q > 0$ (partial to full market) no return $\Delta CW > 0$ $\Delta TW > 0$	Δp ambiguous $\Delta q = 0$ (only return affected) fraction $1 - \mu$ return $\Delta CW > 0$ $\Delta TW > 0$	$\Delta p > 0$ $\Delta q < 0$ fraction $1 - \mu$ return $\Delta CW > 0$ ΔTW ambiguous
	full coverage $c \leq \bar{c}$	$\Delta p < 0$ $\Delta q = 0$ (full market coverage) no return $\Delta CW > 0$ $\Delta TW = 0$	$\Delta p > 0$ $\Delta q < 0$ (full market to partial) fraction $1 - \mu$ return $\Delta CW = 0$ ΔTW ambiguous	

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

Figure 3: Effects of the introduction of a return policy

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 consumers have a 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. This is because 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. In case the firm changes from exclusive to intermediate targeting, the price effect is ambiguous. If the projection bias α is large, the price with the intermediate targeting strategy is low as $\underline{p}_r = \tilde{u}(\bar{s}|\underline{s}) = \underline{u} + (1 - \alpha)\Delta$ is decreasing in α —the utility unmotivated consumers predict for the motivated state is low. This would result in prices adjusting downward. By contrast, if the bias is low, \underline{p}_r is high, and the price with exclusive targeting pre-introduction is lowered as the overprediction of motivated consumers for the unmotivated state is lessened. In this case, the switch to the intermediate targeting strategy would be associated with a price increase.

If the firm targets all consumers before the policy introduction, a return policy in-

creases consumer welfare if even second-period-unmotivated consumers keep the good, as the associated price decrease to 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 exactly offset by the lower market coverage. 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. 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 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](#).

Robustness The above analysis shows that the considered policies lead to different predictions regarding the associated effects on market outcomes and efficiency. [Appendix A.4](#) assesses the robustness of the model in various dimensions. We demonstrate 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. A mandatory institution of this policy can be consumer welfare enhancing. We also show that the inclusion of positive return costs into the model leaves the qualitative predictions unchanged, and discuss the comparison between the policy options. Notably, while a mandatory cooling off period is not always able to prevent negative consumer utility, it can be the consumer welfare maximizing strategy. This is because the incentive to target the full market is stronger with a cooling off period.

3 Data requirements for policy evaluation

This section uses our model analysis to assess the data requirements that allow for unambiguously assessing a policy’s effect on several economic indicators ex-post. This includes the ability to evaluate a policy with respect to the sign change in consumer and total welfare, respectively. Furthermore, we discuss under what 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 $\alpha > 0$.

Regarding aggregate data, we examine which combination of the following variables can be used to assess a policy with respect to different indicators: prices, final quantities (i.e., quantities sold net of return or stepping back), return frequencies in case of a return

policy, frequency of consumers not confirming purchase in case of a mandatory cooling off period, and knowledge of the potential market size. As indicators, we consider changes in consumer welfare, total welfare, whether consumers exhibit nonstandard preferences, and the relative range of marginal cost. Ideally, we would like to use the data to unambiguously determine the respective model case and thus allow for a clear evaluation of the indicator.

3.1 Data requirements in the mandatory cooling off case

Figure 4 provides an illustration of the different potential outcomes of the introduction of a mandatory cooling off policy. We map all of the outcomes into a space that relates the implied changes of the policy with respect to quantity to the implied changes with respect to price.

	final quantity decreases	final quantity constant	final quantity increases
price increases			
price relatively constant	<u>observe stepping back</u> $c > \tilde{c}_c > \tilde{c}$ $p_{pre} = \tilde{u}, p_{post} = \delta \tilde{u}$ $q_{pre} = \mu, q_{post} = \mu^2$ Bias: $\alpha > 0$	<u>observe no stepping back</u> $c \leq \tilde{c} < \tilde{c}_c$ $p_{pre} = \underline{u}, p_{post} = \delta \underline{u}$ $q_{pre} = 1, q_{post} = 1$	
price decreases			<u>observe no stepping back</u> $c > \tilde{c}, c \leq \tilde{c}_c$ $p_{pre} = \tilde{u}, p_{post} = \delta \tilde{u}$ $q_{pre} = \mu, q_{post} = 1$ Bias: $\alpha > 0$

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. The figure also depicts the cost range and prices and final quantities pre and post policy introduction.

Figure 4: Ex-post evaluation of cooling off period

Focusing on price and quantity changes alone already shows first differences in the informational content of these variables for a policy evaluation. First, 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 implied by the model before and after the policy introduction that fits such a movement. By contrast, this is not the case for the directional movement in price. If the price stays relatively constant, this can be because the firm always targets all consumers, or because it always targets only unmotivated consumers. We use the term relatively constant because the discount factor δ implied by the waiting period can lead to a decrease in the price. The decrease

would be the same for both cases and, for a discount factor close to 1, is negligible.¹²

Only observing how many consumers step back from the purchase decision after signing a letter of intent does not allow unambiguously identifying the model case. This is because no consumers stepping back from the purchase decision could be attributed both to the case for which all consumers are targeted with or without policy, or to the case for which all consumers are targeted only following the policy introduction. If the discount factor is close to 1, observing the price change in addition allows assessing the directional change of consumer welfare and total welfare – the stepping back behavior allows distinguishing between the two cases where prices are relatively constant.

In addition to the overall assessment, information about other market fundamentals can also be gathered. Unambiguously recovering all the Δq - Δp cases is equivalent to recovering the firm's relative cost range. In addition, if stepping back is observed, the fraction of motivated consumers μ is simply given by $\frac{q_{post}}{q_{pre}} = \frac{\mu^2}{\mu} = \mu$. Similarly, if quantities increase due to the policy introduction, the pre policy market quantity q_{pre} can be used to recover μ . Finally, in the cases where the market is only partially covered prior to the policy introduction, a positive projection bias must be present, i.e. $\alpha > 0$. In the opposite case, i.e. if the market is fully covered before the policy introduction, which implies that it will remain so after introducing a cooling off period, it is not clear whether consumers are unbiased, $\alpha = 0$, or whether the marginal cost of production is sufficiently low. Our findings are summarized in the following proposition.

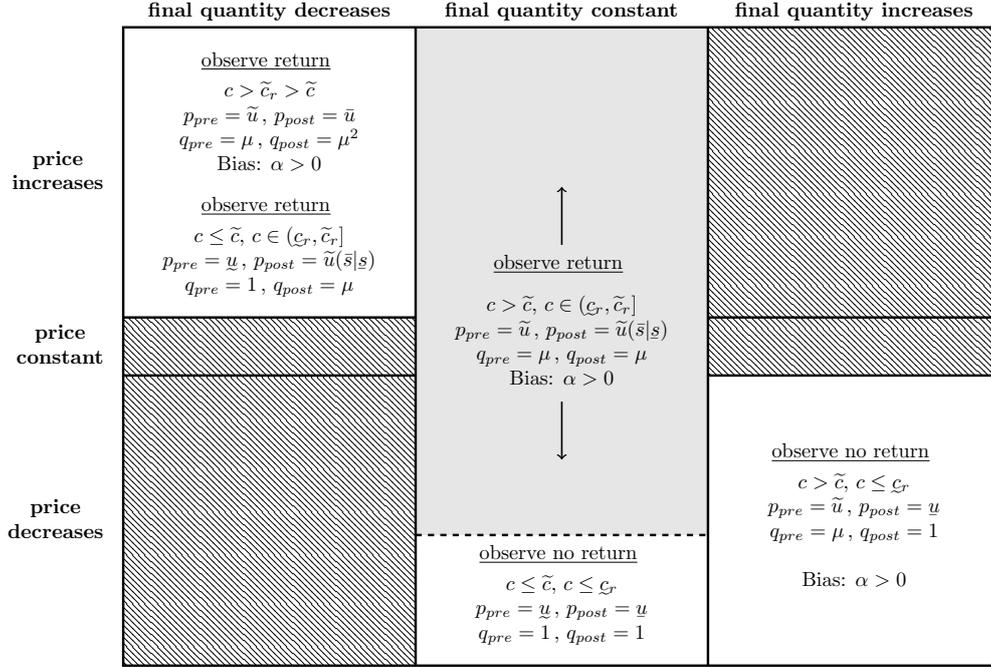
Proposition 4 *Data on quantities or data on both prices and confirmation rates are sufficient to assess the directional change in consumer welfare and the relevant cost range following the adoption of a mandatory cooling off period. In some cases, the existence of a projection bias, i.e. $\alpha > 0$, and the fraction of motivated consumers, μ , can also be recovered.*

Proof. Follows from the previous discussion. ■

¹²Such considerations are not important for the change in quantities after the policy introduction. The actual price behavior depends on the discount factor – both when quantities remain constant and when they decrease, the price stays constant for δ close to 1, i.e. almost no discounting, and decreases otherwise. Given the short time frame associated with cooling off policies in practice, we view this as a reasonable case. The identification and estimation of time preferences is still a relatively difficult task, particularly in the behavioral context; see, for example, [Mahajan et al. \(2018\)](#) for a structural approach.

3.2 Data requirements in the return policy case

Figure 5 illustrates the different outcomes for the introduction of a return policy.¹³ One can easily see that the implied changes in quantity, prices, or their combination, cannot unambiguously distinguish among all the emerging cases. This is because some cases yield the same price-quantity predictions.



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 $\alpha \in [0, 1)$. The figure also depicts the cost range and prices and final quantities pre and post policy introduction.

Figure 5: Ex-post evaluation of return policy

If the quantity decreases post-introduction, both the case in which motivated consumers are targeted exclusively pre and post policy adoption, and the case in which there is a switch from targeting all consumers to pricing at unmotivated consumers' predicted utility for the high motivation state $\tilde{u}(\bar{s}|\underline{s})$ are associated with a price increase. Moreover, it may be the case that the quantity stays constant after the policy introduction, i.e. $\Delta q = 0$, while the price decreases.

The case in which the firm switches from exclusive targeting to intermediate targeting can also be associated with a price increase or prices staying constant, in which case the

¹³For ease of exposition both Figure 5 and the discussion implicitly assume $\alpha < 1$. For $\alpha = 1$, i.e. if consumers always predict their utility to be identical to the one they currently experience, two special cases materialize. First, it is possible that a quantity decrease is associated with constant prices and no return if $c > \tilde{c}_r > \tilde{c}$. Second, $c \leq \min\{\tilde{c}, \underline{c}_r\}$ would lead to constant prices, constant quantities, and no return. In both cases, having the combination of price data, quantity data, and return data available allows a researcher to clearly identify $\alpha = 1$.

combination of quantity and price data would allow for distinguishing this case from the one where the market is always fully covered (as prices always decrease in this case).

This leads to the question of whether additional data are able to distinguish among the different cases such that the latent indicators can be assessed. One candidate is data on return frequencies, which can be gathered by firms and made available to researchers or policymakers after the introduction of such a policy. However, the inclusion of these data is only able to resolve the issue when there is no change in the final quantity following the policy adoption. If all consumers are targeted, there should be no product return by consumers. By contrast, if the firm prices the product at unmotivated consumers' predicted utility for the high motivation state $\tilde{u}(\bar{s}|\underline{s})$, consumers who are again unmotivated in the consumption period will return the product.

When the final quantity decreases following the policy introduction, even the use of return data is insufficient because both relevant cases are associated with positive return rates. One strong assumption that would suffice in this case is that the researcher knows the total market size. Under this assumption, it is possible to distinguish between the case in which the market is fully covered and the case in which only motivated consumers buy the product by comparing the consumed quantities relative to the overall market size.¹⁴

An unambiguous evaluation of the directional consumer welfare change would be possible if the data are able to distinguish between all the different cases. However, if the cases cannot be distinguished, an evaluation is no longer as straightforward. If no data on market size are available, once quantity decreases, the direction of the change in consumer welfare can no longer be unambiguously determined. In one of the cases, consumer welfare increases ($c > \tilde{c}_r$), while in the other it is constant ($c \leq \tilde{c}$ and $c \in (\underline{c}_r, \tilde{c}_r]$). Recall that the neutral effect of the policy hinges on return being costless. If the quantity remains constant, the directional impact on consumer welfare can be determined even without data on market size or return frequencies. This is because even if prices decrease in both cases (which renders them indistinguishable using the available data), they always lead to an increase in consumer welfare, and only differ in magnitude.

Whenever the change in observable 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, μ can be recovered. For example, for $c > \tilde{c}_r > \tilde{c}$, i.e., if a price increase and quantity decrease are observed, μ is given by $\frac{q_{post}}{q_{pre}}$. We sum up our findings in the next proposition.

Proposition 5 *Data on quantities, return frequencies and market size are sufficient to*

¹⁴Modern 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. In the context of our model, this assumption seems to be even stronger because it implies that in some cases all potential consumers already consume the product.

assess 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.■

One notable observation is that the availability of price data when a return policy has been established provides little added value with respect to distinguishing among different cases. In the absence of quantity data, it only allows for to clearly identifying a single case in the one instance where the switch from exclusive to intermediate targeting is associated with a constant price. Otherwise, it only broadly separates cases in which motivated consumers will be targeted from cases in which all consumers will be targeted.

When quantity data are available, price data similarly lead to an additional gain in information only in the case of constant quantities and a sufficiently low bias that the switch from exclusive to intermediate targeting is not associated with a price decrease, which cannot be ex-ante assumed. Overall, this is in contrast to the cooling off policy, in which price data in combination with frequencies of purchase confirmations were a substitute to the availability of quantity data.

Corollary 1 *Price data can never substitute quantity data in evaluating the effects of a return policy.*

We next discuss the extent to which our approach can be applied to other model settings, as well as how to empirically implement the assessment.

3.3 Generalizations and applicability

In principle, the necessary data requirements to be able to assess a policy with respect to an economic indicator can be analyzed for any underlying theoretical model. In the companion note [Michel and Stenzel \(2018\)](#) we formalize a simple framework to generalize how to assess the necessary data requirements for arbitrary models.

Other forms of nonstandard consumer behavior As an example in the context of consumers with potentially nonstandard beliefs, consider the case in which consumers may be fully credulous towards the advice of a seller as in [Inderst and Ottaviani \(2013\)](#). Full credulity refers to consumers always following the advice of a seller before actually receiving a product. A return period in the absence of return costs gives a firm the incentive

to recommend only suitable products to consumers because otherwise those consumers would always give back the product after learning that it is not suitable. By contrast, a mandatory cooling off period does not alter a firm’s targeting incentives because consumers would still believe a firm’s recommendation in the second period and there is no state-dependent utility. Relating these cases to the necessary data requirements, quantity data alone would be sufficient to assess consumer welfare and marginal cost ranges if a return period is instituted and allow for assessing consumer credulity. By contrast, a mandatory cooling off period would never be informative of these benchmarks. An interesting implication from this example is the potential to distinguish between different forms of nonstandard consumer preferences. For example, if a mandatory cooling off period is adopted, and a firm changes its pricing behavior, this would be consistent with consumers exhibiting a projection bias, but not with them being overly credulous.

Effect of competition and 2005 UK Extended Warranty Order Our model focuses on a single firm on the supply side for the benefit of an easy illustration. However, a similar analysis can be conducted with an oligopolistic market structure. Consider, for example, the prominent case in which firms compete with each other for a base product that might lead to add-on costs, or the voluntary purchase of add-on products (see, e.g., [Gabaix and Laibson \(2006\)](#), [Heidhues *et al.* \(2016\)](#), [Michel \(2017\)](#)). A policymaker might be unsure about whether consumers are aware of the implied add-on costs, or whether they ignore them, and thus consider the adoption of a policy to increase transparency in the market for the add-on product. The policy’s implications regarding the changes in quantities and prices then differ depending on whether consumers anticipated these costs in the first place. If they do, add-on products should be priced competitively before the adoption of the policy, and prices and quantities would not change significantly following the policy adoption. In the event that some consumers are myopic, however, added transparency would lead to a reduction in add-on prices due to the competition effect. Depending on the fraction of naive consumers in the population, this would go along with an increase in the base price due to a lower degree of loss-leadership for the base products. Looking only at the base prices in the market can thus be misleading: if the base product price goes up following the policy introduction, this can be a potential sign of an effective policy because it limits the loss-leadership of firms. To really see whether this policy is helpful, one can further look at data from add-on prices and quantities sold both before and after the policy adoption. A decrease in the add-on price would give further insights into firms’ responses to the policy. By contrast, only looking at add-on quantities cannot always lead to an unambiguous evaluation: depending on the change in pricing because of the policy and the usefulness of the add-on, quantities could go both up and down in case a policy is effective.

One such example of a policy introduction is the “The 2005 UK Extended Warranty

Order”. In 2003, the UK’s Competition Commission found that a “complex monopoly exists” for extended warranties sold via retailers in the consumer electronics industry. The Competition Commission argued that prices for extended warranties in this industry have typically not been displayed, cancellation periods for contracts have been short, and consumers have lacked both information and choices about different extended warranty providers. As a result, the legislature passed the Supply of Extended Warranties on Domestic Electrical Goods Order in 2005. The order included a cancellation right for consumers, incorporating the right to shop around for other extended warranty quotations, and a mandatory information leaflet that was to be provided before signing. The policy led to an increase in consumers shopping around for extended warranty contracts instead of primarily buying them at a retail store. It further led to a decrease in the prices of extended warranties. As a measure of value-for-money, [OFT \(2012\)](#) finds that the policy led to an increase in the claims ratio, defined as the incurred claim costs relative to the insurance payments paid, from 20 percent to approximately 50 percent six years after the policy was introduced.

Firms’ anticipation of policy evaluation One facet to account for when tailoring a policy and data gathering such that an ex-post evaluation is feasible is 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. If they are sufficiently patient, they could wait until the review of a policy, after which they coordinate to a more profitable equilibrium. We believe, however, that this fear is largely unfounded. First, even if data are gathered, 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 greatly lowers the incentives for firms to strategically hide their true incentives in the short run for the benefits of higher long-run profitability.

Estimation issues An important question is how an evaluation as outlined in this paper can be implemented empirically. A detailed discussion should depend on the specific policy and institutional setting and is beyond the scope of this paper. We briefly discuss several potential avenues.

First, it is possible to use descriptive statistics for the different observable variables and look at the change in the average movement. Second, a time series or panel data analysis can be conducted, for example to control for additional variables apart from the policy. Third, in case the policy is adopted in only a specific region or industry, the availability of plausibly well-suited control groups that are similar in terms of characteristics and not affected by the policy may allow for conducting a difference in difference analysis. This would provide additional information on the movements of prices and quantities after the policy introduction relative to the control group. In all cases, it is possible to test for

significance of the directional movements of observable variables such as quantities, prices, or return frequencies using a t-test, or to test the joint significance of these changes using an F-test.

Negative model test In the event that the observable data before and after a policy introduction are not consistent with any of the potential 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 of our model, it would “falsify” the model. This can be seen as a negative test for the underlying assumptions. 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. If one observes 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.

4 Conclusion

In light of the increased popularity of consumer policies that are based on a “behavioral justification” such as the presence of consumer biases, a systematic exploration of the necessary data requirements to evaluate such policies seems to be an important avenue.

We contribute to this task by providing a simple approach to analyze such requirements when it is not clear whether consumers experience nonstandard preferences. This is based on analyzing the consequences of a policy introduction in a theoretical model and on tracking which data are able to clearly distinguish between the different equilibrium combinations to credibly make an assessment of implied model parameters and economic indicators. Exploiting the assumption of profit-maximizing supply-side behavior can oftentimes help conducting this type of analysis, which is especially useful when revealed preference approaches on the demand side cannot be used. Such optimality assumptions on the supply side are implicitly made in most behavioral theory models that incorporate strategic firm behavior.

We illustrate our approach by considering a setting in which a firm faces consumers that may experience a projection bias when predicting their future utility. In this setup, we analyze the introduction of two consumer policies. We find that a mandatory cooling off period has considerably lower data requirements to allow for an ex-post assessment than a return policy. Quantity data by themselves, or price data in combination with consumers’ confirmation rates are sufficient to assess whether the mandatory cooling off

period increased or decreased consumer welfare. Furthermore, for a subset of cases the same data allows to assess the directional change in total welfare, whether consumer experience standard preferences or a projection bias, and the fraction of motivated consumers. In contrast, if a return policy is adopted, the data requirements for an unambiguous policy assessment are considerably higher. The minimum requirements to allow for an unambiguous ex-post assessment of the directional change in consumer welfare are the availability of data on quantities, return frequencies, and the overall market size. In this case, we show that price data are redundant, while both quantity data and data on return frequencies are always necessary.

The difference in the data requirements between the two policies indicates the importance of a well-tailored data collection for an accurate policy evaluation, which should ideally be incorporated in the policy design itself. Because of the simplicity of looking at the necessary data requirements from a theoretical viewpoint, we believe that an analysis like ours can easily be applied to many settings featuring a behavioral foundation.

We see the present paper as a starting point for further investigations into how supply-side responses can be informative of underlying market settings and consumer behavior. Furthermore, we believe that effective policy design should already include the question of how a policy can be evaluated ex-post. Whether the requirement of a policy being easily evaluated ex-post constitutes a strong restriction for policy design is an important question for future work.

References

- ARMSTRONG, M. (2015). Search and ripoff externalities. *Review of Industrial Organization*, **47** (3), 273–302.
- BERNHEIM, B. D. (2009). Behavioral welfare economics. *Journal of the European Economic Association*, **7** (2-3), 267–319.
- and RANGEL, A. (2009). Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics. *Quarterly Journal of Economics*, **124** (1), 51–104.
- BOWER, G. (1992). How might emotions affect learning. In S. Christianson (ed.), *The Handbook of Emotion and Memory: Research and theory*, Lawrence Erlbaum Associates, pp. 3–31.
- BUSSE, M. R., POPE, D. G., POPE, J. C. and SILVA-RISSO, J. (2015). The psychological effect of weather on car purchases. *The Quarterly Journal of Economics*, **130** (1), 371–414.

- CAMERER, C., ISSACHAROFF, S., LOEWENSTEIN, G., O'DONOGHUE, T. and RABIN, M. (2003). Regulation for conservatives: Behavioral economics and the case for “asymmetric paternalism”. *University of Pennsylvania Law Review*, pp. 1211–1254.
- CHETTY, R., LOONEY, A. and KROFT, K. (2009). Salience and taxation: Theory and evidence. *American economic review*, **99** (4), 1145–77.
- CONLIN, M., O'DONOGHUE, T. and VOGELSANG, T. (2007). Projection bias in catalog orders. *American Economic Review*, **97** (4), 1217–1249.
- CONSUMER AFFAIRS VICTORIA (2009). Cooling-off periods in victoria: their use, nature, cost and implications. *Research Paper*, (15).
- DELLAVIGNA, S. (2018). *Structural behavioral economics*. Tech. rep., National Bureau of Economic Research.
- , LINDNER, A., REIZER, B. and SCHMIEDER, J. F. (2017). Reference-dependent job search: Evidence from hungary. *The Quarterly Journal of Economics*, **132** (4), 1969–2018.
- FISHER, G. and RANGEL, A. (2014). Symmetry in cold-to-hot and hot-to-cold valuation gaps. *Psychological science*, **25** (1), 120–127.
- GABAIX, X. and LAIBSON, D. (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics*, **121** (2), 505–540.
- GRUBB, M. D. (2015). Overconfident consumers in the marketplace. *Journal of Economic Perspectives*, **29** (4), 9–36.
- HEIDHUES, P. and KŐSZEGI, B. (2010). Exploiting naivete about self-control in the credit market. *American Economic Review*, **100** (5), 2279–2303.
- and KOSZEGI, B. (2018). Behavioral industrial organization. *Handbook of Behavioral Economics*, Elsevier.
- , KŐSZEGI, B. and MUROOKA, T. (2016). Inferior products and profitable deception. *The Review of Economic Studies*, **84** (1), 323–356.
- HINNOSAAR, M. (2016). Time inconsistency and alcohol sales restrictions. *European Economic Review*, **87**, 108–131.
- INDERST, R. and OTTAVIANI, M. (2013). Sales talk, cancellation terms and the role of consumer protection. *The Review of Economic Studies*, **80** (3), 1002–1026.

- LAIBSON, D., REPETTO, A. and TOBACMAN, J. (2008). Estimating discount functions with consumption choices over the lifecycle. *mimeo*.
- LOEWENSTEIN, G., O'DONOGHUE, T. and RABIN, M. (2003a). Projection bias in predicting future utility. *Quarterly Journal of Economics*, **118** (4), 1209–1248.
- , — and — (2003b). Projection bias in predicting future utility. *Quarterly Journal of Economics*, **118** (4), 1209–1248.
- MAHAJAN, A., MICHEL, C. and TAROZZI, A. (2018). Identification of time inconsistent models. *mimeo*.
- MICHEL, C. (2017). Market regulation of voluntary add-on contracts. *International Journal of Industrial Organization*, **54**, 239–268.
- and STENZEL, A. (2018). A simple framework to analyze data requirements for policy evaluation. *mimeo*.
- OFT (2012). Extended warranties on domestic electrical goods. *Office of Fair Trading market study report OFT1403*.
- 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 AUSUBEL, L. (2004). Time inconsistency in the credit card market. *mimeo*.

A Model Analysis

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.

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

$$\underline{u} \equiv \mu\tilde{u}(\bar{s}|\underline{s}) + (1 - \mu)\tilde{u}(\underline{s}|\underline{s}) = \underline{u} + \mu(1 - \alpha)\Delta. \quad (\text{A.2})$$

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, $\underline{u} < E[u]$.

For expositional purposes, we express prices charged by the firm in terms of second-period-consumption units. As consumption occurs at the earliest in the second period, this minimizes the carrying of discount factors which do not qualitatively alter the analysis. Moreover, we make the tie-breaking assumption that in case of equal profits, the firm prefers the strategy which yields the larger market share. Throughout the analysis, we express the firm's targeting and pricing as a function of its marginal cost. The fraction of motivated consumers in each period μ , the projection bias α , and the actual consumption utilities characterized by \underline{u} and Δ are treated as exogenous parameters. 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 for sufficiently high marginal cost.

Notation Throughout the analysis, both a subscript or superscript \sim denotes that pricing is based on predicted utilities (which suffer from the projection bias), while both a subscript or superscript $-$ indicates that pricing is based on actual consumption utilities. A superscript generally denotes firms targeting motivated consumers only, 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 the pricing is based on *actual* utilities.

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.

$$D(p) = \begin{cases} 1 & \text{if } p \leq \underline{u} = \underline{u} + \mu(1 - \alpha)\Delta \\ \mu & \text{if } p \in (\underline{u}, \tilde{u}] = (\underline{u} + \mu(1 - \alpha)\Delta, \underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta] \\ 0 & \text{if } p > \tilde{u} = \underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta \end{cases} \quad (\text{A.3})$$

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

$$\begin{aligned}\underline{\pi} &= \underline{u} - c = \underline{u} + \mu(1 - \alpha)\Delta - c \\ \tilde{\pi} &= \mu(\tilde{u} - c) = \mu(\underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta - c),\end{aligned}\tag{A.4}$$

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

$$\begin{aligned}\mu \cdot (\underline{u} + (\mu + (1 - \mu)\alpha)\Delta - c) &\geq \underline{u} + (\mu - \alpha\mu)\Delta - c \\ \Leftrightarrow c &\geq \underline{u} + \frac{\mu}{1 - \mu}\Delta(1 - 2\alpha - (1 - \alpha)\mu) \equiv \tilde{c}.\end{aligned}\tag{A.5}$$

Note that

$$\begin{aligned}E[u] \geq \tilde{c} &\Leftrightarrow \underline{u} + \mu\Delta \geq \underline{u} + \frac{\mu}{1 - \mu}\Delta(1 - 2\alpha - (1 - \alpha)\mu) \\ &\Leftrightarrow -(2 - \mu)\alpha \leq 0\end{aligned}\tag{A.6}$$

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

$$\begin{aligned}\text{(i)} \quad \partial\tilde{c}/\partial\alpha &= -(2 - \mu)\frac{\mu}{1 - \mu}\Delta < 0 \\ \text{(ii)} \quad \partial\tilde{c}/\partial\underline{u} &= 1 > 0 \\ \text{(iii)} \quad \partial\tilde{c}/\partial\Delta &= \underbrace{\frac{\mu}{1 - \mu}}_{>0} \cdot \underbrace{(1 - 2\alpha - (1 - \alpha)\mu)}_{\geq 0} \\ \text{(iv)} \quad \partial\tilde{c}/\partial\mu &= \underbrace{\frac{1}{(1 - \mu)^2}}_{>0} \Delta \underbrace{[1 - 2\alpha - (1 - \alpha)\mu(2 - \mu)]}_{\geq 0}\end{aligned}$$

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

$$\begin{aligned}\widetilde{CW} &= \mu \cdot (E[u] - \tilde{u}) = -\mu(1 - \mu)\alpha\Delta \leq 0 \\ \underline{CW} &= E[u] - \underline{u} = \mu\alpha\Delta \geq 0,\end{aligned}\tag{A.7}$$

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, regulation 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.¹⁵ 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.

¹⁵The 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 (\underline{p}) to changes in μ and Δ :

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

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 $c < E[u]$.

A.2 Cooling off

We next consider the case of a mandatory cooling off period. 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. Therefore, 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, $S(p)$:

$$D_c(p) = \begin{cases} 1 & \text{if } p \leq \delta \underline{u} = \delta(\underline{u} + \mu(1 - \alpha)\Delta) \\ \mu^2 & \text{if } p \in (\delta \underline{u}, \delta \tilde{u}] = (\delta(\underline{u} + \mu(1 - \alpha)\Delta), \delta(\underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta)] \\ 0 & \text{if } p > \delta \tilde{u} = \delta(\underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta) \end{cases} \quad (\text{A.8})$$

$$S(p) = \begin{cases} 0 & \text{if } p \leq \delta \underline{u} = \delta(\underline{u} + \mu(1 - \alpha)\Delta) \\ \mu(1 - \mu) & \text{if } p \in (\delta \underline{u}, \delta \tilde{u}] = (\delta(\underline{u} + \mu(1 - \alpha)\Delta), \delta(\underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta)] \\ 0 & \text{if } p > \delta \tilde{u} = \delta(\underline{u} + \alpha\Delta + \mu(1 - \alpha)\Delta) \end{cases} \quad (\text{A.9})$$

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

$$\begin{aligned} \tilde{\pi}_c &= \mu^2 \cdot \delta(\tilde{u} - c) = \delta\mu^2(\underline{u} + \alpha\Delta + (1 - \alpha)\mu\Delta - c) \\ \underline{\pi}_c &= \delta(\underline{u} - c) = \delta(\underline{u} + (1 - \alpha)\mu\Delta - c). \end{aligned} \quad (\text{A.10})$$

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

$$\begin{aligned}
& \tilde{\pi}_c \geq \underline{p}_c \\
\Leftrightarrow \delta\mu^2 \cdot (\underline{u} + (\mu + (1 - \mu)\alpha)\Delta - c) & \geq \delta(\underline{u} + (\mu - \alpha\mu)\Delta - c) \\
\Leftrightarrow c & \geq \underline{u} + \frac{\mu}{(1 - \mu)}\Delta \frac{(1 - (1 + \mu)\alpha - (1 - \alpha)\mu^2)}{1 + \mu} \quad (\text{A.11})
\end{aligned}$$

where

$$\begin{aligned}
& \tilde{c}_c > \tilde{c} \\
\Leftrightarrow \frac{(1 - (1 + \mu)\alpha - (1 - \alpha)\mu^2)}{1 + \mu} & > (1 - 2\alpha - (1 - \alpha)\mu) \\
\Leftrightarrow \alpha + (1 + 2\alpha)\mu & > 0, \quad (\text{A.12})
\end{aligned}$$

which holds generically. Similarly, we have that

$$\begin{aligned}
\tilde{c}_c \leq E[u] & \Leftrightarrow \underline{u} + \frac{\mu}{(1 - \mu)}\Delta \frac{(1 - (1 + \mu)\alpha - (1 - \alpha)\mu^2)}{1 + \mu} \leq \underline{u} + \mu\Delta \\
& \Leftrightarrow -(1 + \mu(1 - \mu))\alpha \leq 0 \quad (\text{A.13})
\end{aligned}$$

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

- (i) $\partial\tilde{c}_c/\partial\alpha = \frac{\mu}{1-\mu^2}\Delta(-\mu - (1 - \mu^2)) < 0$
- (ii) $\partial\tilde{c}_c/\partial\underline{u} = 1 > 0$
- (iii) $\partial\tilde{c}_c/\partial\Delta = \frac{\mu}{1-\mu^2}(1 - (1 + \mu)\alpha - (1 - \alpha)\mu^2)$
- (iv) $\partial\tilde{c}_c/\partial\mu = -\frac{\Delta}{(1-\mu^2)^2} \cdot [\mu(1 - \mu^2)(\alpha + (2 - \alpha)\mu) - (1 + \mu^2)((1 - \alpha)(1 - \mu^2) - \alpha\mu)]$

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 $c \in (E[u], \delta\tilde{u}]$.

Cost Range	$c \leq \tilde{c}$	$c \in (\tilde{c}, \tilde{c}_c]$	$c > \tilde{c}_c$
Δ Price	≈ 0 (by δ)	$\delta \underline{u} - \bar{u} < 0$	≈ 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 ($\alpha > 0$)

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 $\alpha > 0$.

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 $R(p)$:

$$D_r(p) = \begin{cases} 1 & \text{if } p \leq \underline{u} \\ \mu & \text{if } p \in (\underline{u}, \tilde{u}(\bar{s}|\underline{s})] = (\underline{u}, \alpha\underline{u} + (1-\alpha)\bar{u}] \\ \mu^2 & \text{if } p \in (\tilde{u}(\bar{s}|\underline{s}), \bar{u}] = (\alpha\underline{u} + (1-\alpha)\bar{u}, \bar{u}] \\ 0 & \text{if } p > \bar{u} \end{cases} \quad (\text{A.14})$$

$$R(p) = \begin{cases} 0 & \text{if } p \leq \underline{u} \\ (1-\mu) & \text{if } p \in (\underline{u}, \tilde{u}(\bar{s}|\underline{s})] = (\underline{u}, \alpha\underline{u} + (1-\alpha)\bar{u}] \\ \mu \cdot (1-\mu) & \text{if } p \in (\tilde{u}(\bar{s}|\underline{s}), \bar{u}] = (\alpha\underline{u} + (1-\alpha)\bar{u}, \bar{u}] \\ 0 & \text{if } p > \bar{u} \end{cases} \quad (\text{A.15})$$

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

$$\begin{aligned}
\bar{\pi}_r &= \mu^2 \cdot (\bar{u} - c) = \mu^2 \cdot (\underline{u} + \Delta - c) \\
\underline{\pi}_r &= \mu \cdot [\underline{u} + (1 - \alpha)\Delta - c] \\
\pi_r &= \underline{u} - c,
\end{aligned} \tag{A.16}$$

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

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

$$\begin{aligned}
\bar{\pi}_r \geq \pi_r &\iff \mu^2\underline{u} + \mu^2\Delta - \mu^2c \geq \underline{u} - c \\
&\iff c \geq \underline{u} - \frac{\mu^2}{(1 - \mu)(1 + \mu)}\Delta \equiv \tilde{c}_r^2
\end{aligned} \tag{A.18}$$

$$\begin{aligned}
\underline{\pi}_r \geq \pi_r &\iff \mu\underline{u} + \mu(1 - \alpha)\Delta - \mu c \geq \underline{u} - c \\
&\iff c \geq \underline{u} - \frac{\mu(1 - \alpha)}{(1 - \mu)}\Delta \equiv \tilde{c}_r^3
\end{aligned} \tag{A.19}$$

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

$$\begin{aligned}
\tilde{c}_r^1 > \tilde{c}_r^2 &\iff 1 - \mu - \alpha > -\frac{\mu^2}{1 + \mu} \\
&\iff \alpha < \frac{1}{1 + \mu} \\
&\iff \tilde{c}_r^2 > \tilde{c}_r^3.
\end{aligned} \tag{A.20}$$

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

$$\tilde{c} \geq \tilde{c}_r^3 \iff 1 \leq (3 - \mu)(1 - \alpha) \iff \alpha \leq \frac{2 - \mu}{3 - \mu} \tag{A.21}$$

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	$-\mu(1 - \alpha)\Delta < 0$	$(1 - \mu)(c - \underline{u}) - \mu\alpha\Delta$	$(1 - \mu)(\underline{u} - c) - (\alpha + (1 - \alpha)\mu)\Delta$
Δ Consumer Welfare	$(1 - \alpha)\mu\Delta > 0$	-	> 0 (from neg to pos)
Δ Total Welfare	-	$(1 - \mu)(c - \underline{u})$	$(1 - \mu)(E[\underline{u}] - c) > 0$ 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	$> 0 \iff 1 - 2\alpha - (1 - \alpha)\mu > 0$	$-\mu(1 - \mu)(\underline{u} + \alpha\Delta - c)$
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 Robustness

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, 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

$$dCW = 0, dTW = (1 - \mu)(c - \underline{u}). \quad (\text{A.22})$$

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.

Comparison between different cooling off policies While there are many instances in which only either a cooling off period, or a return policy is applicable, there may also be scenarios in which both policies would in principle be implementable. As such, it is 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 always 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, this is not necessarily the case.

Proposition A.1 (Cooling off, return policy & consumer welfare)

A mandatory cooling off period can lead to higher consumer welfare and/or total welfare than a return policy.

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.

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

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

$$\begin{aligned}
\tilde{c}_r > \tilde{c}_r^2 &\Leftrightarrow \frac{\mu(1 - \alpha - \mu[\alpha + (1 - \alpha)\mu])}{1 - \mu^2} + \frac{\mu^2}{1 - \mu^2} > 0 \\
&\Leftrightarrow \mu(1 - \alpha - \mu[\alpha + (1 - \alpha)\mu]) + \mu^2 > 0 \\
&\Leftrightarrow \underbrace{(1 - \alpha)}_{\geq 0} - \underbrace{[\alpha + (1 - \alpha)\mu]}_{< 1} \mu \geq -\mu, \tag{A.23}
\end{aligned}$$

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).

■

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. In particular, this may occur whenever the firm targets only motivated consumer in the baseline, i.e. whenever $c > \tilde{c}$. In this case, motivated consumers predicted utilities are fully restricted by the firm's pricing and consumer welfare is negative. While this can always be avoided by the introduction of a return policy, it may also be avoided with cooling off policies provided that the firm would switch its targeting decision to full market coverage, i.e. whenever $c \leq \tilde{c}_c$. Crucially, we establish that it is possible that full market coverage may occur provided that a cooling off policy is in place, while under a return policy the firm would still target only twice-motivated consumers, which directly implies that a cooling off period is preferable from a consumer welfare perspective.

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

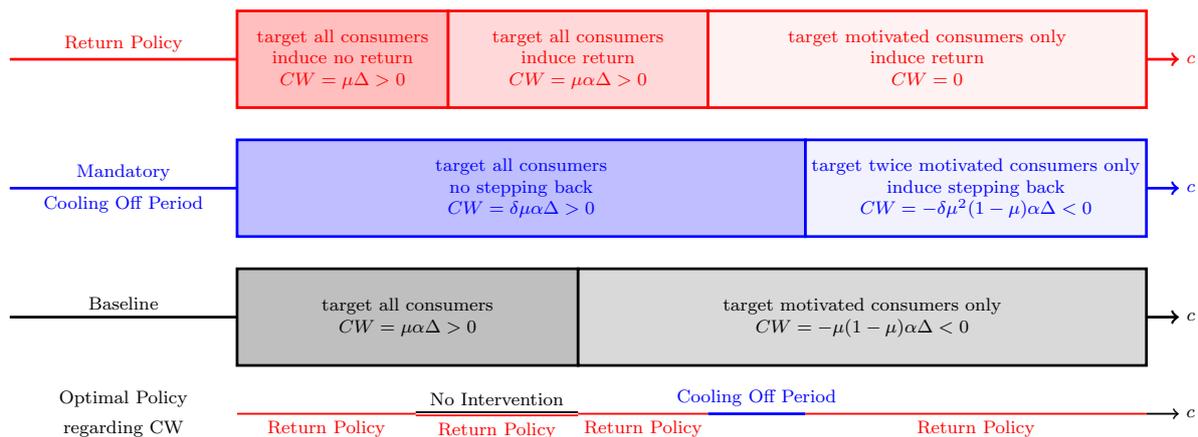


Figure 6: Relevant cost ranges across all policy regimes & Optimal policy

economy also achieves the optimal consumer welfare for low to intermediate c . In those cases, the absence of a policy intervention would arguably be preferred by a social planner.

This shows that cooling off periods may be preferable despite the delayed consumption even in the *best-case-scenario* regarding 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 a cooling off policy, which would preserve negative consumer welfare. This could be avoided entirely with a return policy. Given that the optimal policy depends on the fundamentals of the considered market, and that both policies may be optimal from a consumer welfare perspective, it is crucial to assess policies ex-post.

Costly Return The analysis so far considered the case where return is costless. However, this is unlikely to hold in reality: There are shipping costs of physically returning the good, as well as potential hassle costs of dealing with the associated paperwork and the opportunity costs of devoting time to the return process on the consumers' side. Moreover, the firm has to produce the good and ship it to the consumer, even though consumption does not take place as the good is returned. Both types of costs have a similar impact on the firm's profits and hence its targeting strategy. Consider first the case of anticipated return costs by consumers $r_c > 0$ which are incurred whenever a good is returned instead of consumed. Provided that return costs are not too large, the same three pricing strategies carry over from the setting without return costs. If return costs are extremely high, a consumer may always keep the good even if unmotivated – in that case, pricing by the firm would be based on predicted utilities and hence completely identical to the baseline model. If a consumer anticipates returning the good when unmotivated, the price at which she is willing to initially purchase the good will be lowered by $(1 - \mu)r_c$ – she

needs to be compensated for the anticipated costly return in the unmotivated state. The firm's profits are lowered, similar to the case where a returned product leads to costs for the firm. To model this, let $\eta \in [0, 1]$ be the fraction of the production cost c which the firm incurs whenever a product is returned. $\eta > 0$ for example captures the fact that the firm has been returned a good which was shipped out but not used and which can hence be sold later. Whenever a consumer returns the good, the firm incurs a cost ηc proportional to the initial production cost.

We show below that the main propositions carry over to the setting where the firm internalizes a fraction $\eta > 0$ of the production cost for returned, unused goods, which essentially comprises the case of return costs which consumers incur. In particular, we show that the same targeting strategies hold in equilibrium, albeit with cost thresholds adjusted to reflect the additional costs and lowered profitability.

As in the case of non-internalized return costs, there are three potential targeting strategies. Pricing is as in the case without the return cost, but the firm's profits are lower due to the internalized cost. The profits associated with each strategy are hence given by

(M,R) target only motivated, induce return unless again motivated

$$\pi_{M,R} = \mu(\mu(\bar{u} - c) - (1 - \mu)\eta c) = \mu(\mu(\underline{u} + \Delta) - (\mu + (1 - \mu)\eta)c) \quad (\text{A.24})$$

(A,R) target all consumers, but induce return unless motivated

$$\pi_{A,R} = \mu(\underline{u} + (1 - \alpha)\Delta - c) - (1 - \mu)\eta c = \mu(\underline{u} + (1 - \alpha)\Delta) - (\mu + (1 - \mu)\eta)c \quad (\text{A.25})$$

(A,N) target all consumers, induce no return

$$\pi_{A,N} = \underline{u} - c \quad (\text{A.26})$$

To prefer (M,R) to (A,R), we require

$$\begin{aligned} \mu \cdot [\mu(\underline{u} + \Delta) - (\mu + \eta(1 - \mu))c] &\geq \mu(\underline{u} + (1 - \alpha)\Delta - c) - (1 - \mu)\eta c \\ c &\geq \frac{\mu}{\mu + (1 - \mu)\eta} \underline{u} + \Delta \frac{\mu(1 - \mu - \alpha)}{\mu(1 - \mu) + (1 - \mu)^2 \eta} \equiv \hat{c}_r^1 \end{aligned} \quad (\text{A.27})$$

Note that for $\eta = 0$ we have $\hat{c}_r^1 = \underline{u} + \Delta \frac{1 - \mu - \alpha}{1 - \mu} = \tilde{c}_r^1$. To prefer (M,R) to (A,N), we require

$$\begin{aligned} \mu \cdot [\mu(\underline{u} + \Delta) - (\mu + \eta(1 - \mu))c] &\geq \underline{u} - c \\ c &\geq \frac{1 - \mu^2}{1 - \mu^2 - \mu(1 - \mu)\eta} \underline{u} - \Delta \frac{\mu^2}{1 - \mu^2 - \mu(1 - \mu)\eta} \equiv \hat{c}_r^2 \end{aligned} \quad (\text{A.28})$$

Note that for $\eta = 0$ we have $\hat{c}_r^2 = \underline{u} - \Delta \frac{\mu^2}{1-\mu^2} = \tilde{c}_r^2$. Finally, to prefer (A,R) to (A,N), we require

$$\begin{aligned} \mu(\underline{u} + (1-\alpha)\Delta - c) - (1-\mu)\eta c &\geq \underline{u} - c \\ c &\geq \underline{u} \frac{1-\mu}{1-(\mu+(1-\mu)\eta)} - \Delta \frac{\mu(1-\alpha)}{1-(\mu+(1-\mu)\eta)} \end{aligned} \quad (\text{A.29})$$

Note that for $\eta = 0$ we have $\hat{c}_r^3 = \underline{u} - \Delta \frac{\mu(1-\alpha)}{1-\mu} = \tilde{c}_r^3$.

As before, it can still occur that either $\hat{c}_r^1 \geq \hat{c}_r^2 \geq \hat{c}_r^3$ or $\hat{c}_r^1 \leq \hat{c}_r^2 \leq \hat{c}_r^3$. To see this, define $\Gamma \equiv \mu + (1-\mu)\eta$ where $\mu \in (0,1)$ and $\eta \in [0,1]$ ensures $\Gamma \in (0,1)$. We then have that

$$\begin{aligned} \hat{c}_r^1 \geq \hat{c}_r^2 &\Leftrightarrow \underline{u} \left[\frac{\mu}{\Gamma} - \frac{1-\mu^2}{1-\mu\Gamma} \right] + \Delta \left[\frac{\mu(1-\mu-\alpha)}{(1-\mu)\Gamma} + \frac{\mu^2}{1-\mu\Gamma} \right] \geq 0 \\ &\Leftrightarrow \underline{u}(1-\mu)[\mu-\Gamma] + \Delta\mu[1+\alpha\mu\Gamma-\alpha-\mu], \end{aligned} \quad (\text{A.30})$$

$$\begin{aligned} \hat{c}_r^1 \geq \hat{c}_r^3 &\Leftrightarrow \underline{u} \left[\frac{\mu}{\Gamma} - \frac{1-\mu}{1-\Gamma} \right] + \Delta \left[\frac{\mu(1-\mu-\alpha)}{(1-\mu)\Gamma} + \frac{\mu(1-\alpha)}{1-\Gamma} \right] \geq 0 \\ &\Leftrightarrow \underline{u}(1-\mu)[\mu-\Gamma] + \Delta\mu[1+\alpha\mu\Gamma-\alpha-\mu], \end{aligned} \quad (\text{A.31})$$

$$\begin{aligned} \hat{c}_r^2 \geq \hat{c}_r^3 &\Leftrightarrow \underline{u} \left[\frac{1-\mu^2}{1-\mu\Gamma} - \frac{1-\mu}{1-\Gamma} \right] - \Delta \left[\frac{\mu^2}{1-\mu\Gamma} - \frac{\mu(1-\alpha)}{1-\Gamma} \right] \geq 0 \\ &\Leftrightarrow \underline{u}(1-\mu)[\mu-\Gamma] + \mu\Delta[1+\alpha\mu\Gamma-\alpha-\mu]. \end{aligned} \quad (\text{A.32})$$

Hence the sign of

$$\begin{aligned} &\underline{u}(1-\mu)[\mu-\Gamma] + \mu\Delta[1+\alpha\mu\Gamma-\alpha-\mu] \\ &= -\underline{u}(1-\mu)^2\eta + \mu(1-\mu)\Delta[1-\alpha(1+(1-\eta)\mu)] \end{aligned} \quad (\text{A.33})$$

which is the same as the sign of

$$\Psi \equiv \mu(1-\alpha(1+(1-\eta)\mu))\Delta - (1-\mu)\eta\underline{u} \quad (\text{A.34})$$

determines the order of the thresholds. Note that for $\eta = 0$, we have $\text{sgn}(\Psi) = \text{sgn}(1-\alpha(1+\mu))$. The firm's pricing decision hence carries over from the baseline, albeit accounting for the internalized cost η per produced but returned good. Noting that $\pi_{M,R} \geq 0 \iff c \leq \frac{\mu}{\mu+(1-\mu)\eta}\bar{u}$, the following proposition, which is analogous to [Proposition 3](#), follows immediately.

Proposition A.2 (Firm's pricing given return policy) *There exist thresholds \underline{c}_r and \hat{c}_r which determine the firm's pricing decision in presence of a return policy.*

- (i) *If $c \leq \underline{c}_r$, the firm initially targets all consumers and all consumers keep the good. The firm charges $\underline{p}_r = \underline{u}$ to reap profits $\underline{u} - c$, while consumer surplus is $CW = \mu\Delta$.*

(ii) If $c \in (\underline{c}_r, \hat{c}_r]$, the firm initially targets all consumers, who return the good unless they are motivated in the second period. The firm charges \underline{p}_r to reap profits $\pi_{A,R} = \mu(\underline{u} + (1 - \alpha)\Delta) - (\mu + (1 - \mu)\eta)c$, while consumer surplus is $\underline{CW}_r = \mu\alpha\Delta$.

(iii) If $c \in \left(\hat{c}_r, \frac{\mu}{\mu + (1 - \mu)\eta}\bar{u}\right]$, the firm initially targets only motivated consumers, who return the good unless they remain motivated in the second period. The firm charges \bar{p}_r to reap profits $\pi_{M,R} = \mu(\mu(\underline{u} + \Delta) - (\mu + (1 - \mu)\eta)c)$, while consumer surplus is $\overline{CW}_r = 0$.

$(\underline{c}_r, \hat{c}_r]$ is nonempty if and only if $\text{sgn}(\Psi) = 1$.

Overall, the targeting behavior is hence similar to that without return costs; however, the profitability of strategies which induce on-path return is lowered, which increases the range of marginal cost such that the full market is covered. Note that we never observe targeting of all with return on-path if $\text{sgn}(\Psi) = -1$. This is intuitive once we inspect Ψ more closely.

(i) $-(1 - \mu)\eta\underline{u}$ is unambiguously weakly negative (strictly for $\eta > 0$)

(ii) $\text{sgn}(\Psi) = 1$ thus necessitates that $1 - \alpha(1 + (1 - \eta)\mu) \geq 0$

– the only way to get $1 - \alpha(1 + (1 - \eta)\mu) < 0$ in turn requires

(a) a large α

(b) a relatively low η

(c) a relatively large μ

So if the bias is large, the internalized cost of wasteful production η low, and the fraction of motivated consumers μ relatively large, there never will be targeting of all consumers initially, but reducing return on path. This is because (i) a high fraction of motivated consumers makes the fully exclusive targeting strategy relatively more profitable, (ii) high internalized return costs make the exclusive targeting strategy relatively more profitable as less return is induced on path, and (iii) a high degree of the bias α lowers the profitability of the full targeting strategy with return – the price which can be charged is equal to $\tilde{u}(\bar{s}|\underline{s}) = \underline{u} + (1 - \alpha)\Delta$ which decreases in α ; the larger the bias, the lower the predicted high-motivation utility for unmotivated people.

Comparison of policies Regarding the comparison of policies with positive return costs, first note that we can rewrite

$$\begin{aligned}\hat{c}_r^1 &= \frac{\mu}{\mu(1 - \mu) + (1 - \mu)^2\eta} ((1 - \mu)(\underline{u} + \Delta) - \alpha\Delta) \\ \hat{c}_r^2 &= \frac{1}{1 - \mu^2 - \mu(1 - \mu)\eta} (\underline{u} - \mu^2(\underline{u} + \Delta))\end{aligned}\tag{A.35}$$

Note that if $\Psi \geq 0$, \hat{c}_r^1 is the relevant cutoff. But this cutoff is positive iff $(1 - \mu)(\underline{u} + \Delta) - \alpha\Delta > 0$, in which case it decreases in η . So whenever $\Psi \geq 0$ and $\hat{c}_r^1 = \hat{c}_r > 0$, the fact that $\hat{c}_r|_{\eta=0} < \tilde{c}_c$ suffices to conclude that cooling off is preferred to a return policy for a positive mass of marginal cost. Moreover, if it is negative the point is moot as exclusive targeting always prevails – while the cutoff is decreasing in η in this case, it always remains negative. So in case we have $\Psi \geq 0$, we will always have the same property as previously established.

If $\Psi < 0$, $\hat{c}_r^2 = \hat{c}_r$ is the relevant cutoff, which for $\underline{u} - \mu^2(\underline{u} + \Delta) > 0$ is positive and increasing in η . So it may be the case that we then have $\hat{c}_r^2 > \tilde{c}_c$ for some values of η (recall that for $\eta = 0$, we have $\hat{c}_r^2|_{\eta=0} < \tilde{c}_c$). We can summarize that a return policy may do worse than a cooling off period with regards to consumer welfare even if return costs are internalized. However, for large η , i.e. if the internalized fraction of the production cost is sizeable, there may not exist a cost range such that this is the case. The reasoning behind this is simple: As the internalization occurs only in the scenario where a return policy is in place, it may make exclusive targeting sufficiently unattractive for the firm such that it always prefers less exclusive targeting behavior to avoid the internalization of return costs. Aside from the total welfare loss which is associated with a return policy in the presence of return costs, i.e. $\eta \cdot c > 0$ for each returned good, the result that a cooling off period may be better for consumer welfare than a return policy, despite the fact that the latter avoids negative expected consumer utility, carries over.